Endogenous Vertical Differentiation, Variety, and the Unequal Gains from Trade*

Gunnar Heins†

University of Florida

January 24, 2019

Abstract

How unequal are the gains from trade? In this paper, I develop an empirical framework to quantify the consequences of international trade on welfare of consumers across the income distribution. I first document large vertical differences between domestic and imported varieties across EU member countries with imports in richer countries being consistently cheaper and less sophisticated than varieties offered by home firms and the opposite being true in poorer economies. I then employ a structural model in which consumer demand for higher-quality goods is non-homothetic and firms endogenously choose the quality of their products. The model delivers household-specific price indexes whose parameters can be estimated from the data and incorporates a flexible production side which allows consistent predictions of how countries vertically specialize in response to globalization. It can be brought to the data using random coefficients demand estimation techniques and I infer demand and supply parameters for more than 3,000 highly disaggregated industries. I find the gains from trade to be moderately unequal except in small, less wealthy economies. Further, I show that not accounting for vertical differentiation would overstate the impact of trade on cost-of-living inequality by up to 52%.

---

[*] This paper was the first part of my dissertation at the University of Chicago and many thanks go therefore to my advisors Ralph Ossa, Ali Hortaçsu, Felix Tintelnot, and Jonathan Dingel. I also would like to thank Kerem Cosar, Cristian Dagnino, Jessie Handbury, Amit Khandelwal, Ignacia Mercadal, Eduardo Morales, Derek Neal, Ippei Shibata, as well as seminar participants at Purdue, Rice, Carlos III de Madrid, Tufts University, University of Oslo, BI Norwegian Business School, Einaudi Institute for Economics and Finance, University of Florida, Analysis Group, Cornerstone, as well as the international trade and industrial organization working groups at the University of Chicago for comments and suggestions.

[†] University of Florida, Gainesville, USA; Email: gheins@ufl.edu

1 Introduction

The last two decades have seen an unprecedented increase in exports from developing countries. China's exports alone have increased twentyfold over the last twenty years, and account now for roughly 13% of worldwide trade flows. Given its sheer magnitude, this trend has not only reignited the discussion on whether or not countries benefit from free trade in general, but also the question if trade is an important source of inequality.

There are several convincing arguments why trade matters for inequality: Labor from developing countries for example is arguably a more direct competition to blue collar workers in developed economies and may therefore amplify wage differentials between more and less educated workers through offshoring (see e.g. Hummels et al. (2014)) Also import competition from low-wage countries has been shown to be particularly harmful for less skilled and mobile workers (see e.g. Autor et al. (2014)). As evident for instance from the recent U.S. elections and primaries, these trends appear to have resulted in the growing widespread perception that trade increases inequality on both ends of the political spectrum.

The majority of these arguments and perceptions however are based on income and wages and almost universally ignore that trade also affects prices, the availability of goods and the cost of living. There are however good reasons why trade would also affect the cost-of-living of poor and rich households differently: First, consumption baskets are known to differ systematically with income with lower-income households frequently purchasing cheaper, less sophisticated versions of products relative to wealthier households. Second, countries vary in terms of how good they are at producing those consumption baskets: Germany is e.g. commonly referred to as being good at producing higher-quality cars while low-income countries tend to produce clothes or food more cheaply. As a consequence, trade will affect consumer groups asymmetrically if imports and domestic goods are differently relevant to households across the income distribution.

In this paper, I develop a theoretical and empirical model to quantify this channel. I first present suggestive evidence that domestic and imported varieties of goods indeed systematically differ: First, in line with previous research, there appear to be large unobserved differences between countries in terms of their products' quality : Richer economies export the same type of goods at up to 3 times higher unit values without noticeably lower market shares, suggesting that these varieties have certain characteristics that are superior to others and are preferred by consumers. Second, while imports in richer EU economies are consistently cheaper and of lower quality than varieties offered by home firms, the opposite is true in poorer Eastern European countries. This suggests not only that imported and domestic varieties differ but also that the size and direction of these differences is country-specific and correlated with income levels.

In order to quantify how unequal the gains from trade are across countries, I develop and employ a structural model which is necessary for several reasons: First, in order to measure and aggregate individual welfare across the income distribution and thousands of products, I need to quantify how and how differently households value each variety, which is infeasible without at least some basic structure on the demand side. My theoretical framework delivers household-specific price indexes whose parameters can be estimated from the data and, importantly, can then be used for counterfactuals and the comparison of a scenario with and without trade.

Second, the stylized facts on EU countries are based on unit values which are known to be a potentially misleading measure of product quality. While the observation that low unit values do not correspond to higher market shares in the data is suggestive of vertical differentiation, it is nevertheless important to rigorously disentangle the impact of quality on unit values from that of wages and productivity. My model does so by relying on the idea of Khandelwal (2010) and jointly uses information on unit values and market shares to identify quality from the data.

Third, vertical differentiation is plausibly endogenous. The measured price and quality differentials between foreign and domestic goods are equilibrium outcomes and are realistically affected by trade. In particular, I show that if countries adjust their varieties in response to foreign competition, for which there is ample support in the data (see e.g. Amiti and Khandelwal (2013), Bloom, Draca, and Van Reenen (2016), or Dingel (2017)), the empirical results will be biased when not adequately accounting for supply side adjustments. The model therefore incorporates a flexible production side which allows consistent estimation of how countries vertically specialize in response to globalization.

Motivated by the stylized facts present in the data, I develop a multi-sector model of international trade with three key features: First, I allow consumer demand to be potentially non-homothetic, i.e. demand for higher-quality varieties to vary with income. Second, countries differ in their available production technologies and endogenously choose the quality of their products. Third, in order to be able to make statements on how unequal the effects of international trade are, each country is populated by a distribution of heterogeneous households which will earn different incomes.

I estimate the model on a dataset of trade flows within and into the European Union as well as matched production data in more than 3,000 disaggregated product categories. A major estimation challenge on the demand-side is that household-level consumption data is rarely available, and if so, it is typically limited to a single country or a small set of products. Relying on such datasets would therefore make an estimation of the unequal gains from trade for a wide range of countries infeasible and also limit the ability to empirically determine to which degree exporters respond to variation in the income distribution in each destination.

My main methodological contribution on the demand side is to overcome this limitation by developing a demand framework which can be brought to the data using only commonly available data on market-level expenditure shares, unit values and the income distribution in various countries. In particular, using log-logit preferences, my model delivers a tractable estimation equation which resembles that of discrete choice random coefficients models of consumer demand in the style of Berry, Levinsohn, and Pakes (1995, BLP), which have become increasingly popular in the structural industrial organization literature. This allows me to greatly benefit from recent computational methods in models with consumer heterogeneity and, while still computationally demanding, estimate demand for several thousand distinct product categories reliably.

Once the demand side parameters are known, I back out the production technology of each producer and product that rationalizes the price and quality choices in the data. The empirical framework provides estimates on the elasticity of the (log) marginal cost with respect to quality, which allows me to quantify how firms vertically differentiate in response to market conditions and importantly, foreign competition. The model delivers an intuitive condition for the direction in which firms adjust the quality of their products in response to trade based on (1) their own technology, (2) foreign technology, and (3) the extent to which demand is non-homothetic.

I show that in the presence of consumer heterogeneity, the impact of trade liberalization on vertical differentiation can be captured by a sufficient statistic, the effective price elasticity faced by each producer. This elasticity is an equilibrium outcome and will be affected by trade if domestic and foreign varieties are differently attractive across consumer groups. In the empirical application, I find that trade predominantly results in entry of cheaper, less sophisticated varieties, which makes domestic EU companies disproportionately less attractive to poorer households and lowers the average price elasticity of their remaining consumers. In the model, this reduction in the effective aversion to higher prices leads firms to quality upgrade and I show that this channel is absent in representative household frameworks.

My parameter estimates imply that richer countries have a comparative advantage in producing higher-quality varieties for the majority of products. Specifically, the elasticity of (log) unit costs with respect to quality is lower in richer economies and results in those countries specializing in more sophisticated and expensive varieties of the same type of goods. I also find that demand for higher-quality varieties is disproportionately increasing in income and across the income distribution: In 78% of product categories, an increase in income leads households to put a greater weight on quality relative to the price in their consumption choice.

Based on these estimates, I simulate a counterfactual move to autarky to quantify the gains from trade for households in each country and across the income distribution. I first show that when holding the supply side fixed, the disproportionate inflow of cheaper, less sophisticated varieties results in poorer households benefiting more from trade than richer ones do in each EU member country. Depending on the specification, trade results in an on average 1 to 2.5 percentage points stronger decline in the cost of living for a household at the 20th percentile of the income distribution relative to one at the 80th percentile.

I also document considerable heterogeneity in terms of how unequal the gains from trade are in each country. While poorer households gain almost 3 percentage points more in the richest EU economies, the gains are almost equally distributed in the poorer member states Romania or Lithuania. Since richer countries tend to have a comparative advantage in producing higher quality-goods, but at a high price, poor households in these countries particularly benefit from access to cheaper foreign varieties through trade.

I find that allowing for endogenous quality specialization on the supply side has a significant impact on the results. The estimates imply that the average domestic producer will want to quality downgrade when moving to autarky in order to cover demand from poorer households that was previously met by cheaper foreign firms. These adjustments mitigate how unequal the gains from trade are: When vertical product characteristics are held fixed, the gap between rich and poor consumers averages 2.5 percentage points and can take larger values of 4 to 7 percentage points in rich economies. In comparison, when accounting for quality and price adjustments this gap shrinks to about 1.2 percentage point on average. In total, this translates into the gains from trade being 51.8% less unequal when taking vertical differentiation into account.

Finally, the gap in the welfare gains between poor and rich decreases further and actually reverses on average when accounting for entry of less productive firms in response to trade barriers. Given the inverse relationship between quality and productivity, entry of less productive firms lowers the average quality of domestic firms, thereby providing an additional channel which disproportionately benefits poor consumers. I find that this mechanism lowers the cost-of-living of rich households by 1.8 to 2.7 percentage points more than that of poorer ones. While trade continues to be 4-6 percentage points more pro-poor in richer economies compared to poorer ones, this channel also generates a sizable correlation between country size and inequality in the gains from trade: In particular, I find that trade in small and less wealthy countries magnifies nominal income inequality by almost 4 percentage points but has close to balanced consequences in larger and richer EU economies such as France, Germany, or the UK.

The relationship between trade and inequality is one of the classic topics of interest in international economics and has sparked an extensive theoretical and empirical literature over many decades. The distributional consequences of trade through the consumption channel however have only received sporadic attention in the empirical literature, especially in comparison to a large literature on the impact on earnings of workers with different skills, endowments, and education (see e.g. Goldberg (2015) for an extensive literature survey). Notable exceptions are for example Porto (2006) or Broda and Romalis (2009) who provided early studies on real consumption inequality in the context of international economics.

More recently, Fajgelbaum and Khandelwal (2016) measure how unequally trade affects consumer welfare based on the almost-ideal demand system (AIDS, Deaton and Muellbauer (1980)). The main advantage of using a log-logit demand system here is that it naturally and explicitly allows me to introduce product characteristics such as a product's quality. This enables me to quantify the location of each producer in characteristics-space and how it will change with trade and foreign competition. In Fajgelbaum and Khandelwal's (2016) Armington model, the supply side is fixed in the sense that countries exogenously specialize in producing goods with a certain income elasticity. In contrast, the income elasticity of varieties in my setup is an equilibrium outcome and itself changes with international trade, providing an additional, quantitatively relevant channel through which trade affects inequality and I believe this is the first paper that measures its aggregate quantitative importance.[1]

---

[1] Additionally, while admittedly having disadvantages in other dimensions, it is also slightly more straightforward

Broda and Romalis (2009) and Faber (2014) use detailed individual-level shopping data to study the heterogeneous impact of Chinese and U.S. imports, respectively, on the cost of living of differently rich consumers. While such datasets provide extremely rich information on consumer behavior and the extent to which demand is non-homothetic, they are generally not available for a broad set of countries and cover only a subsection of goods. The framework of this paper has the benefit that it can be applied to standard international trade datasets which makes it applicable to a wide range of countries and particularly well-suited for cross-country comparisons.

The paper also relates to recent work by Faber and Fally (2017), who use a detailed scanner dataset for the U.S. to quantify the unequal gains from trade. In particular, the paper shows that when taking into account the firm-size distribution and the observation that larger, more productive firms tend to sort into producing more income-elastic varieties, trade disproportionately benefits wealthier consumers. While this paper focuses more on within-firm differentiation, I show that Faber and Fally's (2017) findings also hold in the current setting and hence appear to apply more broadly to the manufacturing sector and for a wide range of countries. In addition, I document that the impact of trade on cost-of-living inequality varies quite strongly with country characteristics with trade being substantially more pro-rich in poorer as well as in smaller countries. To my knowledge, this paper is the first one to quantify the relationship between country size and the impact of trade on cost-of-living inequality.

This paper also contributes to a large and growing literature on vertical differentiation (Schott (2004), Hummels and Klenow (2005), Hallak (2006), Khandelwal (2010), Hallak and Schott (2011), Feenstra and Romalis (2014), Dingel (2017)) and quality upgrading (Verhoogen (2008), Amiti and Khandelwal (2013), Bloom et al. (2016)) in international trade. In contrast to those papers, I quantify and endogenize quality for a wide range of products and countries in a framework with consumer heterogeneity and highlight the importance of the effective price elasticity faced by each producer as sufficient statistic for the impact of trade liberalization on vertical differentiation. My estimates allow me to predict the extent and direction in which firms adjust the quality of their products in response to trade and suggest substantial heterogeneity across both countries and products.

Finally, the paper also relates to the literature on demand estimation and non-homothetic consumer demand in international economics and related fields (see e.g. Fieler (2011), Fajgelbaum, Grossman, and Helpman (2011), Hummels and Lugovskyy (2009), Handbury (2013)). My paper contributes to this literature by providing a novel identification strategy and estimates on the extent to which demand is non-homothetic for thousands of goods and a range of countries. It is also a large-scale application of the BLP demand system and delivers estimates on price elasticities and the importance of unobserved heterogeneity for thousands of product categories and more than

---

to introduce an extensive margin or prohibitively high trade cost in the log-logit demand system, in part due to the possibility of AIDS generating negative expenditure shares in empirical applications. More generally, characteristics-space approaches for example also allow projecting the welfare implications of new varieties (see Petrin (2002)). Finally, while AIDS relates to translog utility, the discrete-choice model used here is a generalization of CES preferences. The aggregate welfare results and parameter estimates of this paper are therefore readily comparable to those found e.g. in Broda and Weinstein (2006) or Ossa (2015).

hundred thousand varieties in contrast to studies that focus on a particular industry (see e.g. Nevo (2000), Petrin (2002) or Cosar et al. (2015)). Further, the addition of consumer heterogeneity generates household-specific elasticities of substitution for each product, which differ by income group as well as country and complement previous estimates obtained in representative household frameworks (see e.g. Broda and Weinstein (2006), Soderbery (2015), or Ossa (2015)). Given the increasing popularity of random coefficients models of consumer demand, I believe these estimates will prove useful and relevant beyond the context of the present analysis.

The paper is structured as follows: The next section presents several stylized facts in the data which motivate the analysis and setup of the model. Section 3 covers the theoretical model along with its key predictions. In section 4, I lay out my estimation strategy, describe the data, and discuss identification of the model. Section 5 presents the parameter estimates, model fit, and the counterfactual implications. Finally, sections 6 and 7 assess the robustness of the results and conclude.

2 Motivation

This section documents several empirical facts on trade flows involving EU member states, which motivate the analysis, the setup of my theoretical model, as well as the resulting empirical framework.[2] In particular, the data suggests (1) the presence of systematic vertical differences between domestic and imported varieties of products, (2) large unobserved quality differences between exporters, and (3) a non-homothetic consumer demand.

In this part, I use data on bilateral trade flows and matched production data for the European Union in highly disaggregated 8-digit product categories between 1989 and 2013. The data covers 27 importers and the universe of exporters to the European Union during the given time period. The EU's statistical agency, Eurostat, categorizes products into roughly 10,000 8-digit categories and within these categories, products are homogeneous in the sense of being similar in terms of their product characteristics. I will describe the data along with summary statistics in more detail in section 4.

Figure 1 documents the first set of empirical regularities in the European trade data: First, richer countries in terms of GDP per capita import goods at systematically higher prices. Panel a) summarizes the coefficients β_j^M from a regression of import prices on importer dummies and product fixed effects,

$$\log(p_{jk}^M) = \beta_j^M \mathbb{I}\{\text{Importer} = j\} + \beta_k^M \mathbb{I}\{\text{Product} = k\} + \varepsilon_{jk}, \tag{1}$$

where p_{jk}^M denotes the average import price in country j in a product category k, weighted by trade volume. Each β_j^M can be interpreted as percentage deviation in import prices compared to a reference country, for which I choose France. In panel a), I restrict the sample to exporters

---

[2] Some of these facts have also been documented in other datasets, see in particular Schott (2004), Hummels and Klenow (2005), and Hallak (2006).

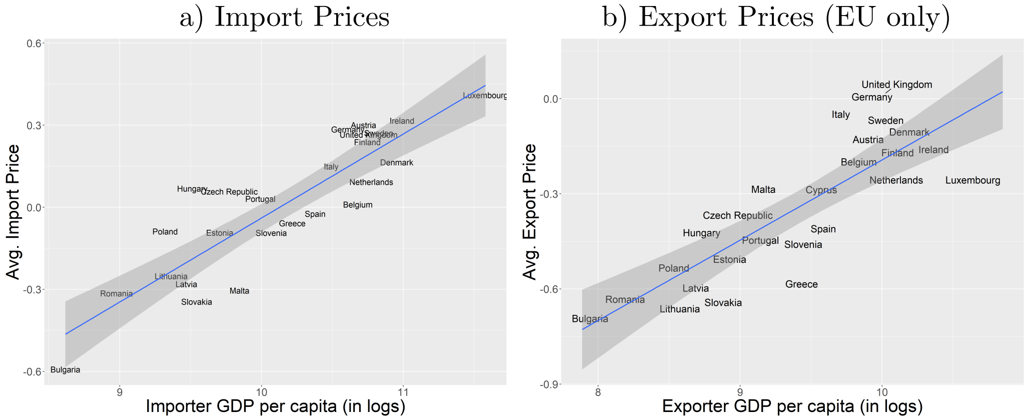
a) Import Prices

b) Export Prices (EU only)

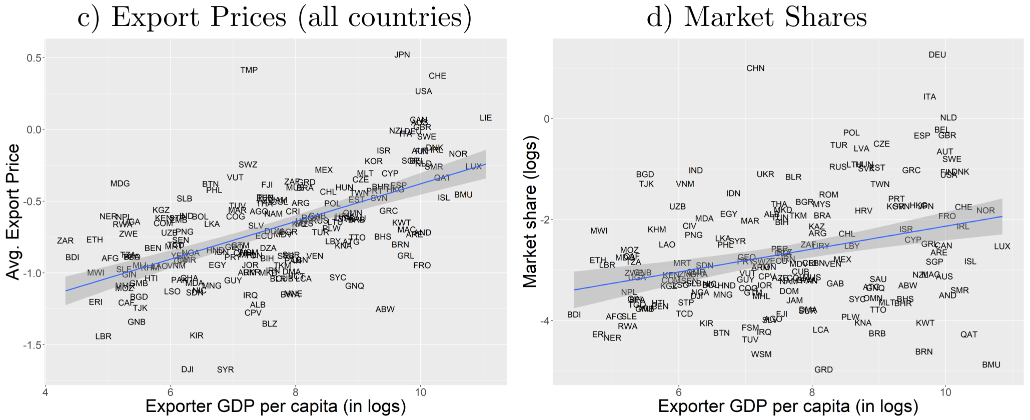
c) Export Prices (all countries)

d) Market Shares

Figure 1: **Import and Export Prices**: Panel a) plots the coefficients β_j^M from regression (1). Each value can be interpreted as the percentage deviation in import prices relative to France. Panels b) and c) plot the coefficients β_j^X from regression (2), which reflect the percentage deviation in export prices relative to France. Plot b) reports the estimates for EU countries only while all estimates are shown in panel c). Panel d) summarizes the estimates of β_j^S from regression (3). All data is for the year 2007.

from outside Europe to minimize the impact of location within Europe and to ensure that the relationship is not driven by a country's proximity to other rich economies.[3]

The plot in Panel a) demonstrates a strong positive relationship between a country's GDP per capita and the average price of its imports: The average import price is for example about 30 - 40% higher in the richer economies Ireland and Luxembourg compared to France while Bulgaria imports at about 60% lower prices. The slope of the fitted line in Panel a) equals 0.30, which implies that a 10% increase in GDP per capita correlates with a 3% increase in the average import price.

Second, Panels b) and c) show that richer countries not only import but also export goods at higher prices. Analogous to Panel a), these figures plot the estimates of the coefficients β_j^X in the

---

[3] As an example, the average price of exports to Poland would be higher when using the full sample since it is geographically close to Germany, which sells high-priced varieties. Even when using the full sample however, there is still a clear positive relationship between a import prices and GDP per capita.

regression

$$\log(p^X_{jk}) = \beta^X_j \mathbb{I}\{\text{Exporter} = j\} + \beta^X_k \mathbb{I}\{\text{Product} = k\} + \varepsilon_{jk}, \tag{2}$$

where p^X_{jk} denotes the average export price of country j weighted by trade volume when exporting product k. In Panel b), I only report the estimates for the 27 EU countries while Panel c) plots the results for all countries. I again exclude France in this regression so that the β^X_j's reflect the average percentage deviations of each exporter's prices from those of French firms. As evident from both panels, there is a similarly strong positive correlation between exporter GDP per capita (in logs) and average export prices (in logs). The highest-priced exports originate in rich countries such as Japan, Switzerland, and the United States and their exports cost roughly 3 times as much as the lowest-priced varieties. The slope of the fitted line in panel c) is 0.21 implying that a 10% increase in GDP per capita correlates with a 2.1% increase in the average export price.[4]

Third, these higher export prices do not translate into lower market shares. As shown in Panel d), evaluating regression (2) with market shares as dependent variable instead of export prices, i.e.

$$\log(ms_{jkm}) = \beta^S_j \mathbb{I}\{\text{Exporter} = j\} + \beta^S_k \mathbb{I}\{\text{Product} = k\} + \varepsilon_{jkm}, \tag{3}$$

reveals no negative relationship between market shares ms_{jkm} of exporter j in country m and j's GDP per capita. In fact, richer countries have on average even slightly higher market shares, despite selling at significantly high unit values.

I interpret the above empirical patterns as evidence for two familiar regularities in the data: First, Panels b) to d) suggest the presence of sizable unobserved quality differences between exporters, even within narrowly defined product categories: The fact that expensive varieties have high market shares implies that they must have certain characteristics which dominate others and are preferred by consumers. This vertical differentiation is also systematically correlated with income: Richer economies tend to produce higher-quality varieties than poorer countries do.

Second, Panel a) suggests that richer economies also import higher-quality varieties. This observation speaks in favor of demand being to some extent non-homothetic in the sense that richer households have a stronger preference for higher-quality versions of imported goods or are more able to afford them.

Figure 2 documents that there is also substantial variation in the extent to which the prices of domestic and imported goods differ across countries: Both figures plot the percentage price difference between domestic and imported varieties for the median product in each country. The left figure compares the price of domestic varieties to that of all foreign imports, while the right one compares it to that of Asian imports only, which account for the majority of growth in imports to Europe over the last decades.

As shown in the left panel of Figure 2, domestic goods are about 50% cheaper than imported varieties in the poorer EU economies but up to 25% more expensive in richer economies. The fitted

---

[4] This number is somewhat larger than the estimate found by Schott (2004) for the U.S. of 0.13. For the largest EU economies however, the numbers are significantly closer with 0.16 in France, 0.16 in Germany and 0.18 in the UK.

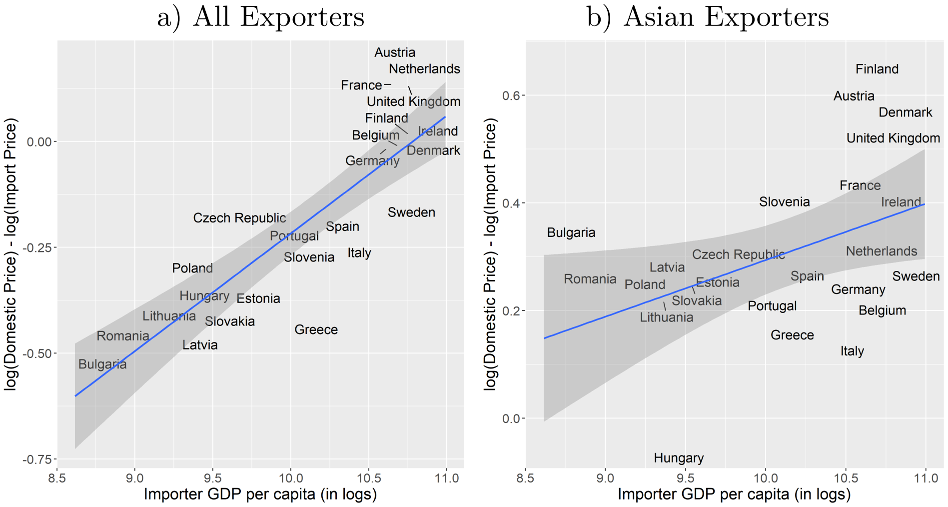

a) All Exporters

b) Asian Exporters

Figure 2: **Domestic and Import Prices**: Panels a) and b) plot the median percentage difference between domestic and import prices for each EU country. Plot a) uses the full sample while plot b) plots the difference in prices between domestic goods and Asian exports. All data is for the year 2007.

line in Panel a) has a slope of 0.28, implying that a 10% increase in GDP per capita translates into a 2.8% larger gap between domestic and import prices. When looking only at Asian imports, all but one country's produced varieties are more costly than its imports: For the median product, domestic varieties are between 10 and 65% more costly relative to imports.

Taken together, Figures 1 and 2 suggest the presence of significant differences between the quality and characteristics of domestic and imported varieties. Even though shipping is costly, foreign goods are frequently cheaper than those offered by home firms, without providing noticeable advantages in market shares. This is especially true for varieties originating from Asian economies. Further, there is considerable heterogeneity in these differences: While imports in richer EU economies are systematically cheaper than those offered by home firms, the opposite is true in poorer Eastern European countries. This suggests not only that imported and domestic varieties differ but also that the size and direction of these differences is country-specific and varies with income levels.

These observations also raise the natural question, if trade affects welfare of poor and rich households differently. As stated earlier and suggested by Figure 1, consumption baskets vary with income, with poorer households consuming a larger share of cheaper, less sophisticated varieties of products. It therefore appears likely that the observed differences between imported and domestic varieties translate into unequal gains for households across the income distribution in each country. The goal of the structural framework presented in the next sections is to answer this question more rigorously and to determine if and for which countries the impact of trade on the cost of living counteracts or amplifies the observed inequality in income.

3 Model

3.1 Model Features

In order to quantify how unequal the gains from trade are, I develop and employ a structural model for several reasons: Most importantly, in order to measure and aggregate individual welfare across the income distribution and thousands of product categories, I need to quantify how and how differently households value each variety, which requires at least some basic structure on the demand side. Further, the suggestive evidence presented above was based on unit values and while the observation that low unit values do not correspond to higher market shares in the data is suggestive of vertical differentiation, it is nevertheless important to rigorously disentangle the impact of quality on unit values from that of wages and productivity. Finally, vertical differentiation is plausibly endogenous. The price differentials presented in the previous section are equilibrium outcomes and are realistically affected by trade. The model therefore incorporates an explicit supply side in order to predict how countries vertically specialize in response to trade.

The resulting multi-sector model of international trade has four key features, which are motivated by the stylized facts of the previous section. First, consumer demand is non-homothetic, i.e. demand for higher-quality varieties is potentially increasing in income. Second, countries differ in their available production technologies and endogenously choose the quality of their products. Third, in order to be able to make statements on how unequal the effects of international trade are, each country is populated by a distribution of heterogeneous households which will earn different incomes. Finally, trade is costly: There are both per-unit trade costs as well as a fixed cost of selling to a market as in Melitz (2003). The demand side of the model builds on Handbury (2013), who estimates a non-homothetic demand system for groceries in the United States. I extend her theoretical framework by adding a structural supply side and the estimation routine to a setting where household-specific expenditures are not available to the researcher.

3.2 Households

I assume that each country is populated by a distribution of households i which are endowed with l_i units of labor. There are two major types of goods in the economy: A set of differentiated, manufacturing goods, $\mathbf{x}$, and services z. Within the manufacturing sector, there are many different products $k = 1, ..., K$ a household can buy, e.g. cars or coffee, and in equilibrium, each of these products will be consumed at a nonzero amount. I assume that all manufacturing goods are tradable and that labor is perfectly mobile across sectors k.

Firms in each country produce differentiated varieties $j = 1, ..., J$ of products, e.g. German, French, and Japanese companies each produce a unique type of car. These varieties may differ in terms of vertical and horizontal characteristics as well as in the price p_{jk} which consumers pay for them. Figure 3 summarizes the consumption decisions of consumers.

In order to allow for both types of product differentiation, I assume that household utility for

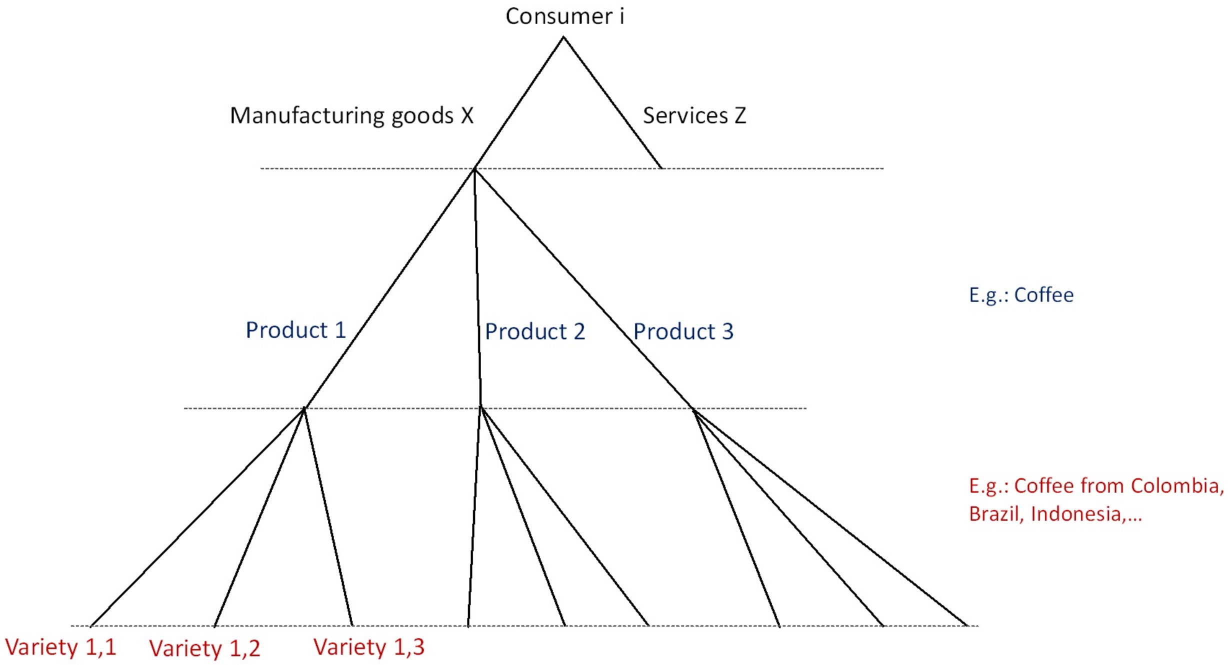

Figure 3: Overview of Consumer Choices

a variety j is given by

$$u_{jk}^{(i)} = x_{jk} e^{\frac{q_{jk}+\varepsilon_{ijk}}{\alpha(z_i)}}.$$

where x_{jk} denotes the quantity consumed. It is multiplied by a taste or demand shifter which depends on the quality of the product q_{jk} as well as an idiosyncratic valuation of the respective variety, ε_{ijk}. q_{jk} captures that varieties are vertically differentiated, i.e. that certain product characteristics are superior to others and are preferred by all consumers. Horizontal differentiation implicitly enters through ε_{ijk}, which allows consumers to have differing tastes for each j, but does not affect the mean utility provided by each variety. Intuitively, q_{jk} captures dimensions such as reliability or horse power of a car, while ε_{ijk} rather reflects differences in color or appearance.

Finally, in order to introduce non-homotheticities in demand, I follow Fajgelbaum, Grossman, and Helpman (2011) and assume that services z and product quality are complements in utility, i.e. that the marginal utility of quality is increasing in the consumption of services, $\alpha'(z_i) < 0$. Under this assumption, richer households will have a higher demand for higher quality varieties than poorer households as long as services are a normal good. Intuitively, this assumption implies for example, that renters of an expensive apartment would benefit more from higher-quality furniture, and vice versa.

I assume that the overall utility which a household i receives from buying product k is given by the sum of the individual components $u_{jk}^{(i)}$:

$$u_k^{(i)} = \sum_{j \in J_k} u_{jk}^{(i)}. \tag{4}$$

In this case, it will be optimal for households to buy only one variety of a product. More specifically, households will choose variety j^* if it dominates all other varieties in terms of utility, i.e. if

$$\begin{aligned} & u_{j^*k}^{(i)} \geq u_{jk}^{(i)} \\ \Leftrightarrow \quad & x_{ij^*k} e^{\frac{q_{j^*k} + \varepsilon_{ij^*k}}{\alpha(z_i)}} \geq x_{ijk} e^{\frac{q_{jk} + \varepsilon_{ijk}}{\alpha(z_i)}}; \qquad \forall j \in J_k, \end{aligned}$$

where J_k denotes the set of all varieties that are available to the household.

As is common in the discrete-choice literature, I assume that the idiosyncratic valuation ε_{ijk} can be captured by a distribution. In particular, I make the common assumption that it follows a type 1 extreme value distribution with location parameter 0 and scale parameter 1. This assumption greatly enhances the tractability of the household decision and delivers a closed-form solution for the choices made by consumers. Specifically, as shown in Appendix C.1, the probability that a consumer with service consumption z_i chooses variety j^*, can be written in the familiar logit form

$$\Pr(i \mapsto j^*) = \frac{\exp[q_{j^*k} - \alpha(z_i) \ln p_{j^*k}]}{\sum_{j \in J_k} \exp[q_{jk} - \alpha(z_i) \ln p_{jk}]} \tag{5}$$

where p_{jk} denotes the price of a variety j of product k.

Equation (5) has a very intuitive interpretation. The numerator captures all characteristics associated with variety j while the denominator reflects those of all competing products. It is therefore straightforward to show that varieties with higher quality and lower prices are more attractive to consumers in general and will therefore have a higher probability of being chosen. Also notice that the idiosyncratic term ε_{ijk} is not present in this equation which implies that horizontal product characteristics do not affect consumer choices on average.

Demand is non-homothetic through the term $\alpha(z_i)$ which premultiplies the price p_{jk}. This implies that households place different weights on varieties' prices in their consumption decision and therefore make different choices. Given that $\alpha(z_i)$ therefore effectively governs how price-elastic households behave, I will refer to it as the price coefficient throughout the paper, which is in line with the terminology used in the structural IO literature. Since $\alpha'(z_i) < 0$, $\alpha(z_i)$ will be lower for high-income consumers, consistent with the prior that poorer households are more income-restricted and therefore more price sensitive than richer households.

The consumer-specific choice probabilities (5) can also be used to derive an expression for the total share spent on a variety in a country. For this purpose, I explicitly introduce a subscript m indicating the respective importing country (market) to highlight the channels through which expenditure shares will differ across countries. By the law of large numbers, the expenditure share of exporter j of product k in market m is equal to

$$\begin{aligned} s_{jkm} &\equiv \frac{p_{jkm}x_{jkm}}{\sum_{j'\in J_{km}} p_{j'km}x_{j'km}} \\ &= \sum_{i\in I_m} \frac{E_{ikm}}{\sum\limits_{i'\in I_m} E_{i'km}} \Pr(i \mapsto j)_m, \end{aligned} \tag{6}$$

i.e. s_{jkm} is an expenditure weighted average of the individual choices, with E_{ikm} being the amount spent on product k by consumer i.

In order to solve for E_{ikm} in equation (6) and to construct aggregate welfare measures later on, I also need to specify how consumers distribute their expenditures across products. For simplicity, I assume that the total utility over manufacturing goods, $U_M^{(i)}(\mathbf{x}, z)$ is of Cobb-Douglas form with

$$U_M^{(i)}(\mathbf{x}, z) = \sum_{k=1,..K} \omega_k \ln u_k^{(i)}.$$

where $\{\omega_k\}_{k=1}^K$ represent consumption weights for all products $k = 1, ...K$ and $u_k^{(i)}$ denotes the utility provided by the chosen variety from equation (4).[5] Under this specification, households will spend a constant fraction of their manufacturing expenditure on each product category, i.e.:

$$E_{ikm} \equiv p_{jkm}x_{ijk} = \omega_k(y_i - p_{zm}z_i), \qquad \text{if } j = j^* \tag{7}$$

$$= 0, \qquad \text{otherwise.} \tag{8}$$

I do not explicitly model the decision between manufacturing goods and services.[6] I do however, consistent with empirical evidence, allow the share that households spend on services, $\gamma(\cdot)$, to vary with income, i.e.

$$p_{zm}z_i = \gamma(y_i)y_i \tag{9}$$

$$p_{jkm}x_{ijk} = \omega_k(1-\gamma(y_i))y_i. \tag{10}$$

I do not place any theoretical restrictions on $\gamma(y_i)$ other than $\gamma(y_i) \in (0, 1)$. However, since richer households typically spend a larger fraction of their income on services and less on manufacturing goods, one would expect γ to be an increasing function in y_i.

Under the above assumptions, equations (5), (6) and (10) can be used to derive the aggregate expenditure share of exporter j in market m as a function of expenditure shares, product

---

[5] In section 6, I evaluate how sensitive the main results are to imposing the Cobb-Douglas assumption on different aggregation levels.

[6] As shown in Handbury (2013), this is fine as long as services are a normal good. Since the share of income spent on services is empirically even increasing in income, this is arguably a realistic assumption.

characteristics and the income distribution:

$$\begin{aligned} s_{jkm} &= \sum_{i \in I_m} \frac{(1-\gamma(y_i))y_i}{\sum\limits_{i' \in I_m} (1-\gamma(y_{i'}))y_{i'}} \frac{\exp[q_{jkm} - \alpha(z_i)\ln p_{jkm}]}{\sum_{j' \in J_{km}} \exp[q_{j'km} - \alpha(z_i)\ln p_{j'km}]} \\ &= \sum_{i \in I_m} \frac{(1-\gamma(y_i))y_i}{\sum\limits_{i' \in I_m} (1-\gamma(y_{i'}))y_{i'}} \frac{\exp[q_{jkm} - \alpha\left(\gamma(y_i)\frac{y_i}{p_{zm}}\right)\ln p_{jkm}]}{\sum_{j' \in J_{km}} \exp[q_{j'km} - \alpha\left(\gamma(y_i)\frac{y_i}{p_{zm}}\right)\ln p_{j'km}]}. \end{aligned} \tag{11}$$

Expenditure shares can vary across markets through essentially three channels: First, the set of available varieties J_{km} may differ, i.e. certain countries do not export to others. Second, the same exporting country may offer different products in different markets or sell at different prices. Hence, conditional on J_{km}, equilibrium prices p_{jkm} and qualities q_{jkm} may vary across markets. Finally, even if the set of available products were the same across countries, expenditure shares will differ through demand being non-homothetic. The observation that richer households have a higher demand for high-quality varieties will e.g. translate on the aggregate into high-quality producers having higher shares in richer countries.

All three channels will also depend on each other through the supply side: The set of available varieties J_{km} will depend on the respective demand in market m, as for example high-quality producers will face a higher demand in rich countries. In order to run meaningful counterfactuals it will therefore be important to understand how demand and supply jointly determine the set of available products in each market.

Equation (11) will be the main structural estimation equation I use on the demand side. The fact that expenditure shares can be expressed as a function of prices p_{jkm}, quality q_{jkm} and the income distribution alone is important, since each of those statistics is either observable or can be readily estimated. This will allow me to recover the distribution of the α_i's in each market and product category, which governs the extent to which demand is non-homothetic and will be an important determinant of inequality in the gains from trade.

3.3 Supply Side

Each country c is populated by a continuum of potential producers which can sell their products at home and abroad. Firms produce horizontally as well as vertically differentiated varieties and while horizontal differentiation is costless and treated as exogenous, the firm's marginal cost $mc_{jm}(q_{jm})$ depends on the level of quality q_{jm} set by the firm.[7] In order to sell to a market, firms additionally have to pay a fixed cost f, which is product-specific. Each firm will therefore only enter a market m if its profits there exceed f, i.e. as long as

$$\frac{p_{jm} - mc_{jm}}{p_{jm}} s_{jm} E_m \geq f, \tag{12}$$

---

[7] All equations of this subsection hold separately for each product category. For notational convenience, I therefore drop the product subscript k throughout this subsection.

where E_m denotes the aggregate expenditure in market m and s_{jm} the share of expenditure spent in country m on firm $j's$ varieties as defined by equation (11). Since s_{jm} depends negatively on entry by other firms, the equilibrium number of firms selling in market m will be bounded and firms will choose to enter market m until (12) holds with equality.

I assume that conditional on entry, firms simultaneously choose prices p_{jm} and quality q_{jm} for each market to which they sell in order to maximize profits. This implies that exporters can in principle tailor their products to the respective markets they sell to and respond to local demand conditions.[8] The optimization problem of each firm can therefore be summarized as

$$(q^*_{jm}, p^*_{jm}) = \underset{p_{jm}, q_{jm}}{\arg\max}\ \pi_{jm} = [p_{jm} - mc_{jm}(q_{jm})] \frac{s_{jm} E_m}{p_{jm}} - f. \tag{13}$$

The profit maximization condition (13) implies that the optimal price and quality choices of each firm must satisfy two first-order conditions:

$$[p_{jm}]: \qquad \frac{p_{jm} - mc_{jm}(q_{jm})}{mc_{jm}(q_{jm})} = \frac{\sum\limits_{i \in I_m} E_{im} s^{(i)}_{jm}}{\sum\limits_{i \in I_m} E_{im} \alpha_i s^{(i)}_{jm} \left(1 - s^{(i)}_{jm}\right)} \tag{14}$$

$$[q_{jm}]: \qquad \frac{\partial mc_{jm}(q_{jm})}{\partial q_{jm}} s_{jm} = (p_{jm} - mc_{jm}(q_{jm})) \sum_{i \in I_m} \frac{E_{im}}{E_m} s^{(i)}_{jm} \left(1 - s^{(i)}_{jm}\right). \tag{15}$$

The first equation (14) characterizes the markup of the firm conditional on the chosen level of product quality q_{jm}. The optimal markup has two intuitive properties: First, it is declining in the price coefficient α_i: The more price-elastic consumers are, the lower will be the price and markup the firm is able to charge. In the limit, as $\alpha_i \to \infty$, firms would charge exactly marginal cost. Second, the markup is increasing in firm size, consistent with the idea that market power allows firms to generate higher profits per unit. As $s^{(i)}_{jm} \to 0, \forall i$, the markup will go to $\sum_{i \in I_m} E_{im} / \sum_{i \in I_m} \alpha_i E_{im}$, while it will be unbounded if $s^{(i)}_{jm} \to 1, \forall i$.

An important difference compared to a representative agent framework is that the price coefficients α_i are weighted by the firm's shares for different consumer groups, $s^{(i)}_{jm}$. This means that the firm will be more responsive to the price elasticity of certain consumer groups, but less to that of others. The effective price elasticity faced by the firm is hence an endogenous variable and depends on the firm's characteristics through $s^{(i)}_{jm}$. A high-quality producer for example will typically have a larger share among richer, less price elastic households. Its effective price elasticity $\sum_{i \in I_m} \alpha_i s^{(i)}_{jm} \left(1 - s^{(i)}_{jm}\right)$ is hence lower and allows it to charge a higher markup.[9]

The first-order condition with respect to quality (15), essentially requires the marginal revenue

---

[8] This is a relatively common assumption made in the literature. As noted by Feenstra and Romalis (2014), Volkswagen selling lower-quality versions of their cars in Latin America is an example of such firm behavior.

[9] Technically, this also requires that the shares $s^{(i)}_{jm}$ are smaller than 50% but this will be the case for the vast majority of observations in the empirical part.

of increasing q_{jm} to equal the additional cost. This additional cost depends on the firm-specific slope of the marginal cost curve with respect to quality $\partial mc_{jm}(q_{jm})/\partial q_{jm}$: A larger increase in the marginal cost associated with an increase in q_{jm} will lead to a lower optimal quality choice, all else equal. The marginal revenue on the other hand depends largely on demand-side factors and competition. As will become clear below, the effective price elasticity of a firm's consumers will also be an important determinant of the demand for quality and in turn of the firm's optimal choice of q_{jm}.

I assume that the marginal cost of serving market m is product- and country-pair-specific. This accounts for shipping costs between countries c and m, tariffs, as well as for any other importer- and exporter-specific factors that affect unit costs. The assumption also implies that firms from c selling to m share a common marginal cost function, $mc_{jm}(\cdot) = mc_{cm}(\cdot)$. In this case, firms $j \in c$ selling to m face the same first-order conditions and will make identical choices when selling to m.[10] As shown in the Appendix, the first-order conditions can then be slightly simplified to

$$\frac{p_{cm} - mc_{cm}(q_{cm})}{mc_{cm}(q_{cm})} = \frac{N_{cm} \sum\limits_{i \in I_m} E_{im} s_{cm}^{(i)}}{\sum\limits_{i \in I_m} \alpha_i E_{im} s_{cm}^{(i)} \left(N_{cm} - s_{cm}^{(i)}\right)} \tag{16}$$

$$\frac{\partial mc_{cm}(q_{cm})}{\partial q_{cm}} \frac{s_{cm}}{N_{cm}} = (p_{cm} - mc_{cm}(q_{cm})) \sum_{i \in I_m} \frac{E_{im}}{E_m} \left[\frac{s_{cm}^{(i)}}{N_{cm}} - \left(\frac{s_{cm}^{(i)}}{N_{cm}} \right)^2 \right], \tag{17}$$

with the share of expenditure in country m spent on goods from c given by

$$s_{cm} = N_{cm} s_{c^* m} = \sum_{i \in I_m} \frac{E_{im}}{E_m} \frac{\exp(q_{cm} + \ln N_{cm} - \alpha_i \ln p_{cm})}{\sum_{j'} \exp(q_{c'm} + \ln N_{c'm} - \alpha_i \ln p_{c'm})}. \tag{18}$$

In anticipation of the empirical model, I assume that the marginal cost can be characterized by the following functional form:

$$mc_{cm} = \exp(m_{0,cm} \cdot q_{cm}^{m_{1,cm}}), \tag{19}$$

where I allow $m_{0,cm}$ and $m_{1,cm}$ to be potentially source-destination specific. The functional form for the marginal costs does in principle not need to be exponential but can for example also be linear or log-linear. There are however several restrictions and desirable properties that guide the choice of this relationship. First, two source-destination specific parameters can be separately identified from equations (16) and (17) and I denote the parameters of $mc_{cm}(\cdot)$ with $m_{0,cm}$ and $m_{1,cm}$.[11]

---

[10] To be exact, if any set of firms optimally choose p_{cm} and q_{cm}, then p_{cm} and q_{cm} will also satisfy the first-order conditions of the remaining firms. In principle, the firm problem might have multiple solutions, i.e. some firms may want to choose a differing p_{cm} and q_{cm} to attract a different consumer group. In practice, I found this not to be the case given my estimates. To test this, I checked whether or not a single company might have an incentive to deviate both using different starting values when solving its first-order condition as well as by using a grid search.

[11] In principle, one can easily allow $m_{0,cm}$ and $m_{1,cm}$ to depend on other observables: Differences in the parameters of this cost functions by exporter c for example will likely be correlated with differences in labor costs in each source country. Shipping costs are a major argument why the coefficients will not only be source- but source-destination-specific as the shipping distance depends on the respective country-pair. Section 5.1 provides a sense on how the

Second, the choice of $mc_{cm}(\cdot)$ should ensure interior solutions for both the optimal price and quality set by each firm, in line with the data. This is important, since if $mc_{cm}(\cdot)$ is too convex or too concave in q_{cm}, firms will optimally want to choose either an infinitively low level of quality or an infinitively high one. Since costs and prices are closely related through equation (14) and the price enters utility in logs while quality does so linearly, an exponential functional form largely avoids these problems in practice. In the empirical application, I found logarithmic or linear relationships to frequently be unable to deliver interior solutions for p_{cm} and q_{cm} and therefore to be inconsistent with the data.

Third, regarding the term in brackets, I follow Feenstra and Romalis (2014), who assume a unit cost of $w_c q_{cm}^{1/\theta}/\varphi_{cm}$ where w_c is the source country's wage, and φ_{cm} its destination-specific productivity. In my specification, w_c and φ_{cm} are absorbed into $m_{0,cm}$ while $m_{1,cm} = 1/\theta$, which reflects the returns to quality, is allowed to vary by source and destination. This specification also has the additional benefit that it implies a constant elasticity of the log marginal cost with respect to quality equal to $m_{1,cm}$, which gives $m_{1,cm}$ a straightforward interpretation.

Fourth, this functional form implies an intuitive, algebraic solution for the optimal quality choice q^*_{cm} given by

$$q^*_{cm} = \left(\frac{1}{\boldsymbol{\alpha}_{cm} m_{0,cm} m_{1,cm}}\right)^{\frac{1}{m_{1,cm}-1}} \tag{20}$$

with

$$\boldsymbol{\alpha}_{cm} = \frac{\sum_{i\in I_m} \alpha_i [s^{(i)}_{c^*m} - (s^{(i)}_{c^*m})^2]}{\sum_{i\in I_m} s^{(i)}_{c^*m} - (s^{(i)}_{c^*m})^2}. \tag{21}$$

where $s^{(i)}_{c^*m} := s^{(i)}_{cm}/N_{cm}$ denotes the expenditure on varieties of an individual firm from country c. I refer to $\boldsymbol{\alpha}_{cm}$ as the effective price elasticity of country c in market m. $\boldsymbol{\alpha}_{cm}$ is essentially a weighted average of household-specific price coefficients α_i with the weights depending on the expenditure shares $s^{(i)}_{c^*m}$ which different groups spend on varieties produced by c: The higher $\boldsymbol{\alpha}_{cm}$ is, the more price-elastic are producer c's consumers on average.

The optimal quality choice is driven by both supply-side and demand side factors: On the supply side, both cost parameters $m_{0,cm}$, and $m_{1,cm}$ affect quality negatively in the empirically relevant case $m_{1,cm} > 1$: First, firms with a comparative advantage in producing quality, i.e. with a lower elasticity of the (log) marginal cost with respect to q_{cm}, $m_{1,cm}$, choose a higher q^*_{cm}. Second, the optimal quality choice is decreasing in $m_{0,cm}$, in line with previous studies that document a positive relationship between productivity and quality (see e.g. Kugler and Verhoogen (2012), and Faber and Fally (2017)).

The main demand-side determinant of q^*_{cm} is the effective price elasticity $\boldsymbol{\alpha}_{cm}$. In particular, $\boldsymbol{\alpha}_{cm}$ affects quality negatively: Since higher-quality varieties are more costly to produce and hence more expensive, there will be less demand for them when households are very price-elastic. The optimal q^*_{cm} is therefore lower.

It is important to note that, in contrast to trade models that are based on CES preferences,

---

estimated parameters correlate with country characteristics.

the effective price elasticity $\boldsymbol{\alpha}_{cm}$ is an endogenous variable and depends both on firm choices and competition: First, the firm's chosen q^*_{cm} feeds back into $\boldsymbol{\alpha}_{cm}$: If country c e.g. produces an expensive, high-quality variety it will have a relatively higher market share among richer consumers which are less price elastic and $\boldsymbol{\alpha}_{cm}$ will be lower in absolute value. Second, and more importantly from a trade perspective, $\boldsymbol{\alpha}_{cm}$ also depends on the composition of both domestic and foreign competition: The presence of many low-quality competitors would for example also lower the firm's effective price elasticity $\boldsymbol{\alpha}_{cm}$, but through diminishing the firm's shares $s^{(i)}_{c^*m}$ among low-income consumers.

Trade liberalization From equation (20) it is also evident that $\boldsymbol{\alpha}_{cm}$ is a sufficient statistic for the direction in which domestic firms vertically differentiate in response to international trade: In particular, the direction in which quality changes with trade depends entirely on whether or not

$$\boldsymbol{\alpha}^{\text{Trade}}_{cm} = \frac{\sum_{i\in I_m} \alpha_i [s^{(i)}_{c^*m} - (s^{(i)}_{c^*m})^2]}{\sum_{i\in I_m} s^{(i)}_{c^*m} - (s^{(i)}_{c^*m})^2} \quad \gtreqless \quad \boldsymbol{\alpha}^{Autarky}_{cm} = \frac{1}{|I_m|} \sum_{i\in I_m} \alpha_i. \tag{22}$$

In autarky, demand from all consumers has to be met by domestic firms and so the effective price elasticity essentially becomes the simple average $\boldsymbol{\alpha}^{Autarky}_{cm} = \frac{1}{I} \sum_i \alpha_i$. Trade introduces foreign varieties, which differ in terms of how attractive they are for each consumer group and therefore affect the distribution of group-specific expenditure shares $s^{(i)}_{c^*m}$ unevenly. If this alters the equilibrium $\boldsymbol{\alpha}_{cm}$ to a value with $\boldsymbol{\alpha}^{\text{Trade}}_{cm} < \boldsymbol{\alpha}^{\text{Autarky}}_{cm}$, which would be the case if predominantly lower-quality varieties enter, trade will induce Home firms to quality-upgrade while the opposite would be true if $\boldsymbol{\alpha}^{\text{Trade}}_{cm} > \boldsymbol{\alpha}^{\text{Autarky}}_{cm}$.

In the empirical application below for example, trade increases the availability of lower-quality products in the EU through greater access to cheaper varieties from Eastern Europe and Asia. Domestic companies hence become disproportionately less attractive to poorer households which lowers the average price elasticity of their remaining consumers. Since the optimal quality choice of firms depends negatively on the price elasticity of a firm's consumers, companies have incentives to specialize in higher-quality varieties when selling to on average richer households.

3.4 Equilibrium and Price Index

The equilibrium can be characterized by a set of quantities $\{x_{ckm}, z_{im}\}$, prices $\{p_{ckm}\}$, and quality choices $\{q_{ckm}\}$ such that households maximize utility, firms maximize profits, and the labor and goods markets clear.

In order to close the model, I assume that services z are homogeneous and produced with constant returns to scale and productivity w_c. In that case, the price of services in country c will be equal to the wage. Further, I allow a fraction of services to be freely traded and consider only equilibria in which these services are produced in each country. These assumptions are equivalent to assuming costless trade in non-manufacturing goods as in Eaton and Kortum (2002) or in a homogeneous outside good as in Chaney (2008) and imply in this case that the equilibrium prices p_{zm} will be equal across countries and can be normalized to 1.

Utility maximization results in the conditions regarding product choice (9) and (10), and variety choice (11). On the supply side, the free entry condition (12) together with the first-order conditions for prices and quality, (14) and (15), characterize firm behavior in equilibrium. Goods market clearing is evident from the expenditure share equation (11): Given prices and quality set by exporters, households demand quantities $\mathbf{x}$, which are readily supplied by exporters. The household budget will hold with equality in equilibrium and Walras' law implies that the labor market has to clear.

In order to quantify the gains from trade in the empirical part, I exploit the structure of the model and construct aggregate, household-specific price indexes and compute their change when moving to autarky. Specifically, as derived in Appendix C.3, the model implies a closed form solution for the manufacturing price index for a household with income y_i living in country m equal to

$$\mathbb{P}_m(y_i) = \exp\left(\sum_k \omega_k \left[1 - \ln\frac{\omega_k}{p^*_{ikm}} - \frac{q^*_{ikm} + \varepsilon_{ik}}{\alpha(\gamma_i y_i)}\right]\right). \tag{23}$$

$\mathbb{P}_m(y_i)$ is individual-specific through $\alpha(\cdot)$ and the variety choices that household i makes, i.e. through q^*_{ikm} and p^*_{ikm} of the chosen varieties of each product k. Hence, consumer welfare ultimately depends on the price-quality combination of each available variety in market m. As $\alpha(\cdot) > 0$, it is easy to see that price increases raise the overall price index while an increase in the quality of a variety, all else equal, lowers $\mathbb{P}_m(y_i)$. Also notice that the price index depends on ε_{ik}, i.e. the idiosyncratic utility draw of household i. It is through this idiosyncratic term, that households benefit from more variety: Each additional available variety comes with a new draw ε_{ijk}, which if high enough, increases consumer utility, even if it is otherwise equal to other available varieties. Hence, the model captures standard love-for-variety effects as for example present in a CES framework.[12]

In order to summarize how unequal the gains from trade are across the income distribution, I follow Fajgelbaum and Khandelwal (2016) and compute the percentage change in $\mathbb{P}_m(y_i)$ when moving to autarky for a household at the 20th percentile of the income distribution and compare it to the change for a household at the 80th percentile.

The move to autarky affects the price index in two ways: First, it forces households to buy domestic varieties which harms those households who preferred the price-quality combination of foreign varieties in the trade-case or had a high idiosyncratic taste for them. Second, price and quality adjustments by Home firms in response to the ban on foreign competition change q_{jk} and p_{jk} of the domestic varieties in autarky.

Notice that since poor and rich household differ in their valuation for quality through $\alpha(\gamma_i y_i)$, both channels can affect households across the income distribution asymmetrically: If foreign varieties for example were on average cheaper and of lower quality than their domestic counterparts,

---

[12] In fact, when firms are equal in terms of p_{jkm} and q_{jkm} and there is a representative consumer, a discrete choice framework implies the same expenditure function as the CES (See Anderson, De Palma, and Thisse (1987)).

poorer households will experience stronger increases in the cost of living. Changes in q_{jk} and p_{jk} by domestic firms can either mitigate these unequal gains or amplify them further: Quality-downgrading actually benefits poorer households by shifting domestic production from higher-quality to cheaper, lower-quality varieties, thereby countering some of the welfare losses associated with losing access to foreign varieties. Quality-upgrading on the other hand would make the gains from trade even more unequal.

4 Data and Estimation

4.1 Data

I use data on trade flows and matched production data for the European Union in highly disaggregated 8-digit product categories between 1989 and 2013. I focus on European data for essentially two reasons. First, Eurostat provides data on domestic production on a high level of disaggregation which can be matched to data on trade flows. As countries generally consume a substantial share of domestic varieties, it is important to account for these when estimating the gains from trade. Second, the assumption that exporters are able to freely choose the quality of their products in the long run is more reasonable for richer economies, which is especially important for the evaluation of a counterfactual move to autarky: Since higher trade barriers in the EU would imply predominantly exit of lower-quality producers, it is plausible that EU countries are capable of producing lower-quality goods (although at potentially high costs). The mirror assumption that e.g. firms from developing economies are able to produce varieties at the quality frontier is arguably a stronger assumption.

For extra-EU trade, the data is provided by the respective trading partners on the basis of customs declarations and covers in principle all imports and exports declared by member states.[13] Information on Intra-EU trade flows is provided on the basis of so called intrastat declarations. Member states have to ensure that at least 97 % of the country's trade value is covered.

I concord data on trade flows and production using the concordance developed by Van Beveren, Bernard, and Vandenbussche (2012). More specifically, I match trade data, which is reported by cn8 categories to prodcom categories (pc8) which are used to classify production data. To account for changes in categories over time, I apply the above concordance to create categories that are consistent over time (pc8plus). I aggregate product categories further whenever a country does not produce a product in any time period. Finally, in order to reduce the impact of implausibly high or low unit values, I exclude those values that are 30 times higher or lower than the median over all exporter-importer trade flows within a product category. More details on the construction of the dataset are provided in Appendix D.

---

[13] Before 2010 it was allowed to exclude transactions whose value and net mass were lower than 1000 Euro or 1000 kg. Given that these transactions will mostly cover smaller transactions, it is less of a concern that these missing observations will significantly affect my results. Additionally, as noted by Eurostat, the trend of customs declarations being more and more done electronically has ensured a very high coverage, even when exporters were not legally required to report certain transactions.

Household income data is taken from Eurostat's database on income and living conditions, which provides data on disposable household income by decile, quartile and the five highest and lowest percentiles for each country and year. I fit these numbers using log-normal distributions with country-time-specific location and scale parameters. As shown in Figure D3 in Appendix D, the fitted values match the actual ones well. Table D4 summarizes the obtained estimates of the parameters of the log-normal normal distribution. The location parameters range from 10.4 for Luxembourg to 7.6 for Romania. The scale parameters which reflect the degree of income inequality within a country tends to be small in the Scandinavian countries (0.46 - 0.52) and bigger in Southern and Eastern Europe with values around 0.66 in Portugal or Greece.

Data on the number of companies by country and industry is taken from Eurostat's *Structural Business Statistics*, which reports the number of enterprises by country and year in distinct 249 Nace Rev. 1 industries. Since the first 4 digits of each pc8 code coincide with those of the Nace classification, it is straightforward to match each pc8 to a Nace code. I then distribute the number of companies in each Nace industry on the respective pc8 codes proportional to production values.

In order to allow the relative demand for services to vary with income, I use National Accounts data on consumption of services and other goods from Eurostat. The percentage of service consumption relative to total consumption ranges from 26.6% in Estonia to 56.6% in Spain. For the estimation, I also rely on data on bilateral geographic distance as well as aggregate country characteristics. I use the Euclidean distance between the most populated cities as provided by CEPII. Population data is taken from the United Nation's World Population Prospects, the CIA's world factbooks, as well as the National Statistics of Taiwan. Data on GDP per capita is taken from the International Monetary Fund's World Economic Outlook Database and the World Bank's World Development Indicators.

For my instrumenting strategy, I use several additional sources of data. To construct a set of alternative instruments, I use data published by the U.S. Census Bureau on cost, insurance and freight (CIF) charges of exporters selling to the U.S between 1989 to 2012. I also use information on imports and exports of EU countries from Comtrade to create a second set of alternative instruments following Feenstra and Romalis (2014). Finally, data on exchange rates is taken from Feenstra, Inklaar, and Timmer (forthcoming). Summary statistics of the final dataset are provided in Table A1 in the Appendix.

4.2 Estimation

As seen in Section 3, the theoretical model delivers three main structural equations which characterize the equilibrium and which are sufficient to quantify the parameters of the model. Specifically, I exploit the closed-form solution for the country-pair-specific expenditure shares

$$s_{cm} = \sum_{i \in I_m} \frac{(1-\gamma_i)y_i}{\sum_{i' \in I_m}(1-\gamma_{i'})y_{i'}} \frac{\exp(q_{cm} + \ln N_{cm} - \alpha(\gamma_i y_i)\ln p_{cm})}{\sum_{j'} \exp(q_{c'm} + \ln N_{c'm} - \alpha(\gamma_i y_i)\ln p_{c'm})}, \tag{ES}$$

the free-entry condition

$$N_{cm} = \frac{(p_{cm} - mc_{cm}) s_{cm} E_m}{f p_{cm}}, \tag{FE}$$

as well as profit-maximizing by firms,

$$(q^*_{jm}, p^*_{jm}) = \underset{p_{jm}, q_{jm}}{\arg\max} \;\; \pi_{jm} = [p_{jm} - mc_{jm}(q_{jm})] \frac{s_{jm} E_m}{p_{jm}} - f, \tag{PM}$$

which implies the set of first-order conditions (16) and (17).[14] Since each equation holds for every product category, I dropped the product subscripts k for notational convenience and will continue to do so for the remainder of the paper, unless stated otherwise.

I quantify the model in two steps: First, I use equation (ES) to estimate the demand side parameters $\alpha(y_i)$ and infer product quality using data on expenditure shares, unit values and the income distribution in each country.[15] In order to separate demand from supply side determinants of expenditure shares, I use an instrument which shifts the supply curve, holding the demand curve constant.

In the second step, given the demand side parameters, I infer marginal costs and the fixed cost of exporting from the first-order conditions for price and quality implied by (PM), as well as the free entry condition (FE).

4.2.1 Step 1: Demand Side

In order to estimate the system of equations defined by (ES), I make two additional functional form assumptions. First, I assume that the price coefficient $\alpha(z_i)$ is of log-linear form and can be written as:

$$\alpha(z_i) = \widetilde{\alpha}_0 + \widetilde{\alpha}_1 \ln z_i.$$

Under this specification, $\widetilde{\alpha}_0$ reflects the average price elasticity across the population in a country, while $\widetilde{\alpha}_1$ governs how much this elasticity varies with z_i and hence ultimately across the income distribution. I have two priors for the parameters $\widetilde{\alpha}_0$ and $\widetilde{\alpha}_1$: For most goods, $\widetilde{\alpha}_0$ will realistically be positive so that households on average prefer cheaper over expensive varieties, all else equal. Additionally, one would expect $\widetilde{\alpha}_1$ to be negative given that poorer households are more income-restricted and typically behave more price sensitive than richer households. Section 5 will show that these predictions are largely supported by the data.

Second, I assume that the share of income spent on services can be written as

$$\gamma_i(y_i) = e^{\gamma_0 + \gamma_1 \ln(y_i)}.$$

---

[14] Notice that equation (ES) uses the result of firms from c selling to m making the same choices. This result is not strictly required however and as shown in the Appendix, also a first-order Taylor approximation for example results in the same set of equations with q_{cm} and p_{cm} then referring to the average quality and price of exporters from c selling to m.

[15] Quality q_{cm} will at this stage be only jointly identified together with the (log) number of firms $\ln N_{cm}$.

This functional form is motivated by evidence that richer countries and household typically spend a larger share of their income on services and less on manufacturing goods.[16]

Under both assumptions, the main estimation equation on the demand side, (ES), then becomes

$$
\begin{aligned}
s_{cm} &= \sum_{i\in I_m} \frac{E_{im}}{\sum_{i'\in I_m} E_{i'm}} \frac{\exp(q_{cm} + \ln N_{cm} - (\widetilde{\alpha}_0 + \widetilde{\alpha}_1 \ln(e^{\gamma_0+\gamma_1 \ln(y_i)} y_i)) \ln p_{cm})}{\sum_{c'} \exp(q_{c'm} + \ln N_{c'm} - (\widetilde{\alpha}_0 + \widetilde{\alpha}_1 \ln(e^{\gamma_0+\gamma_1 \ln(y_i)} y_i)) \ln p_{c'm})}. \\
&= \sum_{i\in I_m} \frac{E_{im}}{\sum_{i'\in I_m} E_{i'm}} \frac{\exp(q_{cm} + \ln N_{cm} - (\alpha_0 + \alpha_1 \ln(y_i)) \ln p_{cm})}{\sum_{c'} \exp(q_{c'm} + \ln N_{c'm} - (\alpha_0 + \alpha_1 \ln(y_i)) \ln p_{c'm})} \qquad (24)
\end{aligned}
$$

with $\alpha_0 \equiv \widetilde{\alpha}_0 + \widetilde{\alpha}_1 \gamma_0$ and $\alpha_1 \equiv \widetilde{\alpha}_1(1+\gamma_1)$.[17] I assume that y_i is log-normally distributed and I allow the parameters of this distribution to be different depending on market m, to reflect that countries differ in terms of the income distribution:

$$
y_i \sim LN(\mu_m, \sigma_m), \quad i \in I_m. \qquad (25)
$$

I obtain μ_m and σ_m by fitting the distribution in each country to data on country-specific income quantiles. As shown in Figure D3, the resulting distributions match these quantiles very well.

The observables in equation (24) are s_{cm}, p_{cm}, and y_i. Quality q_{cm} and the equilibrium number of firms N_{cm} are not directly observed, and even if the parameters α_0 and α_1 were known, only $q_{cm} + \ln N_{cm}$ are identified, not its composition. I do not aim to separably identify these two terms in step 1 but solely pin down the composite term $\xi_{cm} \equiv q_{cm} + \ln N_{cm}$ as well as the parameters α_0 and α_1. In step 2, I use additional data sources on the overall number of firms together with the supply side equilibrium conditions to decompose ξ_{cm} into quality and the equilibrium number of firms in each market.

Equation (24) closely resembles estimation equations delivered by discrete choice random coefficients models of consumer demand in the style of Berry (1994), and Berry, Levinsohn, and Pakes (1995) which have become very popular in the structural empirical IO literature. One notable difference is that this literature typically refers to the term ξ_{cm} as unobserved heterogeneity of a product, while I follow the recent trade literature, particularly Khandelwal (2010), by labeling it quality. Under both interpretations, ξ_{cm} plays the role of a taste shifter: Whenever, conditional on prices, a variety has a higher expenditure share than another one, it must be valued higher by consumers or have certain characteristics which consumers prefer over others.

---

[16] γ_0 and γ_1 can also be backed out relatively easily by using the model prediction that the share of income spent on services in country m equals $p_{zm}z/y = \frac{\sum_{i\in I_m}\gamma(y_i)y_i}{\sum_{i\in I_m} y_i} = \frac{\sum_{i\in I_m} e^{\gamma_0+\gamma_1 \ln(y_i)} y_i}{\sum_{i\in I_m} y_i}$. Using information on each country's share spent on services and the respective income distributions, the distance between right-hand side and left-hand side is minimized at $\gamma_0 = -0.0967$ and $\gamma_1 = 0.0600$ when GDP per capita is denoted in $1,000$ Euros.

[17] I also exclude the expenditure weights $E_{im}/\sum_{i'\in I_m} E_{i'm}$ from the final estimation equation. These weights would imply that richer households are more important for aggregate expenditure shares due to greater total expenditure on each category. Since the weights are very similar across countries however and identification comes from cross-country variation it seems unlikely that they would affect the estimates much.

Instruments Identification of α_0 and α_1 requires a valid instrument for p_{cm}. This is necessary as firm prices will plausibly be correlated with the unobserved characteristics and quality of each variety. In order to overcome this identification problem, I use several approaches to instrument for prices, which lead to similar results. In the baseline specification, I instrument for p_{cm} using the prices set by a country when selling in other markets. The idea of these instruments, brought forward by Hausman (1996) and prominently used for the BLP demand system by Nevo (2001), is that the prices charged in other markets will be correlated with the price set in market m through common cost components such as the costs of inputs. On the other hand, these prices are not affected by local demand and hence satisfy the exclusion restriction. Intuitively, since this instrument captures shifts in producers' costs, it affects and moves the supply curve which in turn allows points on the demand curve to be identified.[18]

As noted by Nevo (2001), there are no obvious guidelines on the set of markets and periods to use when creating these instruments. I construct the instrument as an average over all other markets and periods to minimize the number of observations for which this instrument cannot be used.

A challenge when estimating demand for a large number of products, is that the exclusion restriction might not be satisfied for every product. A potential concern for the validity of the Hausman (1996) - type of instruments for example would be the presence of global demand shocks. I approach this issue by using several alternative instruments and evaluate the robustness of the results to different identification strategies.

A natural alternative is the use of shipping costs (see for example Khandelwal (2010), and Feenstra and Romalis (2014)) which provide exogenous variation in the cost of selling in a particular country. I use two approaches and data sources to construct measures of shipping costs. The first one follows Feenstra and Romalis (2014) and makes use of the fact that exporters report free-on-board (FOB) trade values while those reported by the importing country include cost, insurance, and freight (CIF) incurred by the seller. Using data on both exporter- and importer-reported trade flows, I construct FOB and CIF unit values and use the difference as an instrument for prices charged. Appendix D.5 provides more details on how these instruments were computed. As shown in Table D5, there is a robust, positive first-stage correlation between the instrument and the observed unit values.

I also use a second approach to create a measure of shipping costs based on U.S. trade data. The U.S. Census Bureau publishes data on U.S. imports, explicitly stating CIF charges. I use this information to compute the average charges per unit and kilometer (km) for each product category k and period t and use it as a proxy for the EU equivalent. I then create the instrument as

$$\mathbb{Z}^{(2)}_{cmkt} = \widehat{\text{Cost per unit}} \text{ \& km}_{kt} \cdot \text{Distance}_{c,m}.$$

---

[18] I also follow BLP and Nevo (2001) and make the common implicit assumption that quality and entry can not adjust as flexibly as prices and are predetermined in the current quarter. In particular, this assumption implies that a cost shock may lead firms to change prices but will not trigger an immediate redesign of the product. In that case, exogenous variation in the marginal costs faced by firms will be sufficient to identify α_0 and α_1.

where $\text{Distance}_{c,m}$ denotes the Euclidian distance in km between countries c and m. Appendix D.5 provides more details on this instrument as well along with the first-stage results.

Finally, I use variation in exchange rates over time as additional instrument. While exchange rates are equilibrium objects, individual product categories will arguably have a small impact on the aggregate exchange rate and so changes in exchange rates provide reasonably exogenous shifts in the costs of a firm when selling to market m. To construct this instrument, I use data from the Penn World Tables (Feenstra, Inklaar, and Timmer (forthcoming)) which provides information on each country's exchange rate vis-a-vis the U.S. Dollar between 1950 and 2010 and compute the percentage deviation of each country's exchange rate in period t from its mean.

Estimation Formally, equation (24) implies a set of equilibrium conditions which can be written more compactly as

$$s_{cm} = s(p_{cm}, Y_m, \xi_{cm}; \theta) \tag{26}$$

where s_{cm} are observed expenditure shares and $s(\cdot)$ an expenditure share function defined by equation (24) based on prices, the income distribution Y_m, the composite term $\xi_{cm} = q_{cm} + ln(N_{cm})$, as well as a set of parameters to estimate $\theta = \{\alpha_0, \alpha_1\}$.

As pointed out earlier, the model delivers the same demand system as discrete choice random coefficients models of consumer demand such as BLP which are notorious for being challenging to estimate: $s(\cdot)$ is a sum over hundreds of nonlinear terms and cannot be simplified through reasonable functional-form assumptions on the income distribution. It is therefore necessary to simulate expenditure shares for each consumer in each market and period to evaluate $s(\cdot)$ in each estimation step.

Further, as discussed in more detail by Knittel and Metaxoglou (2008) or Dubé, Fox and Su (2012), GMM estimators of this demand system that have been traditionally used in applied work, frequently converge to different solutions and are sensitive to the start values used in the estimation routine. While this might not be as much of an issue when estimating demand for one type of product, as is typically the case in IO studies, it is certainly a concern here, given the scale of the application in this paper with several thousand product categories.

In order to make an application of this scale feasible, I heavily benefit from a recent approach developed by Dubé, Fox and Su (2012) who set up the problem as a mathematical program with constraints (MPEC, Su and Judd (2012)). Dubé, Fox and Su (2012) show that MPEC successfully recovers the true parameter values in virtually all synthetic datasets they consider. In addition, it provides substantial reductions in computing times, making the estimation up to 40 times faster in their simulations. I still heavily parallelize the computational load which is possible given that demand for each product category can be estimated separately from each other as well as independently of the supply side. Finally, the use of a C++ based code for BLP increases the speed of the estimation further.[19] While the demand side estimation routine in the baseline specification

---

[19] I make extensive use of the R package Rcpp (Eddelbuettel and Francois (2011)) which allows integration of R and C++. Especially evaluating the Hessian and Jacobian matrices in C++ provides additional speed gains. Further, I

still takes a little more than 10 hours[20], reductions in computing time are absolutely essential in this context.

In practice, as is common in the literature, I define the error term as ξ_{cm}, i.e. the impact of the composit term $\xi_{cm} \equiv q_{cm} + \ln N_{cm}$ on household utility and construct conditional moment conditions,

$$E(\xi_{cm} | \mathbb{Z}_{cm}, \mathbb{X}_{cm}). \tag{27}$$

$\mathbb{Z}_{cm}$ denotes the instrument matrix and $\mathbb{X}_{cm} = \{s_{cm}, p_{cm}, Y_m\}$ summarizes the data consisting of expenditure shares, unit values, and the draws from the income distribution. The GMM estimator in this context can then be written as

$$\theta^* = \arg\min_{\theta} Q(\xi(\theta)) = \arg\min_{\theta} g(\xi(\theta))' W g(\xi(\theta)) \tag{28}$$

with W being a weighting matrix and

$$g(\xi(\theta)) = \frac{1}{M} \sum_{m=1}^{M} \sum_{c=1}^{C} \xi_{cm} \cdot h(\mathbb{Z}_{cm}, \mathbb{X}_{cm}). \tag{29}$$

I follow Nevo (2001) and set the weighting matrix equal to $(\mathbb{Z}'\mathbb{Z})^{-1}$ with $\mathbb{Z} = \{\mathbb{Z}_{11}, ..., \mathbb{Z}_{CM}\}$. Function $h(\cdot)$ includes the actual vectors z and x along with their squared and cubed values.

While the actual solution algorithm closely follows Dubé, Fox and Su (2012), I adjust their code to allow for shifts in the distribution of price coefficients based on observables, which is necessary given the particular way in which income affects demand in the model. In practice, I use 100 income draws from the respective distribution for each market when simulating the right-hand side of equation (24) to approximate the income distribution in each country.[21] I allow the income distributions (25) to vary both by importing country as well as by period to reflect cross-country and time series variation in μ_m and σ_m.[22]

In principle, equation (26) can be estimated by solely using data on prices, expenditure shares and income. However, in order to increase the precision of my estimates, I include a set of producer fixed effects which capture the average values of ξ_{cmk} over all periods.[23] I also follow Khandelwal (2010) and add the logarithm of an exporting country's population as a proxy for hidden vari-

---

completely rely on open source software, such as the optimizer IPOPT (Wächter and Biegler (2006)) which allows me to overcome server limits, such as the number of available MATLAB or KNITRO licenses.

[20] under uninterrupted use of 100 CPUs.

[21] I chose 100 draws mainly to reduce the computational burden of the estimation. Increasing the number of draws for example to 1,000 increases the computing time and memory requirements by an order of magnitude but did not appear to alter the results significantly in several test subsamples.

[22] I also demean the resulting draws for $\ln(y_i)$ by subtracting the average log income draw over all markets. This does not affect the resulting estimates for the price coefficients but facilitates the interpretation of the resulting values for α_1.

[23] The predicted values of $\xi_{cmk} = q_{cmk} + \ln(N_{cmk})$, which are later used on the supply side and for counterfactuals, are then given by $\xi_{cmk} = \bar{\xi}_{cmk} + \Delta\xi_{cmk}$ with $\bar{\xi}_{cmk}$ denoting the fixed effect and $\Delta\xi_{cmk}$ the inferred remaining unobserved heterogeneity. For computational reasons, I did not include a fixed effect for each country but instead for 5 groups for each product where I define groups based on country GDP per capita. Increasing the number of groups had only small effects on the estimates but significantly affected computing times.

eties. Finally, I include bilateral distance to allow for differences in demand for varieties from geographically close or home countries, that are not captured by differences in prices.

Because of the computational challenge that comes with estimating random coefficient models for thousands of product categories, I also restrict the time period covered to the years between 2003 and 2005, i.e. to 12 quarters in total. I select this time period primarily to be reasonably far from the recession in the early 2000s and the Great Recession starting 2007. In practice, the estimation routine can still occasionally struggle to find a solution in cases in which there are many exporters with extremely small expenditure shares. In those cases, I delete observations with an expenditure smaller than 0.0015 and rerun the estimation.

Cobb-Douglas Weights One can easily show that utility maximization implies

$$\omega_k = \frac{\sum_i E_{ik}}{\sum_k \sum_i E_{ik}}, \tag{30}$$

where the numerator equals a country's spending on product category k and the denominator a country's total spending on all goods. Since both are observable in the data, constructing these weights is straightforward.

In practice, I consider three different specifications in terms of utility weights. In the baseline specification, I assume ω to be product-country-specific, i.e. countries are allowed to differ in terms of their preferences for different goods. I also explored two alternative specifications in which I allow ω to be product-time specific and product-time-country specific, respectively, to e.g. account for changing consumption patterns over time. I found that the results across these specifications are very similar. As shown below, the results are also robust to weighting all goods equally.

Appendix D.5.4 provides summary statistics for the Cobb-Douglas weights as well as correlations across countries and time periods for the respective specifications. Consumption across products is fairly evenly distributed with the largest product category having an associated weight of 3.2%. Only 6 categories have values of above 1%. As evident from Figure D5, the weights are also strongly correlated across countries in the specification that allows for country-specific values. For the 4 biggest European economies for example, the correlation ranges from 0.59 to 0.84. As shown in Figure D6, the correlation across time periods is even stronger with values above 0.98 between years.

4.2.2 Step 2: Supply Side and Counterfactuals

Once the demand side parameters in Step 1 are estimated, i.e. the price coefficients α_0 and α_1 as well as $q_{cm} + \ln N_{cm}$, they can be used to back out the supply side parameters of the model using conditions (FE) and (PM). At this stage, the main unknowns are the marginal and fixed cost of each firm, mc_{cm} and f, as well as the extent to which each firm's unit cost varies with quality q_{cm}.

Intuitively, Step 2 uses the results from Step 1 to determine each firm's production technology that rationalizes the observed price and quality choices in the data. As seen in Section 3.3, the profit maximization condition (PM) implies that the optimal price and quality choices of each firm must satisfy the following two first-order conditions:

$$\frac{p_{cm} - mc_{cm}}{mc_{cm}} = \frac{N_{cm} \sum\limits_{i \in I_m} E_{im} s_{cm}^{(i)}}{\sum\limits_{i \in I_m} \alpha_i E_{im} s_{cm}^{(i)} \left(N_{cm} - s_{cm}^{(i)} \right)} \tag{31}$$

$$m_{1,cm} \frac{\ln(mc_{cm})}{q_{cm}} \frac{s_{cm}}{N_{cm}} = \frac{p_{cm} - mc_{cm}}{mc_{cm}} \sum_{i \in I_m} \frac{E_{im}}{E_m} \left[\frac{s_{cm}^{(i)}}{N_{cm}} - \left(\frac{s_{cm}^{(i)}}{N_{cm}} \right)^2 \right]. \tag{32}$$

Together with the free entry condition (FE) these conditions define a system of equations which uniquely pins down the three unknowns mc_{cm}, f, and $m_{1,cm}$.

I solve the above system of equations as follows: I first use equation (31) together with (FE) to back out the fixed cost of exporting. For that purpose, I use data on the number of French firms in each industry and solve for the vector of costs ($\{mc_{cm}, f\}$) at which equation (31) is satisfied and the predicted number of French firms selling in France in the model matches that in the data.[24] Once the fixed cost of exporting is known, equations (31) and (32) can then be used to back out each firm's marginal cost and the degree to which it varies with quality, $m_{1,cm}$.[25]

In order to quantify the gains from trade, I then evaluate the aggregate, household-specific price indexes derived in Section 3,

$$\mathbb{P}_m(y_i) = \exp\left(\sum_k \hat{\omega}_k \left[1 - \ln \frac{\hat{\omega}_k}{p_{ikm}^*} - \frac{q_{ikm}^* + \varepsilon_{ik}}{\hat{\alpha_0} + \hat{\alpha_1} \ln(y_i))} \right] \right), \tag{33}$$

and compute their change when moving to autarky. More specifically, I simulate the gains from trade by comparing the current equilibrium to a counterfactual scenario in which all firms are prohibited from exporting to any market m other than their home country. In practice, I exclude these countries from the choice set of consumers which is equivalent to imposing prohibitively high trade costs.

The variables which are then being reoptimized are domestic prices p_{cm}, qualities q_{cm}, and the number of firms N_{cm}. I use equation (20) to compute the optimal quality in autarky which requires only the estimates on the distibution of the price coefficients $\{\hat{\alpha}_i\}_m$ in each country as well as the production parameters $\hat{m}_{0,cm}$ and $\hat{m}_{1,cm}$. The markup equation (31) and the free entry condition (FE) can then be used to back out the new p_{cm} and N_{cm}.

As described above, y_i and ε_{ik} are random variables and follow the distributions described above. To approximate $\mathbb{P}_m(y_i)$ numerically, I use the same 100 draws from the income distribution that were used for the demand estimation and evaluate ε_{ik} using $1,000$ draws from a type 1 extreme value distribution for each consumer and variety. I hold these draws constant before and after the

---

[24] I assume that each French firm also sells in France, which is consistent with the observation in Eaton, Kortum, and Kramarz (2011) that only a small share of less than 0.3% of French firms does not sell domestically.

[25] $m_{0,cm}$ can then also be inferred from the functional form for mc_{cm} and must equal $m_{0,cm} = \ln(mc_{cm})/q_{cm}^{m_{1,cm}}$.

policy change.

Finally, in order to summarize how unequal the gains from trade are across the income distribution, I follow Fajgelbaum and Khandelwal (2016) and compute the percentage change in $\mathbb{P}_m(y_i)$ when moving to autarky for a household at the 20th percentile of the income distribution and compare it to the change for a household at the 80th percentile. When constructing the respective price indexes, I use country-period specific expenditure weights ω_k, which are inferred from the data as described in section 4.2.1.

5 Results

The results are presented in three steps: Section 5.1 summarizes the parameter estimates on both demand and supply side. I then focus on an example category to provide a sense of the data, the estimation procedure, and especially the intuition behind the moving parts of the model. Given that the optimal price and quality choices of firms do not have closed form solutions, this example also sheds light on the main mechanisms driving vertical differentiation in the model and its welfare implications.[26] Section 5.3 then covers the counterfactual results for the full sample.

5.1 Parameter Estimates

Table 1 summarizes the estimates of the demand parameters for the full sample along with predicted demand-side price elasticities. In most categories, the estimates are consistent with the initial priors: α_0 is positive in 88.6% of the cases, which implies a negative price elasticity for the majority of product categories. I also find that α_1, which governs to which degree households differ in terms of their effective demand for quality relative to price, is negative for 78% product categories. Hence, as one would expect, higher-income consumers behave less price-elastically compared to lower-income households.

The top part of Table 1 plots the distribution of the own price elasticity (left-hand side) as well as the elasticity of substitution (right-hand side) over all product categories based on the parameter estimates. The own price elasticity captures how price elastic the consumers of a product are and can be computed as

$$\epsilon_{jm} = \frac{\partial x_{jm}}{\partial p_{jm}} \frac{p_{jm}}{x_{jm}} = -\left(1 + \sum_{i \in I_m} \frac{E_{im}}{E_m} \frac{\alpha_i s_{jm}^{(i)}(1 - s_{jm}^{(i)})}{s_{jm}}\right)$$

for firm j in market m.[27] Notice that these elasticities vary by variety as well as by market through differences in the price coefficients α_i and consumer-specific market shares $s_{jm}^{(i)}$. Intuitively, ϵ_{jm} depends on a weighted average of the price coefficients with the weights being determined by the shares $s_{jm}^{(i)}$. For example, since lower-quality producers tend to sell more to poorer consumers with

---

[26] Equation (20) implicitly characterizes the optimal quality choice but is not a closed-form solution since the effective price elasticity $\boldsymbol{\alpha}_{cm}$ also depends on q_{cm}^*.

[27] See Appendix C.5.

larger α_i's, these producers will typically also face larger own-price elasticities on aggregate.[28] For the same reasons, in contrast for example to a CES framework, also the elasticity of substitution is variety- and market-specific.[29]

Table 1: Price Coefficients and Median Elasticities

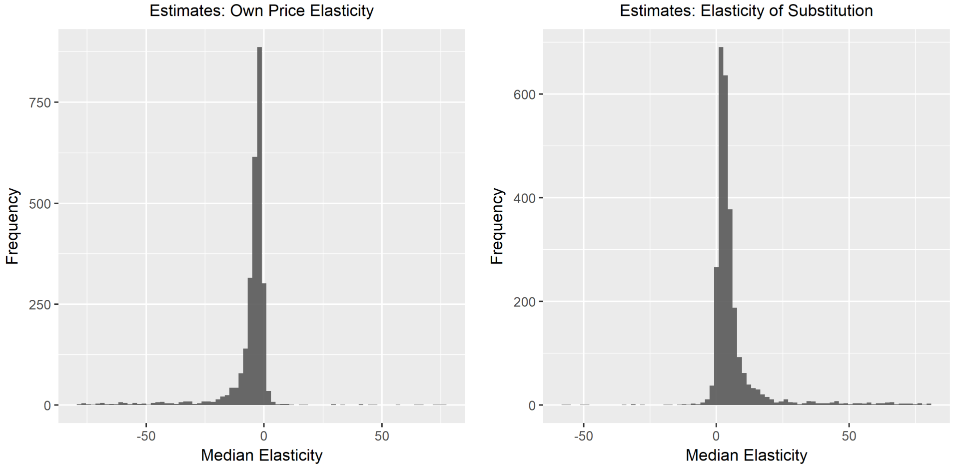

Estimates: Price Coefficients:

α_0 (Share positive)	88.61%
α_1 (Share negative)	77.99%

Estimates: Elasticities:

	Own Price Elasticity	Elasticity of Substitution
10th Percentile	-0.81	0.72
25th Percentile	-1.80	1.92
Median	-3.61	3.80
75th Percentile	-7.57	7.92
90th Percentile	-83.66	87.39

The left figure plots the median own price elasticities over all producing countries for each product category. The figure on the right hand side plots the corresponding elasticities of substitution. Elasticities with an absolute value above or below 100 are excluded.

The bottom part of Table 1 reports summary statistics of the estimates of both own-price elasticities and the elasticities of substitution. Since these estimates are producer- and market-specific, I report the median over all elasticities within a product category as a measure of ϵ_{jm} for each product category. The median own-price elasticity is -3.61 which is consistent with previous empirical applications of nested logit frameworks. Handbury (2013) for example estimates a median own-price elasticity of -2.09 for a wide range of groceries sold in the United States. Berry, Levinsohn and Pakes (1995) find slightly higher price elasticities for cars, ranging from -3.0 to -6.5 depending on the respective car brand. Nevo (2001) estimates elasticities for ready-to-eat

---

[28] As long as market shares are below 50% which is the case for the majority of observations.

[29] The elasticities of substitution presented in Figure 1 are computed numerically. Specifically, to obtain an estimate for a particular producer j and market m, I increase the prices of all other varieties in m by a small percentage, leaving j's price constant, and calculate the associated change in the relative share of j.

cereal brands between -2.3 to -4.3.

The estimates of the elasticity of substitution follow those for the own-price elasticities very closely. For the median product, the implied elasticity equals 3.80, which is consistent with the CES-based estimates by Broda and Weinstein (2006) for U.S. trade data who estimate slightly lower elasticities between 2.2 and 3.7 depending on the level of disaggregation. A plausible explanation for this difference might be that the U.S. is richer than the majority of Western European countries and significantly richer than Eastern European ones. Therefore one would expect U.S. households to be somewhat less price-elastic than their European counterparts. The median elasticity for the 5 richest EU members, which are closer to the U.S. in terms of income per capita equals for example 3.19 and falls well within the range estimated by Broda and Weinstein (2006).[30]

Appendix Figure D4 highlights the importance to instrument for prices: For the majority of product categories, the estimated price coefficients are larger in absolute value when using 2SLS compared to OLS. This was the case in 78% of product categories when using the Hausman (1996)-type instruments and in 69% of cases when using shipping costs. As evident from panels a) and b) of Figure D4, the distribution of the difference between 2SLS and OLS coefficients, $\alpha_0^{(2SLS)} - \alpha_0^{(OLS)}$ has more mass to the right of 0, implying more negative price coefficients in the 2SLS case. Panels c) and d) show further that the majority of price coefficients is negative in both cases, with the main difference being that the OLS estimates are much more concentrated between -5 and 0.

While the estimates imply the price coefficient α_i to be positive in the majority of cases, there are still certain categories for which this does not hold. This is not unusual when estimating demand for a large number of product categories and less of a concern when the characteristics of varieties are treated as exogenous as in Broda and Weinstein (2006), Khandelwal (2010) or Handbury (2013). The addition of a supply-side however does require demand to be downward sloping and sufficiently elastic in order to deliver interior solutions. More specifically, the market- and product-specific elasticity of substitution needs to be greater than 1.

In the cases in which this condition is not satisfied I therefore leave the equilibrium values of p_{cm} and q_{cm} unchanged in the counterfactual while still allowing for entry. Alternatively, I also deleted those observations with little change in the overall results.[31] Especially small positive values of the price elasticity, by entering in the denominator, create the additional numerical issue of occasionally generating very large values for the price index of certain categories k. I therefore also removed observations in which the price coefficient of group i exceeds 0.5.

Table 2 reports summary statistics of the estimates for product quality for each European exporter. Specifically, these values reflect the coefficients of a regression of the quality estimates on producer dummies, controlling for product-importer-period fixed effects. I find that the highest

---

[30] These estimates might also simply differ since Broda and Weinstein (2006) estimate the elasticity of substitution in a CES framework. Since the discrete-choice framework of this paper nests the CES however (see Anderson, De Palma and Thisse (1987)), one would still expect a certain degree of overlap in the respective estimates despite the methodological difference.

[31] I also experimented with using the same strategy as Broda and Weinstein (2006) and restrict the coefficients α_0 and α_1 in the estimation such that the resulting elasticity of substitution equals 1. In practice, this procedure did always deliver a solution in which the constraint was binding, i.e. in which the elasticity is exactly 1.

Table 2: Quality Estimates

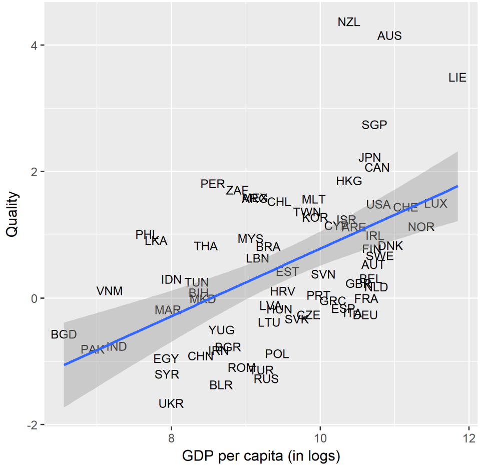

Highest and Lowest Quality Producers (World):

High		Low	
New Zealand	4.375	Romania	-1.095
Australia	4.161	Turkey	-1.132
Liechtenstein	3.500	Syria	-1.201
Singapore	2.747	Russia	-1.270
Japan	2.230	Belarus	-1.369
Canada	2.075	Ukraine	-1.662

Highest and Lowest Quality Producers (EU only):

High		Low	
Luxembourg	1.507	Czech Republic	-0.261
Ireland	0.991	Slovakia	-0.326
Denmark	0.837	Lithuania	-0.376
Finland	0.785	Bulgaria	-0.767
Sweden	0.675	Poland	-0.874
Austria	0.538	Romania	-1.095

All values represent the coefficients of a regression of the quality estimates on producer dummies as well as product-importer-period fixed effects. All values are relative to French product quality, which has been normalized to 0. Countries with less than 5,000 observations are excluded.

quality-producers are New Zealand and Australia, while Belarus and Ukraine are at the bottom end of the quality distribution. Within the European Union, Luxembourg, Ireland, and the Scandinavian economies produce higher-quality varieties while Eastern European countries tend to specialize in cheaper, lower-quality versions of products.

As evident from Table 2, the quality estimates are highly correlated with GDP per capita of the respective country, implying that richer economies have a comparative advantage in producing higher-quality goods. This result is in line with findings by Khandelwal (2010), Hallak and Schott (2011), and Feenstra and Romalis (2014). The estimates are also positively correlated with prices and the inferred marginal cost: The cross-country correlation between average qualities and average prices is 0.84. The correlation between quality and the marginal cost is 0.86, implying that higher quality-varieties are more costly to produce, which translates into higher prices.

Table 3: Cost Estimates

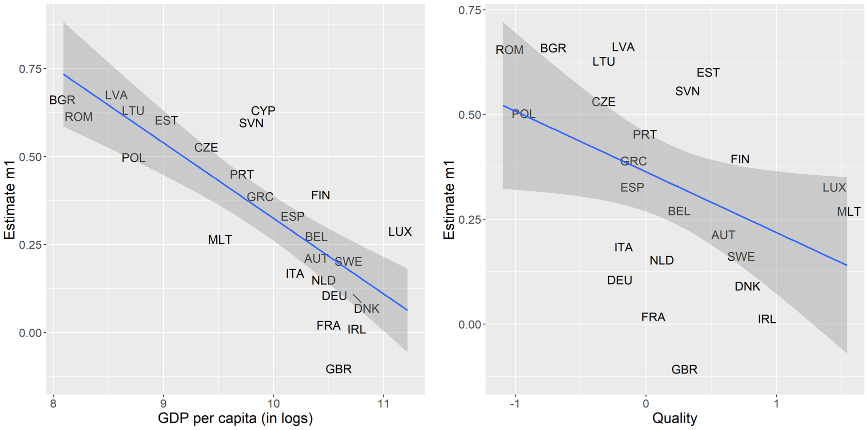

Marginal cost (20 largest economies):

Switzerland	5.248	Netherlands	-0.804
Japan	5.123	Germany	-0.867
United States	3.798	Spain	-1.396
Australia	3.643	Korea	-1.698
Canada	3.542	Italy	-1.711
Mexico	0.572	Brazil	-1.788
United Kingdom	0.407	Russia	-2.025
Saudi Arabia	0.109	Indonesia	-2.506
France	0.000	India	-3.072
Belgium	-0.115	China	-5.075

The top figures plot the estimates of the cost parameter m_1 against GDP per capita (in logs, left) and the average product quality (right) of the respective countries. All m_1 values in the plot represent the coefficients of a regression of the m_1 estimates on producer dummies as well as product-importer-period fixed effects. The numbers in the bottom table report the inferred marginal cost for the 20 largest countries in terms of GDP. All values are relative to those for France, which have been normalized to 0.

Table 3 summarizes the cost estimates for the 20 largest economies. Also these values reflect the coefficients of a regression of the respective cost estimates on producer dummies, controlling for product-importer-period fixed effects. The values in the bottom table show that the marginal cost estimates are largely in line with the quality estimates from Table 2: Australia and Japan were for example among the highest-quality producers and are also among the highest cost producers. China, India, and Indonesia have the lowest inferred unit costs among large countries.

As described previously, the marginal costs are largely driven by three factors: (1) The produced quality level, (2) the overall productivity of a country in a product category, and (3) trade costs. The importance of the latter is evident by the fact that 5 out of the 6 most expensive exporters to Europe are from outside the continent. Canada, Australia and Japan, for example, are additionally among the highest-quality producers in the sample, which further contributes to their rank among the highest-cost producers.

Also within Europe, the cost pattern mirrors the pattern seen in the quality results: Ireland and Luxembourg were for example the highest-quality producers and are also among the highest cost producers. Eastern European countries on the other hand produce their varieties at significantly lower unit costs. While the highest quality producers tend to be also the most expensive ones, there are several notable exceptions: The UK for instance ranks 12th in quality but has the 4th highest cost among EU producers. Slovenia on the other hand ranks 10th in quality out of 27 countries but has the 5th lowest unit cost.

The top left plot in Table 3 highlights how m_1, which governs the marginal cost of increasing quality, varies by country. There is a clear negative correlation between this parameter and a country's GDP per capita, which implies that richer economies have a cost advantage in producing higher-quality varieties. As shown in the plot on the right hand side, higher values for m_1 also translate into smaller actual quality choices made by firms in the respective countries.

5.2 An Example Category

In order to demonstrate the main features and implications of the model and the counterfactual analysis, I begin with an example category before generalizing in the next section. My example category is *Floor-cloths; dish-cloths; dusters and similar cleaning cloths; knitted or crocheted; life-jackets; life-belts and other made up articles.*[32] I chose this category mainly because it is a representative example for the main results on prices, quality and welfare, which hold for the majority of product categories.

Table 4 shows summary statistics for the example category. The 10 biggest suppliers cover over 87% of the European market, with the UK, Germany, and Hungary being the largest producers in terms of expenditure shares. The average price differs significantly across producers, with India or China selling at less than a quarter of the price charged by the most expensive EU countries, the UK and Denmark. The pattern that higher-wage countries sell at higher prices is common across the majority of product categories, which also makes this example category representative in this regard.

It is also evident that price differences alone cannot fully explain the variation in expenditure shares: The U.S. for example has a higher share than India, despite selling at higher prices and being roughly equally close to most European markets. As suggested in column (3), the model rationalizes this with differences in quality: In this example category, the U.S. produces goods of higher quality than Indian firms do. More generally, richer economies tend to produce higher quality varieties than poorer economies do: The richest EU economies Ireland, Denmark and Sweden are among the highest quality producers while Romania, Bulgaria, China and India are located at the other end of the spectrum. There are several notable exceptions however: Latvia or Slovenia which are relatively poor EU members are among the highest-quality producers which shows that

---

[32] Eurostat classifies this under the prodcom category 17402590.

Table 4: Summary Statistics: Example Category

Country	(1) Avg. Price	(2) EU Exp. Share	(3) Avg. Quality	(4) $\hat{m}_1$	(5) GDP per capita	(6) $\alpha(y_{(20)})$	(7) $\alpha(y_{(80)})$	(8) Domestic Share
EU countries:								
France	21.31	4.89	8.77	1.21	41.85	-2.66	-2.17	51.91
Netherlands	9.61	3.04	7.78	1.46	47.84	-2.68	-2.09	62.97
Germany	13.36	16.20	7.53	1.21	40.46	-2.62	-2.10	74.85
Italy	13.20	1.66	7.68	1.97	36.03	-3.00	-2.17	37.42
UK	27.93	26.20	9.18	1.28	46.87	-2.73	-2.02	89.54
Ireland	21.46	0.61	10.05	1.84	59.41	-2.65	-2.02	67.99
Denmark	31.75	0.52	10.77	1.38	57.17	-2.48	-2.08	47.29
Portugal	24.92	10.69	9.36	1.41	21.88	-3.36	-2.70	95.37
Spain	10.81	4.04	7.32	1.30	32.17	-3.00	-2.37	71.26
Belgium	20.21	1.76	9.92	1.39	43.49	-2.64	-2.22	63.82
Sweden	18.59	2.89	9.91	1.44	50.37	-2.53	-2.12	74.64
Finland	15.72	2.25	9.30	1.42	46.50	-2.60	-2.08	80.29
Austria	26.50	1.55	10.00	1.40	45.25	-2.60	-2.09	61.52
Estonia	12.41	0.22	10.17	1.54	16.40	-4.08	-3.27	77.69
Latvia	18.69	0.21	10.43	1.51	12.97	-4.19	-3.59	86.07
Czech Republic	10.11	0.52	7.76	1.80	17.54	-3.60	-3.08	27.45
Hungary	25.62	11.23	10.14	1.42	13.52	-3.71	-3.24	97.45
Romania	15.66	0.39	8.17	3.69	7.92	-4.47	-3.77	6.23
Bulgaria	8.30	0.05	7.50	1.89	5.52	-4.34	-3.69	31.39
Slovenia	19.32	0.25	11.28	1.55	23.59	-3.11	-2.61	63.75
Largest Non-EU exporters:								
China	6.49	6.12	5.07		2.64			
United States	28.94	0.67	9.45		47.96			
India	8.46	0.41	5.47		1.08			

This table shows summary statistics for the example category during the most current period used for the estimation (4th quarter of 2005).

the relationship between q_{cm} and a country's GDP per capita is not imposed by the model or mechanically generated by the estimation, but rather a product of each country's combination of observed market shares and prices.

Column (4) of Table 4 summarizes the estimates of the cost parameter m_1. In this example category, France, Germany, and the UK have the lowest marginal cost of producing quality while Romania, Italy and Bulgaria have a disadvantage in producing higher quality, which in turn results in, on average, lower values of q_{cm} for those countries.

The final two columns demonstrate the extent to which demand is non-homothetic. $\alpha(y_{(20)})$ denotes the price coefficient of a household at the 20th percentile of the income distribution of the respective country, while $\alpha(y_{(80)})$ is the equivalent for the 80th percentile. These numbers are a direct product of the estimated demand parameters $\alpha_0 = 2.76$ and $\alpha_1 = 0.69$ in this product category and imply that households across the income distribution behave differently price elastic and in turn will differ in terms of their preferred varieties. There is also significant variation in these price elasticities across countries, with households in the poorer EU countries being almost twice as price elastic as consumers in richer economies. This cross-country variation is due to differences in the location and variance of the income distribution across EU member states (see Appendix Table D4).

Table 5: Counterfactual Price and Quality Adjustments: Example Category

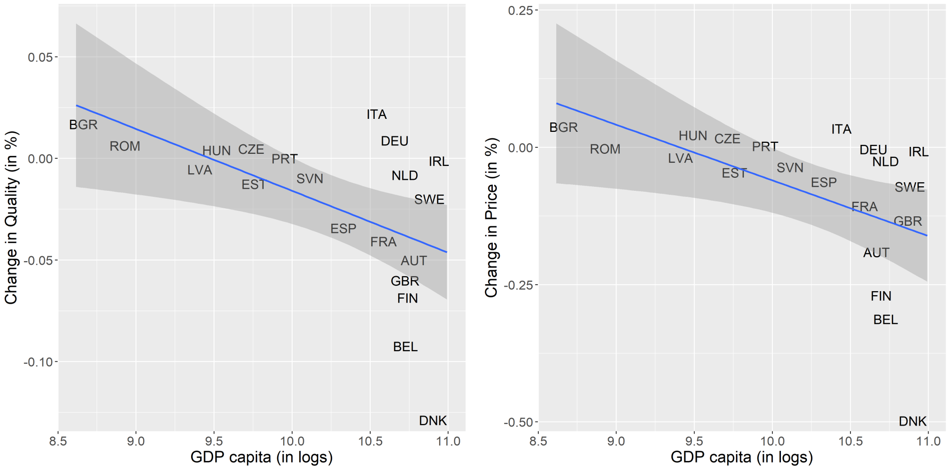

	Full Model						Repr. HH	
Country	q_{old}	q_{new}	Δq (in %)	p_{old}	p_{new}	Δp (in %)	Δq	Δp
France	6.95	6.70	-3.61	15.21	13.68	-10.06	0.00	0.00
Netherlands	6.19	6.15	-0.70	8.89	8.72	-1.94	0.00	0.00
Germany	4.59	4.62	0.71	6.98	7.08	1.42	0.00	0.00
Italy	1.98	2.01	1.30	2.08	2.11	1.40	0.00	0.00
UK	6.51	6.21	-4.51	13.38	11.72	-12.37	0.00	0.00
Ireland	6.52	6.51	-0.07	6.71	6.70	-0.21	0.00	0.00
Denmark	11.27	9.88	-12.32	71.39	36.95	-48.25	0.00	0.00
Portugal	10.25	10.19	-0.62	15.22	14.89	-2.17	0.00	0.00
Spain	7.00	6.84	-2.28	10.40	9.78	-5.96	0.00	0.00
Belgium	9.63	8.80	-8.66	28.21	19.64	-30.39	0.00	0.00
Sweden	8.35	8.23	-1.45	17.31	16.40	-5.21	0.00	0.00
Finland	10.44	9.84	-5.77	32.30	24.91	-22.88	0.00	0.00
Austria	9.55	9.22	-3.42	27.11	23.46	-13.45	0.00	0.00
Estonia	12.60	12.42	-1.45	12.69	12.04	-5.11	0.00	0.00
Latvia	12.85	12.80	-0.38	11.87	11.71	-1.32	0.00	0.00
Czech Republic	5.63	5.70	1.19	3.24	3.31	2.22	0.00	0.00
Hungary	8.99	8.99	0.01	7.89	7.89	0.02	0.00	0.00
Romania	2.94	2.97	1.15	1.49	1.51	1.40	0.00	0.00
Bulgaria	6.08	6.15	1.18	2.86	2.91	2.07	0.00	0.00
Slovenia	10.49	10.40	-0.81	14.37	13.94	-3.05	0.00	0.00

This table summarizes the optimal price and quality choices when moving from the current equilibrium to autarky. Each value refers to the choice made by domestic companies in each country. The *representative household* model removes income inequality by counterfactually setting each household's income equal to the country average. Δq and Δp in the 2 rightmost columns show how domestic firms would adjust p_{cm} and q_{cm} in that model.

Tables 5 and 6 summarize the counterfactual price and quality adjustments associated with a move to autarky. In order to highlight the importance of product differentiation, I first hold the productivity of new entrants on the same level as that of incumbent domestic firms (i.e. I assume that m_0 and m_1 are country-specific). As shown in Table 5, in this case, for the majority of countries, the estimates predict that domestic producers lower both quality and price in the no trade-case relative to the current equilibrium, with an average reduction in quality of 2.03% and a 7.69% decline in prices.

These adjustments are a product of both vertical differences between domestic and imported varieties and consumer heterogeneity: In this example category, trade particularly increases the availability of cheaper products in most EU economies through greater access to cheaper, less sophisticated varieties from non-EU countries (e.g. China and India) as well as EU producers (e.g. Hungary and Germany). Domestic companies hence become disproportionately less attractive to poorer households which lowers the average price elasticity of their remaining consumers. Since the optimal quality choice of firms depends negatively on the price elasticity of a firm's consumers through equation (20), companies have incentives to specialize in higher-quality varieties when selling to on average richer households.

A move to autarky on the other hand reverses this specialization. Since the majority of EU firms specialized in higher-quality varieties and richer households have a greater demand for those products, trade resulted in EU firms having a relatively high market share among richer households. The effective price elasticity is hence lower with trade than in autarky, when demand from all consumers has to be met by domestic firms (in which case $\boldsymbol{\alpha}_{c^*m}$ is just a simple average):

$$\boldsymbol{\alpha}_{cm}^{\text{Trade}} = \frac{\sum_{i \in I_m} \alpha_i [s_{c^*m}^{(i)} - (s_{c^*m}^{(i)})^2]}{\sum_{i \in I_m} s_{c^*m}^{(i)} - (s_{c^*m}^{(i)})^2} \quad < \quad \boldsymbol{\alpha}_{cm}^{Autarky} = \frac{1}{|I_m|} \sum_{i \in I_m} \alpha_i.$$

Because of the direct mapping between the average price elasticity of a firm's consumers and the optimal q_{jm} in equation (20), it follows immediately that EU firms will want to partially shift production to cheaper, lower-quality varieties when moving to autarky.

In addition to the average supply side response, Table 5 also highlights significant cross-country variation in terms of how domestic firms respond to trade barriers. As evident from the two top plots, the declines in the average q_{cm} and p_{cm} are largely driven by richer high-quality producers. Since those firms are disproportionately attractive to high-income households, the effective price elasticity of their consumers will be low. The difference between $\boldsymbol{\alpha}_{cm}^{\text{Trade}}$ and $\boldsymbol{\alpha}_{cm}^{Autarky}$ is therefore particularly large for those countries which results in stronger quality-downgrading. This is the reason why domestic firms in Denmark, Belgium, or Finland are predicted to lower quality by between 6% and 12% and prices by one quarter to one half.

The results also show that not all countries will optimally offer lower-quality varieties in autarky. The low-quality producers Romania, Italy, or Bulgaria for example will want to moderately increase both q_{cm} and p_{cm}: Varieties from these countries are more attractive to lower-income households and so $s_{c^*m}^{(i)}$ is comparably high for those consumer groups. As a consequence the effective price elasticity of these countries' consumers is relatively high and actually exceeds $\boldsymbol{\alpha}_{cm}^{Autarky}$.

It is important to note that a model without consumer heterogeneity would not be able to generate these results. As shown in the two rightmost columns of Table 5, in the case of a representative household, the optimal price and quality choices are independent of trade. Since the consumers of each firm are identical in that case, the effective price elasticity faced by firms is unaffected by the move to autarky, and so will be the optimal q_{cm} and p_{cm}.[33]

Table 6: Price and Quality Adjustments: Variable Productivity

Country	q_{old}	q_{new}	Δq (in %)	p_{old}	p_{new}	Δp (in %)
France	6.95	4.51	-35.07	15.21	6.52	-57.13
Netherlands	6.19	5.18	-16.35	8.89	6.56	-26.23
Germany	4.59	3.81	-17.00	6.98	5.34	-23.52
UK	6.51	5.59	-14.08	13.38	9.50	-29.00
Portugal	10.25	9.98	-2.69	15.22	14.16	-6.96
Spain	7.00	5.64	-19.44	10.40	6.93	-33.38
Belgium	9.63	7.13	-25.96	28.21	11.94	-57.70
Sweden	8.35	7.33	-12.21	17.31	12.56	-27.40
Finland	10.44	8.57	-17.91	32.30	17.21	-46.71
Hungary	8.99	8.91	-0.96	7.89	7.75	-1.72

The above values summarize the optimal price and quality choices when moving from the current equilibrium to autarky, using the estimates of γ_k to infer the difference in productivity between incumbent and entering domestic firms. Each value refers to the choice made by domestic companies in each country. For brevity, only the 10 largest exporters in terms of EU expenditure share are included in the table.

The results in Table 6 additionally allow for differences in productivity between incumbent and entering domestic firms which is important given the direct relationship between a firm's quality choice and the productivity parameters m_0 and m_1 predicted by the model (see equation (20)). In order to quantify the counterfactual productivity of entrants, I follow Feenstra and Romalis (2014) and Chaney (2008) who use industry-specific estimates on $\zeta^k = \gamma/(\sigma - 1)$ to distinguish the shape parameter of the Pareto distribution γ from the elasticity of substitution σ. I use my estimates on the average elasticities of substitution by industry k to back out γ for each product category.[34] These parameter values are then used to infer how the average productivity of domestic firms changes when moving to autarky.[35]

---

[33] This unresponsiveness of prices also highlights the connection between log-logit and CES preferences (See Anderson, De Palma, and Thisse (1987)). With CES preferences and a representative household as e.g. in a Melitz (2003) model, each firm's prices are a direct product of the marginal costs and the constant elasticity of substitution. Adding consumer heterogeneity in combination with vertical differentiation make both endogenous and hence both will be affected by trade. If either is removed however, the price setting decision of firms collapses to that of standard CES-based trade models.

[34] Since the estimates of ζ^k are based on the 3-digit SITC classification, I crosswalk them to HS6 product categories first and then match the resulting values to the prodcom categories used in this paper. I use the same ζ for all pc8 categories that belong to a particular SITC category.

[35] More specifically, in order to quantify the change in productivity of domestic firms, I first compute the additional number of firms that enter in the fixed-productivity model. I then use the estimated γ_k's to infer the unit cost mc_{cm} of these entrants as well as the change in the average unit cost mc_{cm} of domestic firms. In principle, one may rationalize the lower productivity of entering firms with either a higher m_0 or m_1. I assume that entering firms have a lower baseline productivity m_0 and that the marginal cost of producing quality is the same for all domestic firms. Notice that in both cases, less productive firms will sort into producing less sophisticated varieties, which is consistent

As evident from Table 6, allowing for productivity amplifies the reductions in quality and prices seen in Table 5. Since m_1 exceeds 1 for each producer, there is a direct negative relationship between the optimal quality choice q_{cm} and m_0 through equation (20), which in this example category results in sizable quality downgrading and reductions in prices. The magnitude of these changes varies significantly with each country's domestic share: Portugal and Hungary for example are both comparably large producers in the EU and have substantial domestic shares as well (97 and 95%). There would hence be only limited entry following a move to autarky and in turn only a moderate change in the average productivity. The opposite holds for Belgium and France who are comparably small suppliers in this specific industry.

The supply side responses described in Tables 5 and 6 are particularly important for the extent to which the gains from trade differ by consumer and across the income distribution. To demonstrate this point, Table 7 summarizes the impact on consumer welfare in this specific product category when EU countries counterfactually shift to autarky. In order to compute the gains from trade, I use the ideal price index implied by the model (see equation (23)), which in the case of one good simplifies to

$$\mathbb{P}_m^{(k)}(y_i) = \exp\left(1 + \ln(p_{ikm}^*) - \frac{q_{ikm}^* + \varepsilon_{ik}}{\alpha(y_i)}\right), \tag{34}$$

and compute it with and without trade.

As described earlier, the move to autarky affects $\mathbb{P}_m^{(k)}(y_i)$ in two ways: First, it forces households to buy domestic varieties which harms those households who preferred the price-quality combination of foreign varieties or had a strong idiosyncratic taste for them. Second, the price and quality adjustments seen above change q_{km} and p_{km} of the domestic varieties in autarky. Since poor and rich household differ in their valuation for quality through $\alpha(y_i)$, both adjustments will also affect households across the income distribution differently.

Table 7 summarizes the predicted impact of moving to autarky on household welfare. In order to highlight the importance of vertical differentiation, I compute the change in $\mathbb{P}_m^{(k)}(y_i)$ both with and without supply side adjustments. The first three columns show the average welfare consequences for all countries on (1) households at the 20th percentile of the income distribution ($\Delta\mathbb{P}_m^{(k)}(y_{(20)})$), (2) households at the 80th percentile ($\Delta\mathbb{P}_m^{(k)}(y_{(80)})$), and (3) the gap between the two $\Delta\mathbb{P}_m^{(k)}(y_{(20)}) - \Delta\mathbb{P}_m^{(k)}(y_{(80)})$, when domestic producers are solely allowed to enter in response to the removal of foreign competitors. Columns (4) to (6) present the gains from trade when firms additionally adjust the characteristics of their varieties as seen in Table 5. Finally, the rightmost columns document the analogous results for the variable-productivity model described above.

---

with empirical evidence (see e.g. Faber and Fally (2017)).

Table 7: Gains from Trade: Example Category

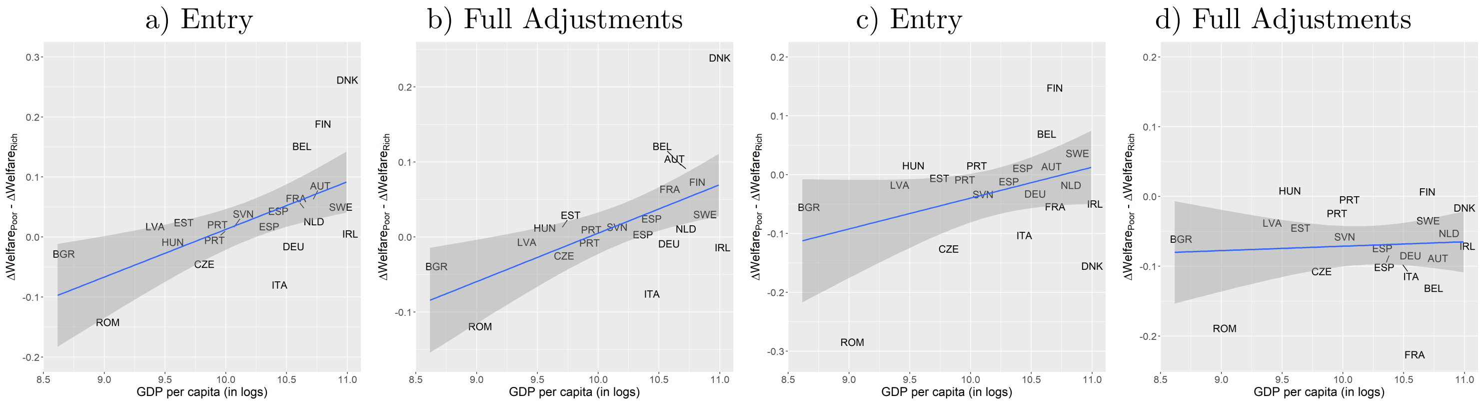
a) Entry b) Full Adjustments c) Entry d) Full Adjustments

Country	Fixed Productivity						Variable Productivity					
	Entry only			Full Adjustments			Entry only			Full Adjustments		
	$\Delta P_{(20)}$	$\Delta P_{(80)}$	Gap	$\Delta P_{(20)}$	$\Delta P_{(80)}$	Gap	$\Delta P_{(20)}$	$\Delta P_{(80)}$	Gap	$\Delta P_{(20)}$	$\Delta P_{(80)}$	Gap
Austria	2.04	-4.28	6.32	4.30	-4.81	9.12	35.87	35.59	0.28	13.23	22.99	-9.77
Belgium	6.44	-9.78	16.21	3.87	-5.70	9.57	59.96	51.90	8.06	15.49	27.41	-11.91
Bulgaria	-1.95	1.96	-3.90	-2.07	2.75	-4.82	10.18	16.86	-6.68	5.56	12.61	-7.04
Czech Republic	-1.93	3.65	-5.59	-0.62	1.03	-1.66	30.50	44.19	-13.69	11.22	22.90	-11.68
Denmark	14.45	-10.65	25.11	12.87	-10.18	23.05	441.14	457.75	-16.61	29.42	31.95	-2.54
Estonia	1.75	-1.76	3.51	0.53	-0.78	1.31	18.50	17.97	0.53	10.11	13.66	-3.55
Finland	8.23	-9.55	17.78	4.61	-3.77	8.38	38.53	22.63	15.90	17.37	15.71	1.66
France	2.39	-2.44	4.83	2.24	-3.23	5.47	38.77	42.99	-4.22	7.95	29.68	-21.73
Germany	-0.89	-0.40	-0.49	0.51	0.42	0.09	6.09	8.23	-2.14	2.23	9.73	-7.49
Hungary	0.19	-0.04	0.23	-0.01	-0.38	0.37	0.90	0.77	0.13	0.77	0.87	-0.10
Ireland	-0.29	0.27	-0.56	-1.60	0.68	-2.27	10.26	14.10	-3.85	6.59	12.61	-6.03
Italy	-3.13	3.76	-6.89	-2.66	4.07	-6.72	2.74	11.95	-9.21	0.99	10.78	-9.79
Latvia	0.19	-0.48	0.66	-0.02	-0.24	0.22	5.63	6.24	-0.60	3.58	6.43	-2.85
Netherlands	1.86	-1.80	3.66	1.33	-0.64	1.97	22.38	23.04	-0.66	11.62	17.81	-6.19
Portugal	0.26	-0.68	0.94	0.28	0.16	0.12	2.27	1.89	0.38	1.34	2.82	-1.48
Romania	-5.96	9.32	-15.28	-6.45	6.33	-12.78	26.39	55.99	-29.60	3.52	23.35	-19.83
Slovenia	0.59	-1.10	1.69	0.73	0.43	0.30	25.14	26.99	-1.85	12.72	17.52	-4.80
Spain	1.28	-1.94	3.22	0.44	-1.03	1.47	16.50	16.56	-0.07	7.06	15.54	-8.48
Sweden	2.49	-2.69	5.19	0.46	-3.41	3.87	19.19	16.66	2.53	11.04	14.29	-3.26
UK	2.43	-3.55	5.98	0.98	-1.15	2.14	9.40	5.49	3.91	3.40	10.52	-7.12
Average			**3.13**			**1.96**			**-2.87**			**-7.20**

This table summarizes the percentage change in the price index as defined by (34) when moving from the current equilibrium to autarky in each country. $\Delta P_{(20)}$ denotes the percentage change in the price index for a household at the 20th percentile of the income distribution while $\Delta P_{(80)}$ denotes the same statistic for the 80th percentile. *Gap* is defined as $\Delta P_{(20)} - \Delta P_{(80)}$. The *Entry only* model keeps each firm's quality and price fixed when moving to autarky, while the full model allows for adjustments on both margins.

There are four main takeaways from Table 7: First, when productivity is held fixed, the rich-poor gap is positive for the majority of countries which means that poorer households tend to benefit more from trade in this product category than richer ones do in the European Union. Depending on the model, trade results in an on average 1.96 to 3.13 percentage points stronger decline in the price index for poorer households. This result is due to most EU countries specializing in higher-quality varieties and hence imports being more attractive to poorer than to wealthier households.

Second, there is considerable heterogeneity in terms of how unequal the gains from trade are in each country. Poorer households gain between 8 and 23 percentage points more in richer EU economies such as Denmark, Finland or Austria. In Romania and Bulgaria on the other hand, especially richer households benefit from trade. As shown in the two top figures, there is a noticeable positive correlation in both models between a country's GDP per capita and the extent to which trade is pro poor: The rich-poor gaps are larger in richer economies such the Scandinavian countries but quite modest or negative in Eastern Europe. Also these findings are a direct product of the productivity and quality estimates shown in Table 4 as well as the extent to which demand is non-homothetic. Since richer countries tend to have a comparative advantage in producing higher quality-goods, but at a high price, poor households in these countries particularly benefit from access to cheaper foreign varieties through trade.

Third, accounting for quality specialization has a significant impact on the results. In the entry-only model, the average gap equals about 3 percentage points and can take larger values of up to 25 percentage points in Denmark. In comparison, when accounting for quality and price adjustments, this gap shrinks to just about 2 percentage points on average. The reason for this result lies in the direction of the supply-side adjustments seen above: Since EU firms will optimally target the average household in autarky, it is best to lower both q_{cm} and p_{cm}. This quality downgrading increases the availability of less sophisticated but cheaper varieties which predominantly benefits low-income households and hence mitigates the extent to which the gains from trade are unequal.

In addition to the average rich-poor gap, also the majority of country-specific gaps become more moderate. In turn, as visible in the two leftmost plots in Table 7, the predicted supply side adjustments also affect the extent to which the unequal gains from trade vary with a country's income level. In the entry-only model shown on the left side, there is a strong positive relationship between a country's GDP per capita and the extent to which the gains from trade are unequal: Regressing the rich-poor gap on a country's GDP per capita (in logs) for example, implies a sizable positive slope of 0.088. Under the full model, the regression line between the two variables becomes much flatter and the slope is reduced to 0.065.

Finally, allowing for differences in productivity between incumbent and entering domestic firms lowers the welfare gap between rich and poor households further. As shown in columns (7) to (12) it actually reverses the gap on average to -2.87 in the entry-only model and -7.20 in the full one. This result is largely a consequence of the sizable price and quality adjustments described in Table 6. Particularly those countries with an initially small domestic share and hence an extensive counterfactual entry of less productive firms such as Romania or the Czech Republic see strong declines in p_{cm} and q_{cm} when moving to autarky which disproportionately benefits less wealthy households. As shown in the rightmost 2 plots, also the relationship between a country's gdp per capita and the welfare gap appears to be less pronounced in this case. Controlling for each country's import share does however restore the strong correlation between income per capita and the extent to which trade favors poorer consumers: On average, the gap is about 6.5 percentage points larger in richer economies than in poorer ones when controlling for size. Further, vertical differentiation continues to be important in this model as well and in this example category lowers the welfare gap by 4.3 percentage points.

5.3 Counterfactuals

This section extends the analysis of the previous subsection and computes the gains from trade for the full sample of products, using the demand and supply side parameter estimates described above. As for the example product, I quantify these gains by predicting the counterfactual welfare of each consumer group under autarky and compare it to that in the current trade equilibrium. I simulate household welfare in autarky in three scenarios: (1) Firms can enter and exit, (2) firms can additionally adjust prices, and (3) firms can adjust both the price and quality of their products.

As I did for the example category, I begin by constructing the percentage point difference in the welfare gains from trade between poor and rich households for each product category k using equation (34),

$$\Delta \text{Gap}_{kmt} = \frac{\mathbb{P}^{(k)}_{mt}(y_{(20)})_{\text{counter}}}{\mathbb{P}^{(k)}_{mt}(y_{(20)})_{\text{actual}}} - \frac{\mathbb{P}^{(k)}_{mt}(y_{(80)})_{\text{counter}}}{\mathbb{P}^{(k)}_{mt}(y_{(80)})_{\text{actual}}},$$

where m denotes a market, and t the respective time period. I then regress these gaps on importer dummies:

$$\Delta \text{Gap}_{kmt} = \delta_m \cdot \mathbb{I}\{\text{Importer} = m\} + \varepsilon_{kmt}. \tag{35}$$

The coefficients of interest are the δ_m, which can be interpreted as the average welfare gap in importing country m over all products and time periods. The larger this number is, the more will international trade benefit poor consumers relative to richer ones.

In section 6, I augment this regression with product as well as time fixed effects. However, since these dummies then absorb some of the variation in the rich-poor gaps, the estimates of δ_m would not reflect the absolute rich-poor gap in each country anymore in this case but rather capture average deviations in the gap from the product means. Since I am interested in both the absolute and relative magnitude of ΔGap_{kmt}, I do not include these controls in the baseline regressions. As I show later, the resulting cross-country patterns as well as the extent to which the supply side affects the welfare estimates also change only little with the inclusion of these dummy variables.

Tables 8 and 9 summarize the results of regression (35). As in the example category, I first report the results when holding productivity of entrants at the same level as that of incumbent domestic firms to highlight the importance of vertical differentiation. The left figure of Table 8 plots the average rich-poor gap in the gains from trade for each EU country for the entry-only

Table 8: Counterfactuals: Heterogeneous Gains from Trade - Fixed Productivity


a) Entry

b) Entry & Price Adj.

c) Full Adjustments

	$\Delta\%$ Welfare$_{\text{poor}}$ $-$ $\Delta\%$ Welfare$_{\text{rich}}$ ($\Delta\%\mathbb{P}_{(20)} - \Delta\%\mathbb{P}_{(80)}$)						
	Entry	Prices	Full Adj.		Entry	Prices	Full Adj.
France	3.71	3.32	1.81	Austria	4.36	3.93	2.25
Netherlands	3.95	3.56	1.72	Malta	3.31	3.38	2.14
Germany	1.83	1.55	0.89	Estonia	0.45	0.03	-0.24
Italy	1.24	0.94	0.31	Latvia	0.88	0.71	0.57
UK	4.70	4.48	1.95	Lithuania	-0.03	-0.31	-0.27
Ireland	3.63	3.38	0.93	Poland	1.58	1.31	0.84
Denmark	6.86	6.16	4.05	Czech Rep.	1.72	1.51	1.19
Greece	1.53	1.65	0.64	Slovakia	2.52	2.26	1.27
Portugal	1.88	1.66	0.59	Hungary	1.47	1.25	0.99
Spain	2.02	1.93	0.98	Romania	-0.26	-0.30	-0.06
Belgium	2.59	2.32	1.20	Bulgaria	1.33	1.50	0.93
Luxembourg	2.89	3.03	0.60	Slovenia	1.20	1.12	1.24
Sweden	2.00	1.61	1.00	Cyprus	6.69	3.98	2.75
Finland	2.69	2.64	1.84	**Average**	**2.47**	**2.17**	**1.19**

The above figures and table show the dummies δ_m from regression (35). The left figure is based on the model without price and quality adjustments. The figure in the center and the columns labeled *Prices* present the results for a model which only allows for entry and price adjustments when moving to autarky. The rightmost figures and columns show the result for the full model. Outlier Luxembourg is not shown in the plots.

model without price and quality adjustments. The middle figure shows the results when firms can additionally adjust the prices of their goods. Finally, the right figure as well as columns (4) and (8) in Table 8 summarize the results for the full adjustment model.

The results are largely in line with those seen in the previous section. In the fixed-productivity model, the average ΔGap_{kmt} is positive for all countries which means that poorer households benefit more from trade than richer ones do in each EU member country. Depending on the model, trade results in an on average 1 to 2.5 percentage points stronger decline in the cost of living for poorer households. There is also again considerable heterogeneity in terms of how unequal the gains from trade are in each country. While poorer households gain almost 3 percentage points more in

Table 9: Counterfactuals: Heterogeneous Gains from Trade (Variable Productivity)

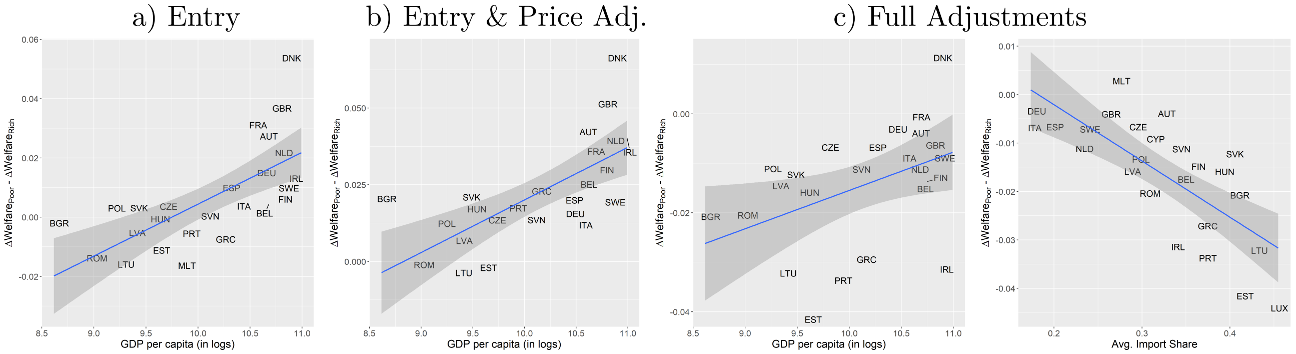

a) Entry

b) Entry & Price Adj.

c) Full Adjustments

	Entry only			Prices	Full Adjustments		
	$\Delta\%\mathbb{P}_{(20)}$	$\Delta\%\mathbb{P}_{(80)}$	Gap	Gap	$\Delta\%\mathbb{P}_{(20)}$	$\Delta\%\mathbb{P}_{(80)}$	Gap
France	8.04	5.17	2.87	3.40	5.68	5.87	-0.18
Netherlands	7.64	5.27	2.37	4.04	6.13	7.12	-1.00
Germany	4.40	3.11	1.29	1.66	4.20	4.42	-0.22
Italy	3.26	2.71	0.56	1.04	2.96	3.77	-0.80
UK	7.93	4.46	3.47	4.95	6.46	6.98	-0.52
Ireland	8.05	6.56	1.49	4.03	6.51	9.53	-3.02
Denmark	16.16	10.58	5.57	6.82	12.77	11.52	1.25
Greece	5.58	6.13	-0.55	2.11	6.31	9.14	-2.83
Portugal	7.78	8.14	-0.35	1.92	6.60	9.85	-3.25
Spain	5.10	4.31	0.80	2.11	4.10	4.90	-0.79
Belgium	9.85	9.41	0.44	2.70	7.46	9.09	-1.63
Luxembourg	5.93	7.48	-1.55	3.63	8.73	13.25	-4.52
Sweden	5.41	4.61	0.80	1.74	4.83	5.61	-0.78
Finland	8.73	8.27	0.46	2.79	7.29	8.65	-1.37
Austria	10.99	8.40	2.59	4.40	7.70	8.21	-0.51
Malta	7.18	8.62	-1.44	3.33	5.44	5.04	0.40
Estonia	6.09	7.41	-1.32	-0.01	6.72	10.76	-4.04
Latvia	3.85	4.19	-0.34	0.86	4.75	6.22	-1.47
Lithuania	4.62	6.02	-1.39	-0.55	5.34	8.67	-3.33
Poland	5.60	5.49	0.11	1.42	5.44	6.66	-1.22
Czech Rep.	6.83	6.27	0.56	1.62	6.10	6.89	-0.79
Slovakia	8.15	7.66	0.49	2.29	8.10	9.21	-1.11
Hungary	7.46	7.60	-0.14	1.46	7.20	8.68	-1.48
Romania	3.99	5.57	-1.58	-0.30	4.16	6.32	-2.16
Bulgaria	7.43	7.82	-0.39	1.85	7.50	9.69	-2.20
Slovenia	6.04	6.20	-0.16	1.17	7.09	8.33	-1.24
Cyprus	8.88	6.37	2.51	5.10	7.20	8.24	-1.04
Average	**7.07**	**6.44**	**0.64**	**2.43**	**6.40**	**7.88**	**-1.48**

The above figures and table show the dummies δ_m from regression (35). The left figure is based on the model without price and quality adjustments. The figure in the center and the columns labeled *Prices* present the results for a model which only allows for entry and price adjustments when moving to autarky. The rightmost figures and columns show the result for the full model. Differences in productivity between incumbent and entering domestic firms are inferred from the estimates of γ_k as described in section 5.2. Outlier Luxembourg is not shown in the plots. The last plot additionally excludes Denmark.

Denmark, the gains are almost equally distributed in Romania or Lithuania. More generally, there is a noticeable positive correlation in both models between a country's GDP per capita and the extent to which trade is pro poor: The average estimates of ΔGap_{kmt} are larger in richer economies such as Austria, the UK, or the Scandinavian countries but quite modest in Eastern Europe.

These results are once again a direct consequence of both specialization and a non-homothetic demand: As seen in section 5.1, demand for quality is increasing in income for the majority of goods. Since richer EU economies tend to specialize in more expensive, higher-quality varieties across product categories, their imports are comparably more attractive to households at the bottom end of the income distribution. In poorer EU economies on the other hand, imports and domestic varieties are on average similarly relevant for all income groups with the result that trade has relatively equal effects in those countries.

As evident from the rightmost columns of Table 8, accounting for quality specialization is also quantitatively important in the full sample. In the entry-only model, the average ΔGap_{kmt} equals about 2.5 percentage points and can take larger values of 4 to 7 percentage points in rich economies such as Denmark, the UK or Austria. In comparison, when accounting for quality and price adjustments this gap shrinks to just about 1.2 percentage points on average. In total, this translates into the gains from trade being $1 - 1.19/2.47 = 51.8\%$ less unequal when taking the supply side into account. As shown in the top 3 plots, these supply side adjustments also mitigate the extent to which the unequal gains from trade vary with a country's income level. In the entry-only model, the correlation between the average ΔGap_{kmt} and a country's GDP per capita is 0.67 and the slope of the fitted line in panel (a) is 0.014. These numbers drop to 0.53 and 0.004 in the full model.

Intuitively, accounting for the supply side also here takes into account that EU products today are partially unattractive or unavailable to poor consumers because of the specialization patterns seen in section 5.1. The results highlight that determining how this specialization would change in a counterfactual move to autarky is quantitatively important when estimating the unequal gains from trade. As was the case in the example category, the finding that these supply side adjustments reduce the welfare gap between rich and poor is due to the majority of domestic producers optimally lowering the quality of their products in autarky in order to target both richer and poorer households.

The results presented in Table 8 also provide a sense to what extent the rich-poor gap decomposes into a price- and quality-component. Price adjustments do partially close the rich-poor gap as well, although in itself considerably less than quality-adjustments. On average across countries, poorer households gain 2.17 percentage points more from trade than rich ones do in the price-adjustment model compared to 2.47 in the entry-only case, a decline in ΔGap_{kmt} of about 12%.

Finally, Table 9 summarizes the result of the model that additionally allows for heterogeneity in productivity across firms, in which case the welfare gap is again smaller. In contrast to the example category however, trade is on average still pro-poor in the entry-only model and only switches sign in the full model. In this case, endogenous vertical differentiation is hence particularly important for the question which consumer group benefits more from trade: On average, richer households gain about 1.48 percentage points more from trade than poorer households compared to 0.64 percentage points less in the entry-only model. Also here, trade is, depending on the specification, 4-6 percentage points more pro-poor in richer economies compared to poorer ones. In rich economies such as Germany, France or Austria, trade has again fairly balanced consequences and actually favors poor households in Denmark. In the poorer EU economies Romania and Lithuania, trade benefits the rich more throughout all models.

As suggested in the example category, the impact of heterogeneity in productivity is particularly strong for small countries. The rightmost figure of Table 9 shows a strong, negative correlation between a country's average import share and the welfare gap between poorer and richer households. This relationship is a product of both the quality choice of firms and the extent of predicted counterfactual entry: Since entry of domestic firms will be particularly strong in countries with large factual import shares, these countries will also experience sharp declines in the average productivity of firms. Given the inverse relationship between quality and productivity, the average quality will therefore also see the biggest declines in smaller countries which translates into the welfare implications described above.

The finding that trade might disproportionately benefit poor consumers when taking the firm-size distribution and productivity differentials into account is in line with recent evidence from scanner data for the U.S. by Faber and Fally (2017). The results presented in Table 9 hence show that these findings appear to also apply more broadly to the manufacturing sector and a wide range of countries and additionally vary quite strongly with country characteristics. In particular, trade in small and less wealthy countries magnifies nominal income inequality by almost 4 percentage points but has close to balanced consequences in larger and richer EU economies such as France, Germany, or the UK.

Table 10: Counterfactuals: Heterogeneous Gains from Trade - Price Index

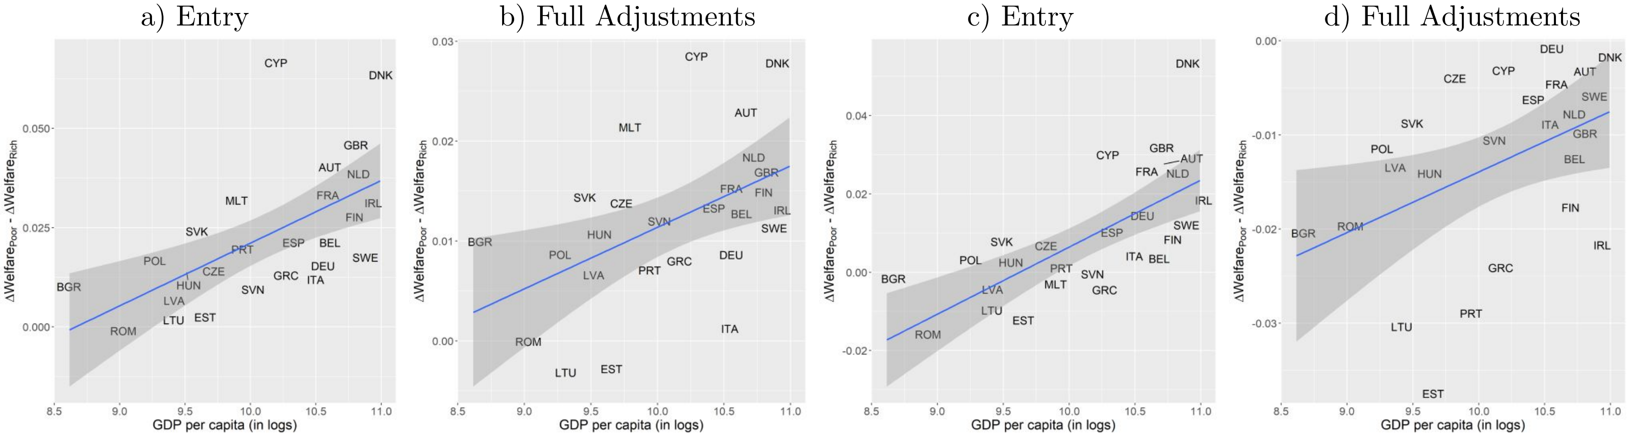

a) Entry b) Full Adjustments c) Entry d) Full Adjustments

	Fixed Productivity						Variable Productivity								
	Gap			Gap			Entry	Full Adjustments				Entry	Full Adjustments		
	Entry	Full		Entry	Full		Gap	$\Delta\mathbb{P}_{20}$	$\Delta\mathbb{P}_{80}$	Gap		Gap	$\Delta\mathbb{P}_{20}$	$\Delta\mathbb{P}_{80}$	Gap
FRA	3.49	1.60	AUT	4.20	2.36	FRA	2.76	4.63	5.01	-0.39	AUT	2.77	7.14	7.39	-0.25
NLD	3.68	1.91	MLT	3.00	2.21	NLD	2.34	5.77	6.47	-0.70	MLT	0.38	5.10	4.73	-0.49
GER	1.72	0.94	EST	0.42	-0.21	GER	1.26	3.86	4.02	-0.16	EST	-1.05	5.44	9.11	-3.67
ITA	1.17	0.05	LVA	0.85	0.59	ITA	0.59	1.92	2.89	-0.97	LVA	-1.37	4.49	5.86	-0.26
GBR	4.40	1.61	LTU	-0.01	-0.25	GBR	3.37	4.74	5.64	-0.91	LTU	-1.16	5.03	8.14	-3.11
IRL	3.30	1.38	POL	1.49	0.93	IRL	1.66	5.31	7.40	-2.09	POL	0.12	5.03	6.10	-1.07
DNK	6.17	2.85	CZE	1.58	1.45	DNK	5.53	9.12	9.39	-0.26	CZE	0.48	5.44	5.91	-0.47
GRC	1.49	0.87	SVK	2.23	1.51	GRC	-0.26	5.57	7.90	-2.33	SVK	0.59	7.45	8.25	-0.80
PRT	1.76	0.78	HUN	1.37	0.99	PRT	-0.08	6.21	9.03	-2.82	HUN	0.07	6.82	8.19	-1.37
ESP	1.94	1.25	ROM	-0.28	-0.08	ESP	0.82	3.60	4.13	-0.53	ROM	-1.41	3.91	5.95	-2.04
BEL	2.30	1.35	BGR	1.21	0.92	BEL	0.53	7.07	8.40	-1.33	BGR	-0.35	7.08	9.20	-2.12
LUX	3.05	1.03	SVN	1.11	1.27	LUX	-0.69	7.81	11.09	-3.28	SVN	-0.12	6.35	7.49	-1.13
SWE	1.93	1.06	CYP	6.83	2.78	SWE	0.97	4.62	5.29	-0.67	CYP	2.81	6.59	6.98	-0.39
FIN	2.59	1.42	**AVG**	**2.33**	**1.21**	FIN	0.69	4.93	6.63	-1.70	**AVG**	**0.80**	**5.59**	**6.91**	**-1.32**

All values represent percentage changes in the ideal price index $\mathbb{P}_{\mathrm{m}}(y_i)$ as defined by (23) from a move from the current trade equilibrium to autarky. $\Delta\mathbb{P}_{20}$ refers to the percentage change in welfare of households at the 20th percentile of the income distribution while $\Delta\mathbb{P}_{80}$ denotes the corresponding change for the 80th percentile. *Gap* refers to the difference $\Delta\mathbb{P}_{20} - \Delta\mathbb{P}_{80}$. The 2 figures on the left and the *Fixed Productivity* columns show the respective values when productivity of entering firms is held constant. The figures and results on the right hand side refer to the model that allows for variation in productivity. Outlier Luxembourg is not shown in the plots.

Price Indexes Instead of relying on regression (35), it is also straightforward to exploit the structure of the model and use the aggregate, household-specific consumer price indexes as defined in equation (23) to compute the gains from trade. Table 10 summarizes the results which are largely comparable to those found when using regression (35): Trade again benefits poorer households more than richer ones in each EU country when productivity is held fixed with gaps of 2.33 and 1.21, respectively. Supply side responses hence mitigate the extent to which the gains from trade are unequal by 48%. Also the correlation between the welfare gaps and each country's GDP per capita falls from 0.67 to 0.53 and quality adjustments continue to be quantitatively more important than price adjustments. In the variable-productivity model the welfare gaps equal 0.8 and -1.32 percentage points and are hence very close to the values obtained in regression (35) of 0.64 and -1.48.

As shown in Appendix Table B3, the results are also largely unaffected when using expenditure weights ω_k in equation (23). In order to construct these respective price indexes, I infer country-period specific expenditure weights ω_k from the data using equation (30). In this case, the gap in the gains from trade between poorer and richer consumers equals 1.96 and 1.03, respectively, in the model that holds productivity fixed. Product differentiation hence reduces the extent to which the gains from trade are unequal by virtually the same percentage as when using uniform weights (47% compared to 48%). Accounting for entry by less productive firms again reverses the sign of the gap. A noticable difference is that there is only a very small difference in the average gap when comparing the entry and the full model (-1.21 and -1.26). This is however largely driven by one country, Malta, for which the gap in the entry case exceeds -22. Excluding only Malta would already significantly increase the gap to -0.46. More generally, the introduction of expenditure weights does result in several very large and very small gaps, particularly for small countries. The simple average is therefore quite sensitive to such cases.

6 Robustness

This section summarizes a set of robustness checks which evaluate in how far the baseline results depend on certain assumptions or features of the data. Table 11 gives a brief overview of how sensitive the results are to varying specifications while a more detailed summary is provided in Appendix B.

Instruments As mentioned above, as a first robustness check, I employ a set of alternative instruments based on measures of shipping costs to test to what extent the results might depend on the instrument choice. These costs are largely determined by factors out of the control of firms such as the distance to the importing country, fuel prices, or the cost of freight insurance and hence an arguably exogenous driver of exporter prices.

Data Appendix D discusses the the validity of the alternative instruments in more detail. As shown in Tables D5 and D6, there is a robust, positive first-stage correlation between each instru-

Table 11: Robustness Checks: Summary

	Fixed Productivity			Variable Productivity	
	Entry	Full	Ratio	Entry	Full
Baseline	2.47	1.19	0.48	0.64	-1.48
Alternative Instrument	1.44	0.98	0.68	0.25	-0.80
Weighted Price Index	1.96	1.03	0.53	-1.26	-1.21
Excluding Intermediates	2.23	1.22	0.55	0.89	-1.30
6-Digit Aggregation	2.39	0.80	0.33	-0.09	-2.46
4-Digit Aggregation	1.59	0.65	0.41	-0.75	-2.25
2-Digit Aggregation	1.05	0.42	0.40	-0.24	-1.12

Each value reflects the welfare gap in the gains from trade between households at the 20th percentile compared to those at the 80th percentile of the income distribution. All numbers are simple averages over all EU countries in the sample. The *Fixed Productivity* columns show the respective values when productivity of entering firms is held constant. The *Variable Productivity* columns refer to the model that allows for variation in productivity across firms and between entering and incumbent domestic firms.

ment and unit values across product categories. As in the baseline specification, the estimated price coefficients are also larger in absolute value compared to OLS for the majority of product categories (69% of cases compared to 78% in the baseline). Panel b) of Figure D4 shows graphically that the distribution of the difference between 2SLS and OLS coefficients, $\alpha_{2SLS} - \alpha_{OLS}$ has more mass to the right of 0, implying less negative price coefficients in the OLS case. Panel d) shows that the 2SLS distribution of price coefficients has a predominantly negative support and more mass on the left tail compared to the OLS estimates, which are much more concentrated between -5 and 0.

Table B1 presents the main results when using shipping costs as instruments for prices. More specifically, the table summarizes the results of regression (35) when both measures of shipping costs are included in the instrument matrix, although the estimates are similar if each one is used individually. The qualitative and quantitative results are largely in line with those obtained in the baseline specification: When productivity is held fixed, trade benefits poorer consumers more than richer ones, with a difference of 1.44% in the entry-only case, and 0.98% in the full model. While these numbers are slightly lower compared to the baseline results of 2.47% and 1.19%, product differentiation is still quantitatively important and implies 46.9% less unequal gains than without adjustments.

Also here, the productivity channel in combination with vertical differentiation reverses the sign of the gap, in this case to -0.8. As in the baseline case, the entry-only model would have still predicted a moderate positive gap of 0.25. These figures are reassuring that the main findings are not driven by the instrument choice.

Intermediates Another potential concern is related to the fact that parts of the products in the sample are intermediate products rather than final goods. Those are in turn not directly bought by consumers but rather by companies using them in the production of final goods. In order to assess this issue in my counterfactuals, I remove intermediates from the data and simulate the gains from trade for the remaining subsample. In order to classify products, I use the UN Conference on Trade and Development's (UNCTAD) classification of products, which assigns each HS6 category one of 4 types (consumer goods, capital goods, intermediate goods, and raw minerals).[36]

As summarized in Table B2, the estimates do not differ substantially from the baseline specification and show that the gap between rich and poor households is still present. I find that the difference in welfare gains equals 2.23% and 0.89% in the entry-only models, respectively, which is in line with the results for the full sample (2.47% and 0.64%). Also the impact of product differentiation is comparable to before: Allowing for price and quality adjustments results in 1.01 and 0.41 percentage point reductions in the welfare gap measure.[37]

Aggregation The next robustness check addresses the issue that the estimates of the gains from trade are known to be sensitive to the aggregation level. Ossa (2015) for example shows that disaggregating the data by 3 digit industry instead of using a single elasticity of substitution and aggregate data magnifies the gains from trade by approximately factor 3. While the focus of this paper is less on measuring aggregate welfare gains from trade, it is still plausible that the aggregation level also matters for the degree to which the gains from trade are unequal.

To address this issue, I aggregate the data up to the 2-, 4-, and 6-digit industry level and reestimate the model. Appendix Figure B1 and Table B4 show how the welfare results are affected by changing the level of industry disaggregation. Qualitatively, the results are similar to those at the 8-digit level: Trade lowers the price index of poorer consumers by 2.46, 2.25, and 1.12 percentage points less than that of richer households, respectively, compared to 1.48 found in the 8-digit case.

Quantitatively, the gains from trade tend to be less unequal when aggregating. In the fixed productivity model, the gaps decline from 2.47 and 1.19 in the 8-digit case to 1.05 and 0.42 when aggregating the data up to 2-digit industries. Also in the variable productivity model the welfare gaps are smaller in absolute terms with values of 0.24 and 1.12, which is consistent with larger predicted aggregate gains on the 8-digit level partially carrying over to the extent to which they are unequal.

Regardless of the aggregation level however, not accounting for price and quality adjustments does overstate the unequal gains from trade in comparable magnitudes: Allowing for endogenous product differentiation results in a 60 to 66 percent less unequal distribution of the gains from trade depending on the aggregation level, which is in line with the results in the 8-digit case (52%).

---

[36] Since my data is on a different level of disaggregation (pc8/cn8 versus HS6), I assign each cn8 code a category based on the first 6 digits, which coincide with those of HS6 categories. When aggregating up to pc8plus categories, I assign pc8plus categories a group based on EU trade volumes: If the respective pc8plus category consists predominantly of HS6 imports in group 1 for example, I assign this category group 1 as well.

[37] I also ran the above regressions for the other groups and found that the above patterns are to some extent also present for the group of intermediates. This should not necessarily come as a surprise given that higher-quality final producers tend to also use higher-quality intermediates (see e.g. Manova and Zhang (2012)). One should therefore expect to not only see a higher demand for higher-quality final goods but also for higher-quality intermediates in richer countries due to specialization. Further, access to higher-quality intermediates would then carry through to greater access to higher-quality final goods and have plausibly similar implications for inequality.

Generally, the results also tend to be quite a bit more noisy the more the data is aggregated, which can be expected for two reasons: First, particularly 2-digit industries frequently encompass a large number of different goods with a nontrivial extent of heterogeneity in terms of characteristics and production costs.[38] Aggregate unit values on this level are therefore a much more noisy measure of prices. Second, with a small number of industries, a single sector can have a much bigger impact on the estimated gains from trade.

Product Fixed Effects The final robustness check addresses the ideal price index as defined by equation (23). A potential concern with this index is that aggregation over many product categories might hide that only very few product categories drive the results. One might also be worried that the cross-country variation in the rich-poor gaps is driven only by differences in the consumption of product categories and not varieties within these categories.

In order to address these questions, I regress the welfare gaps on the product-importer level, ΔGap_{kmt}, not only on importer dummies but also product and time fixed effects:

$$\Delta\text{Gap}_{kmt} = \delta_m \cdot \mathbb{I}\{\text{Importer} = m\} + \delta_k \cdot \mathbb{I}\{\text{Product} = k\} + \delta_t \cdot \mathbb{I}\{\text{Period} = t\} + \varepsilon_{kmt}. \tag{36}$$

Table B5 summarizes the results of this regression. Including product dummies essentially absorbs any variation in the rich-poor gaps that is due to the respective product category k, which mitigates for example the impact of individual product categories with extreme gaps in each country. The coefficients on the importer dummy, δ_m, therefore do not reflect the absolute welfare gaps in each country in this case, but rather average deviations from the product means.

The estimates are nevertheless informative in so far as they confirm the cross-country pattern found above, with trade in richer EU countries disproportionately benefiting poor consumers more than in poorer countries. As evident from Table B5, there continues to be a positive correlation between country GDP per capita and the average ΔGap_{kmt}, both in the entry-only and full models. In terms of magnitudes, the welfare gap in the richer economies Austria and the UK is for example about 6 percentage point bigger than in Romania and Bulgaria in the entry-only models, which is quantitatively similar to the results obtained from regression (35). The inclusion of fixed effects does also not alter the finding that price and quality adjustments result in smaller cross-country variation in the rich-poor gap. When using the full model, the difference between rich and poor countries falls to 1.5 to 2 percentage points.

7 Conclusions

In this paper, I developed a structural framework to study the unequal gains from international trade under endogenous vertical differentiation. Based on the observation that consumption baskets systematically differ with household income and that countries are differently good at producing

---

[38] As an example, one particular 2-digit industry summarizes firms that produce tractors, passenger cars, bicycles, and baby carriages.

these baskets, I built a multi-sector model and quantified how trade affects the cost of living for different income groups. The model can be brought to the data using widely available data on trade flows, domestic production and the income distribution, making it applicable to a wide range of countries, and the paper provides estimates of demand and supply parameters for more than 3,000 industries.

Overall, I find the gains from trade to be fairly moderately unequal. While there is significant cross-country variation in terms of which groups benefits more from trade, essentially two channels mitigate the extent to which the gains from trade are unequal. On one hand, trade induces domestic firms to specialize in varieties for which there is less foreign competition. The inflow of cheaper Chinese varieties for example is disproportionately beneficial to poor households, but also induces quality-upgrading by domestic firms which has the opposite welfare effects. Additionally, exit of less productive, low-quality producing domestic firms counteracts these direct effects even further. Taken together, the model actually predicts that in the majority of EU countries, it is richer households who experience a stronger decline in the cost of living, particularly in small economies.

A natural question is how the results above compare with those found in the literature on the impact of trade and import competition on workers with differing skill and education levels. While studies by Autor, Dorn, and Hanson (2013), Autor et al. (2014) or Hummels et al. (2014) have emphasized the heterogeneous impact of trade on nominal income, they have been silent on the implications for real income inequality. Autor et al. (2014) for example find that low skilled workers in the U.S. experienced a 7.8% decline in earnings as consequence of extensive import competition from China compared to high-skilled workers who saw virtually no changes in wages. This number dominates the relatively moderate gaps found in this paper, which e.g. equal in the similarly rich Scandinavian countries between -1.37 and 1.25. Qualitatively, also a model without the differentiation and productivity-channels would have come to this conclusion, but would have predicted much smaller effects on real income inequality.

Appendices

A Summary Statistics

Table A1: Summary Statistics

	N	Mean	Std.Dev	Min	Max
Trade and Production:					
Observations per product category	3,103	4,258	3,016	4	21,995
Expenditure Shares	13,212,843	0.0890	0.2060	2.54e-09	0.9999
Exporters per product	3,103	51.79	21.48	1	169
Importers per product	3,103	24.79	3.36	1	27
Unit Values	13,212,843	10.45	185.13	2.57e-06	177,996.5
Total quarterly demand (1000 EUR)	1,177,936	30,580.23	238,892.5	0.0001496	3.28e+07
Import Share	719,478	0.2854	0.2635	1.92e-07	0.9999
Number of Companies	3,085	46.01	178.41	1	3558.42
Country Characteristics:					
Exporter Population	13,212,843	110,782.3	276,896.8	9.65	1,319,625
Bilateral Distance (in km)	13,212,843	2,752.48	3,319.40	6.69	19,586.18
Exporter GDP per capita	13,212,843	28,096.4	17,341.5	99.16	129,870.2
Median Disp. Income (EUR)	27	15,078	10,949	1,904	42,909
Gini Coefficient - Income	27	0.30	0.04	0.23	0.37
Instruments:					
Hausman (1996) IV	13,212,843	10.49	153.62	0.0000105	58,904.06
Feenstra, Romalis (2014) IV	13,212,613	-.0260	1.0091	-17.1723	19.5264
Shipping Cost per unit	13,212,613	3.1822	500.0253	4.62e-08	272,315.7
Shipping Cost per Euro	13,212,613	0.0011	0.02452	1.70e-08	12.75
Δ Exchange Rates	13,212,843	0.1579	0.5708	-0.8074	1.8293

B Robustness

Table B1: Robustness: Alternative Instrument

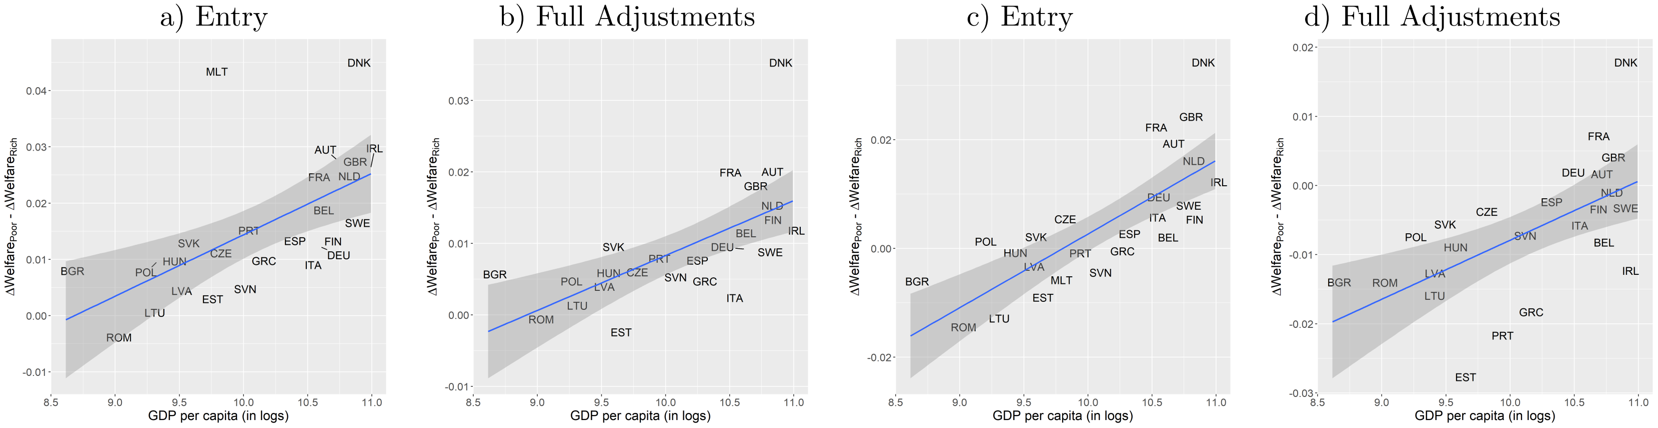
a) Entry b) Full Adjustments c) Entry d) Full Adjustments

	Fixed Productivity					Variable Productivity									
	Gap			Gap			Entry	Full Adjustments				Entry	Full Adjustments		
	Entry	Full		Entry	Full		Gap	$\Delta\mathbb{P}_{20}$	$\Delta\mathbb{P}_{80}$	Gap		Gap	$\Delta\mathbb{P}_{20}$	$\Delta\mathbb{P}_{80}$	Gap
FRA	2.59	1.91	AUT	2.79	1.92	FRA	2.11	6.23	5.62	0.62	AUT	1.80	8.50	8.45	0.05
NLD	2.61	1.43	MLT	4.46	2.65	NLD	1.74	6.51	6.71	-0.20	MLT	-0.46	9.99	9.21	0.78
GER	1.21	0.92	EST	0.43	-0.14	GER	0.81	4.50	4.38	0.12	EST	-0.77	7.89	10.6	-2.67
ITA	0.77	0.14	LVA	0.57	0.30	ITA	0.44	2.98	3.46	-0.48	LVA	-0.38	5.11	6.28	-1.16
GBR	2.76	1.71	LTU	0.18	0.23	GBR	2.30	7.01	6.70	0.31	LTU	-1.16	7.45	9.15	-1.70
IRL	2.64	1.09	POL	0.94	0.57	IRL	1.09	7.60	8.73	-1.13	POL	-0.01	5.55	6.40	-0.85
DNK	4.63	3.63	CZE	1.23	0.70	DNK	3.55	14.08	12.20	1.88	CZE	0.41	6.32	6.80	-0.48
GRC	1.10	0.59	SVK	1.16	0.85	GRC	-0.18	6.76	8.48	-1.72	SVK	0.08	8.48	8.94	-0.46
PRT	1.39	0.69	HUN	1.09	0.68	PRT	-0.22	7.37	9.43	-2.07	HUN	0.06	7.35	8.34	-0.99
ESP	1.15	0.71	ROM	-0.51	-0.16	ESP	0.40	4.36	4.7	-0.34	ROM	-1.58	4.35	5.86	-1.50
BEL	1.74	1.24	BGR	0.68	0.47	BEL	0.33	8.38	9.10	-0.72	BGR	-0.73	7.94	9.45	-1.50
LUX	-1.61	-1.03	SVN	0.60	0.63	LUX	-3.42	13.69	19.70	-6.02	SVN	-0.32	7.28	8.10	-0.82
SWE	1.52	0.98	CYP	1.33	2.58	SWE	0.92	5.28	5.50	-0.22	CYP	-0.50	10.63	10.64	-0.01
FIN	1.44	1.28	**AVG**	**1.44**	**0.98**	FIN	0.40	8.10	8.34	-0.24	**AVG**	**0.25**	**-0.80**		

The above figures and table show the dummies δ_m from regression (35). The columns $\Delta\mathbb{P}_{20}$ and $\Delta\mathbb{P}_{80}$ additionally document the percentage change in welfare of households at the 20th and 80th percentile of the income distribution, respectively. The 2 figures on the left and the *Fixed Productivity* columns show the respective values when productivity of entering firms is held constant. The figures and results on the right hand side refer to the model that allows for variation in productivity across firms and between entering and incumbent domestic firms. Outlier Luxembourg is not shown in the plots.

Table B2: Robustness: Excluding Intermediates

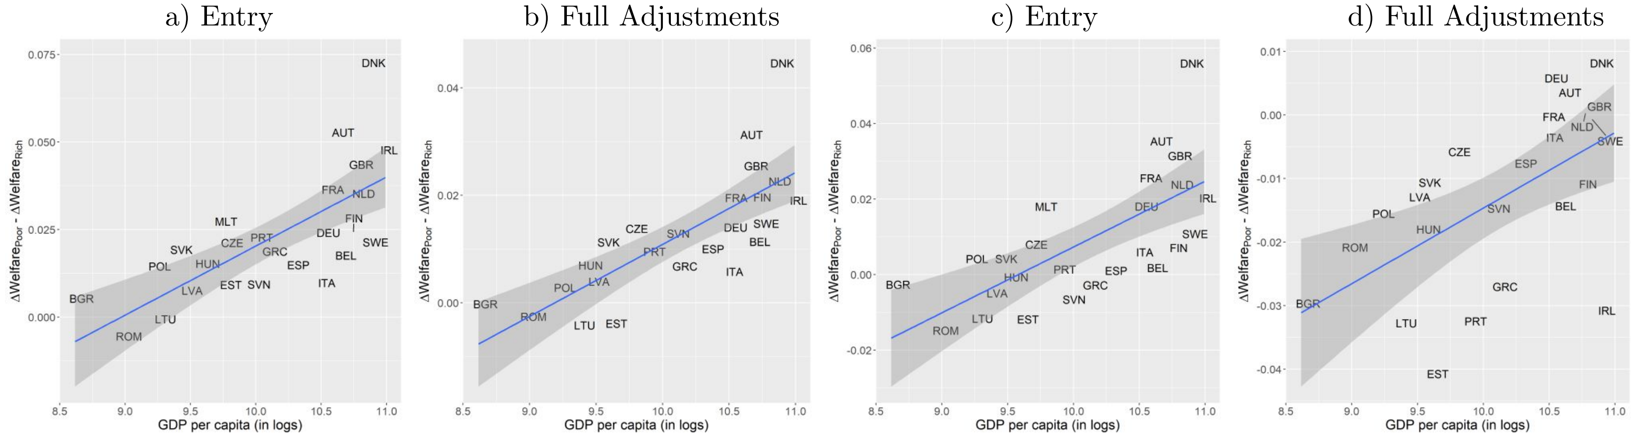

a) Entry b) Full Adjustments c) Entry d) Full Adjustments

	Fixed Productivity				
	Gap			Gap	
	Entry	Full		Entry	Full
FRA	3.84	2.09	AUT	5.48	3.26
NLD	3.73	2.30	MLT	2.53	1.48
GER	2.20	1.54	EST	0.81	-0.25
ITA	1.18	0.71	LVA	0.96	0.53
GBR	4.15	2.41	LTU	0.15	-0.55
IRL	4.57	1.77	POL	1.24	0.42
DNK	7.46	4.60	CZE	1.92	1.23
GRC	1.62	0.81	SVK	2.08	0.99
PRT	2.49	1.09	HUN	1.39	0.84
ESP	1.29	0.86	ROM	-0.76	-0.40
BEL	2.01	1.27	BGR	0.33	-0.16
LUX	-1.45	-0.32	SVN	1.14	1.15
SWE	2.33	1.61	CYP	5.14	1.51
FIN	2.44	2.12	**AVG**	**2.23**	**1.22**

	Variable Productivity								
	Entry	Full Adjustments				Entry	Full Adjustments		
	Gap	$\Delta\mathbb{P}_{20}$	$\Delta\mathbb{P}_{80}$	Gap		Gap	$\Delta\mathbb{P}_{20}$	$\Delta\mathbb{P}_{80}$	Gap
FRA	2.74	6.01	5.93	0.08	AUT	3.72	8.56	8.33	0.22
NLD	2.58	6.58	6.56	0.02	MLT	1.99	1.90	1.40	0.50
GER	1.61	4.64	4.18	0.46	EST	-1.00	6.91	10.86	-3.96
ITA	0.41	3.88	4.35	-0.47	LVA	-0.31	4.88	6.28	-1.40
GBR	2.95	6.69	6.64	0.05	LTU	-0.97	5.30	8.69	-3.39
IRL	2.22	7.13	10.09	-2.97	POL	0.22	5.86	7.52	-1.66
DNK	5.79	13.09	12.16	0.93	CZE	0.99	6.87	7.55	-0.69
GRC	-0.09	6.54	9.12	-2.58	SVK	0.60	7.60	8.78	-1.18
PRT	-0.05	7.50	10.63	-3.13	HUN	0.12	7.61	9.29	-1.68
ESP	0.30	3.87	4.52	-0.65	ROM	-1.68	4.06	6.25	-2.19
BEL	0.40	7.95	9.27	-1.32	BGR	-0.45	7.33	10.41	-3.08
LUX	-4.41	4.27	7.88	-3.61	SVN	-0.43	7.76	9.34	-1.58
SWE	1.26	5.36	5.43	-0.07	CYP	5.09	4.77	5.20	-0.43
FIN	0.52	7.82	9.02	-1.20	**AVG**	**0.89**	**6.32**	**7.62**	**-1.30**

The above figures and table show the dummies δ_m from regression (35). The columns $\Delta\mathbb{P}_{20}$ and $\Delta\mathbb{P}_{80}$ additionally document the percentage change in welfare of households at the 20th and 80th percentile of the income distribution, respectively. The 2 figures on the left and the *Fixed Productivity* columns show the respective values when productivity of entering firms is held constant. The figures and results on the right hand side refer to the model that allows for variation in productivity across firms and between entering and incumbent domestic firms. Outlier Luxembourg is not shown in the plots.

Table B3: Counterfactuals: Heterogeneous Gains from Trade - Weighted Price Index

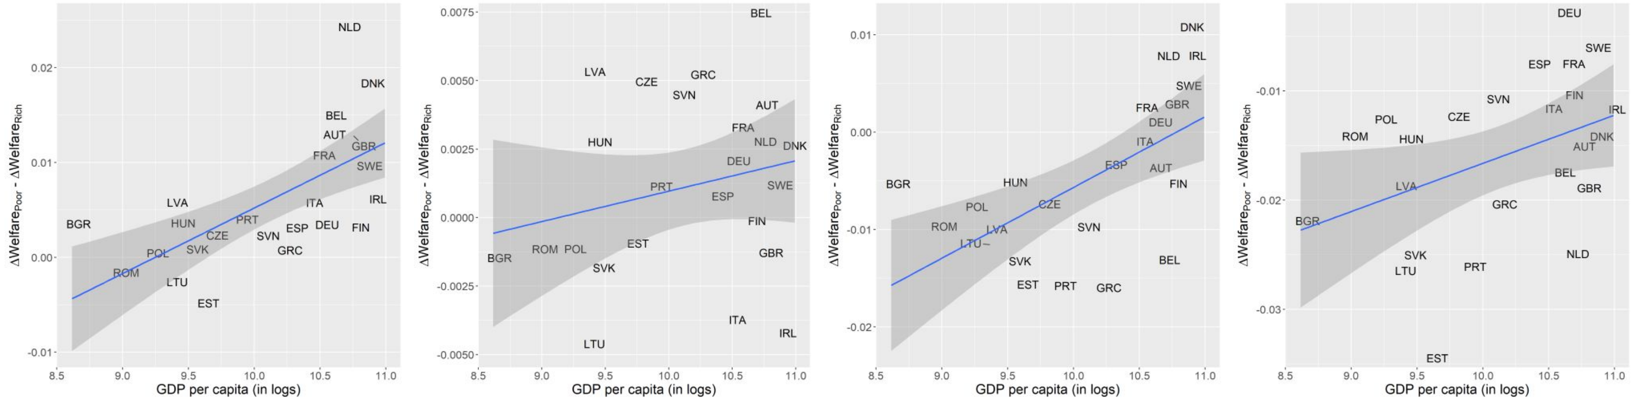
a) Entry b) Full Adjustments c) Entry d) Full Adjustments

	Fixed Productivity				
	Gap			Gap	
	Entry	Full		Entry	Full
FRA	1.09	0.36	AUT	1.37	0.38
NLD	2.50	0.30	MLT	8.70	7.60
GER	0.42	0.18	EST	-0.41	-0.12
ITA	0.50	-0.40	LVA	0.65	0.56
GBR	1.28	-0.16	LTU	-0.34	-0.49
IRL	0.69	-0.39	POL	0.12	-0.14
DNK	1.91	0.29	CZE	0.31	0.47
GRC	0.15	0.49	SVK	0.20	-0.16
PRT	0.32	0.14	HUN	0.29	0.30
ESP	0.38	0.05	ROM	-0.23	-0.14
BEL	1.41	0.72	BGR	0.28	-0.17
LUX	2.25	5.68	SVN	0.30	0.42
SWE	1.04	0.15	CYP	27.34	11.85
FIN	0.39	0.01	**AVG**	**1.96**	**1.03**

	Variable Productivity								
	Entry	Full Adjustments				Entry	Full Adjustments		
	Gap	$\Delta\mathbb{P}_{20}$	$\Delta\mathbb{P}_{80}$	Gap		Gap	$\Delta\mathbb{P}_{20}$	$\Delta\mathbb{P}_{80}$	Gap
FRA	0.32	2.82	3.64	-0.82	AUT	-0.44	5.78	7.35	-1.57
NLD	0.86	4.35	6.77	-2.42	MLT	-22.05	8.02	5.78	2.24
GER	0.03	2.42	2.77	-0.35	EST	-1.49	2.09	5.47	-3.38
ITA	-0.17	1.53	2.76	-1.23	LVA	-0.92	4.08	5.88	-1.80
GBR	0.37	3.90	5.72	-1.82	LTU	-1.16	3.38	6.09	-2.71
IRL	0.86	4.79	6.03	-1.24	POL	-0.69	2.54	3.73	-1.19
DNK	1.15	5.70	7.05	-1.35	CZE	-0.81	4.38	5.68	-1.30
GRC	-1.52	7.23	9.20	-1.97	SVK	-1.25	6.65	9.09	-2.43
PRT	-1.50	3.44	5.99	-2.54	HUN	-0.59	4.17	5.54	-1.37
ESP	-0.26	1.21	2.03	-0.82	ROM	-1.05	2.70	4.18	-1.48
BEL	-1.39	5.00	6.68	-1.68	BGR	-0.61	3.78	6.04	-2.26
LUX	-2.20	11.03	10.53	0.50	SVN	-1.05	7.11	8.25	-1.14
SWE	0.40	4.98	5.65	-0.68	CYP	1.56	11.22	8.12	3.09
FIN	-0.46	2.80	3.77	-0.97	**AVG**	**-1.26**	**4.71**	**5.91**	**-1.21**

All values represent percentage changes in the ideal price index $\mathbb{P}_{\mathrm{m}}(y_i)$ as defined by (23) from a move from autarky to the current trade equilibrium. $\Delta\mathbb{P}_{20}$ refers to the percentage change in welfare of households at the 20th percentile of the income distribution while $\Delta\mathbb{P}_{80}$ denotes the corresponding change for the 80th percentile. *Gap* refers to the difference $\Delta\mathbb{P}_{20} - \Delta\mathbb{P}_{80}$. The 2 figures on the left and the *Fixed Productivity* columns show the respective values when productivity of entering firms is held constant. The figures and results on the right hand side refer to the model that allows for variation in productivity across firms and between entering and incumbent domestic firms. Outliers Luxembourg, Malta, and Cyprus are not shown in the plots.

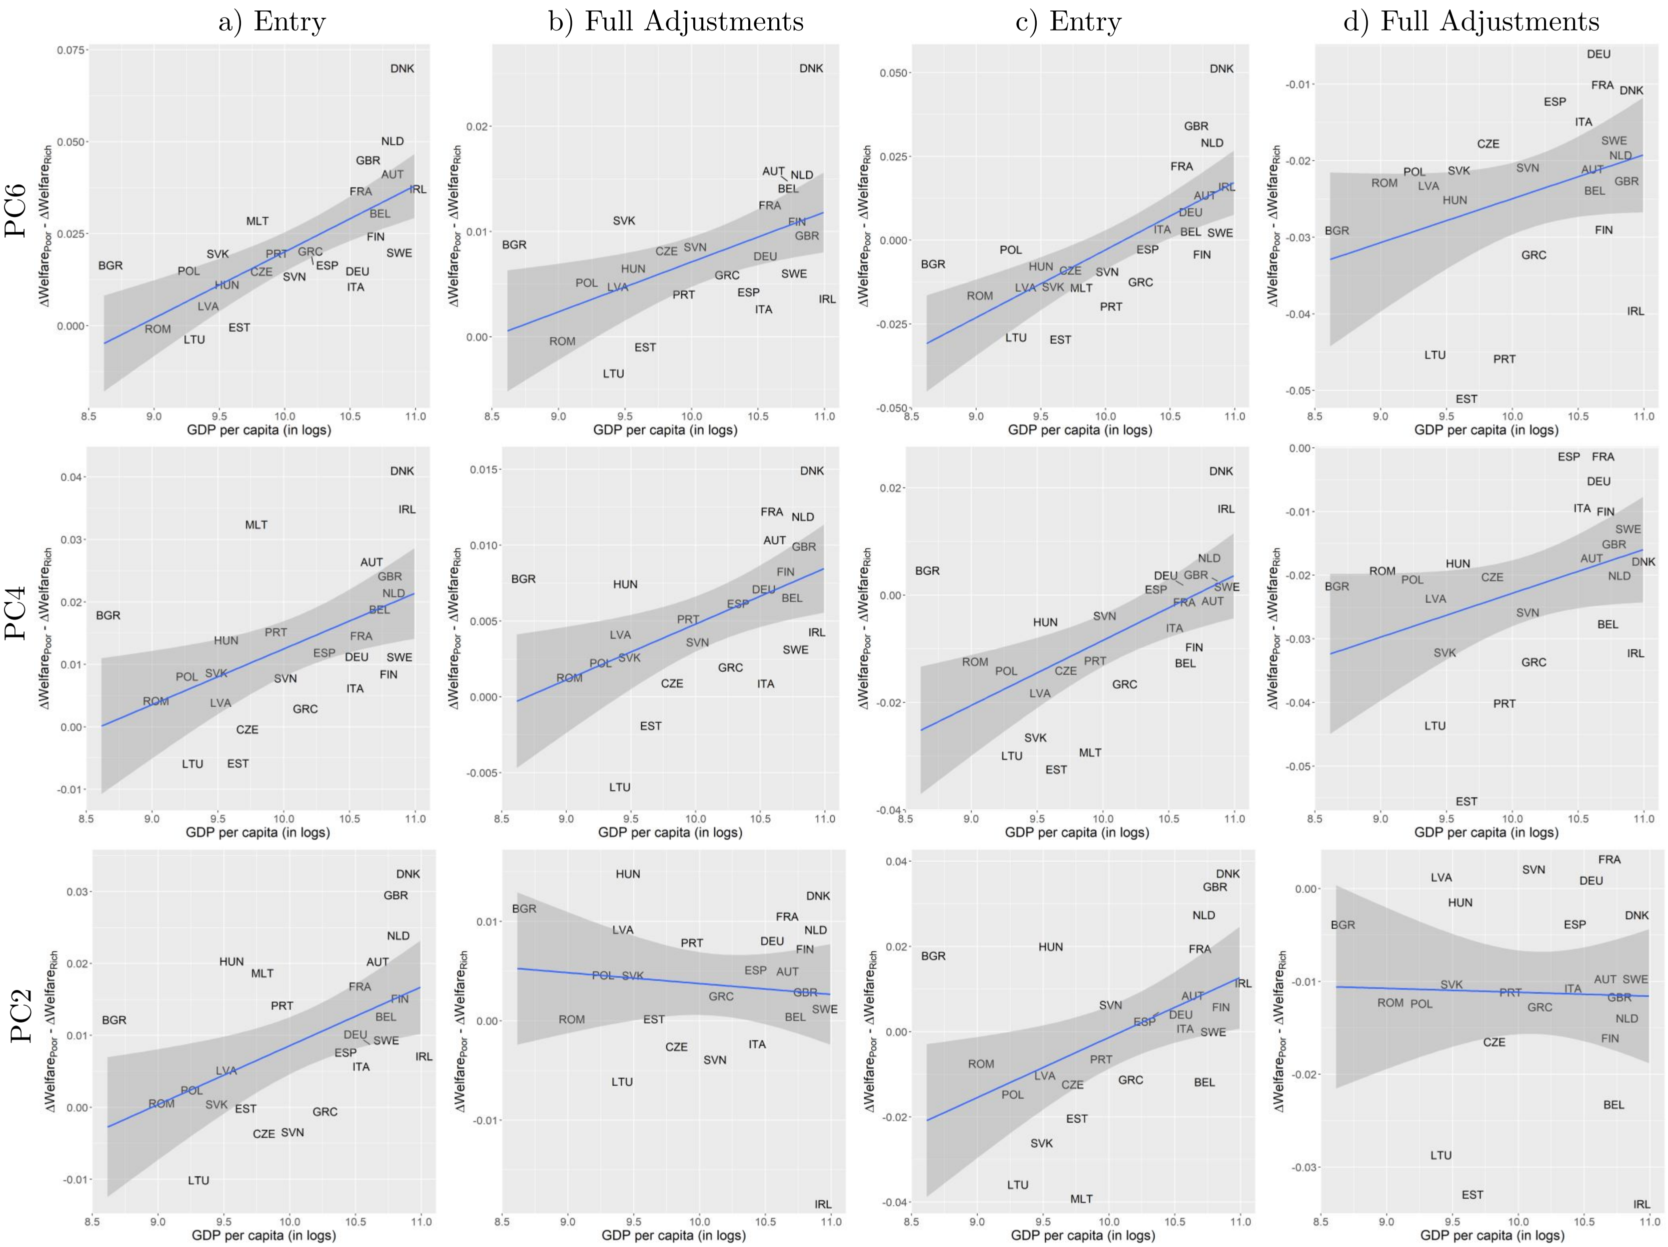

Figure B1: **Robustness: Varying Aggregation Levels**: The above figures plot the dummies δ_m from regression (35) against each country's GDP per capita (in logs) for varying aggregation levels. The 2 leftmost figures show the respective values when productivity of entering firms is held constant. The figures on the right hand side refer to the model that allows for variation in productivity across firms and between entering and incumbent domestic firms. Outlier Luxembourg is not shown in the plots.

Table B4: Robustness: Varying Aggregation Levels

	PC6				PC4				PC2			
	Fixed Prod.		Variable Prod.		Fixed Prod.		Variable Prod.		Fixed Prod.		Variable Prod.	
	Entry	Full	Entry	Full	Entry	Full	Entry	Full	Entry	Full	Entry	Full
FRA	3.47	1.32	2.42	-1.10	1.34	1.18	0.06	-0.24	1.59	0.95	1.78	0.24
NLD	4.82	1.47	2.70	-2.02	2.02	1.14	0.56	-1.89	2.29	0.85	2.91	-1.32
GER	1.67	0.84	0.63	-0.70	1.01	0.76	0.19	-0.63	0.87	0.88	0.57	0.17
ITA	1.26	0.19	0.12	-1.39	0.50	0.04	-0.75	-1.06	0.47	-0.16	-0.04	-1.00
GBR	4.64	0.91	3.20	-2.19	2.53	0.94	0.52	-1.62	2.84	0.36	3.24	-1.09
IRL	3.90	0.43	1.82	-3.86	3.60	0.48	1.75	-3.11	0.81	-1.77	0.96	-3.32
DNK	7.20	2.62	5.35	-0.98	4.22	1.54	2.45	-1.67	3.35	1.33	3.88	-0.21
GRC	1.66	0.65	-1.46	-3.13	0.41	0.24	-1.52	-3.24	-0.16	0.32	-0.96	-1.20
PTG	1.75	0.47	-1.75	-4.49	1.63	0.56	-1.09	-3.90	1.52	0.86	-0.48	-1.04
ESP	1.45	0.50	-0.06	-1.32	1.07	0.66	-0.04	-0.25	0.67	0.44	0.46	-0.46
BEL	3.26	1.38	0.05	-2.29	1.77	0.61	-1.13	-2.88	1.16	-0.03	-1.34	-2.40
LUX	1.30	-0.94	-3.04	-6.57	1.95	0.29	-4.53	-6.84	0.08	-0.47	-6.31	-2.55
SWE	1.79	0.53	0.00	-1.64	1.24	0.36	0.32	-1.39	0.83	0.19	0.18	-1.05
FIN	2.62	1.16	-0.21	-2.81	0.96	0.87	-0.84	-0.88	1.58	0.65	0.69	-1.53
AUT	3.93	1.48	1.13	-2.10	2.76	1.08	0.05	-1.81	2.12	0.57	1.01	-1.05
MLT	3.05	1.50	-1.63	-0.48	3.35	1.65	-3.06	-0.73	1.97	1.24	-3.75	0.09
EST	0.15	-0.03	-2.75	-5.01	-0.47	-0.14	-3.11	-5.44	-0.12	-0.05	-2.21	-3.22
LTV	0.74	0.55	-1.20	-2.30	0.51	0.46	-1.69	-2.25	0.42	0.99	-1.21	0.05
LTU	-0.17	-0.41	-2.68	-4.63	-0.47	-0.64	-2.85	-4.47	-0.91	-0.68	-3.41	-2.95
POL	1.30	0.51	-0.50	-2.05	0.69	0.27	-1.55	-1.93	0.11	0.53	-1.30	-1.16
CZE	1.28	0.75	-0.68	-1.87	0.07	0.04	-1.27	-1.91	-0.26	-0.19	-1.06	-1.73
SVK	2.15	1.17	-1.17	-2.22	0.98	0.31	-2.52	-3.10	-0.06	0.53	-2.44	-0.95
HUN	0.92	0.58	-1.00	-2.41	1.27	0.79	-0.64	-1.92	1.93	1.56	1.83	-0.07
ROM	-0.28	-0.11	-1.87	-2.38	0.30	0.08	-1.38	-2.06	-0.04	-0.06	-0.92	-1.31
BLG	1.45	0.81	-0.92	-3.00	1.68	0.73	0.33	-2.29	1.12	1.06	1.62	-0.47
SVN	1.53	0.78	-0.73	-2.18	0.90	0.31	-0.25	-2.70	-0.24	-0.46	0.80	0.13
CYP	7.71	2.49	1.78	-1.29	7.00	2.85	1.72	-0.60	4.32	1.82	-0.87	-0.96
AVG	**2.39**	**0.80**	**-0.09**	**-2.46**	**1.59**	**0.65**	**-0.75**	**-2.25**	**1.05**	**0.42**	**-0.24**	**-1.12**

The table summarizes the dummies δ_m from regression (35) for varying aggregation levels. The *Fixed Productivity* columns show the respective values when productivity of entering firms is held constant. The numbers in the *Variable Productivity* columns refer to the model that allows for variation in productivity across firms and between entering and incumbent domestic firms.

Table B5: Robustness: Fixed Effect Regressions

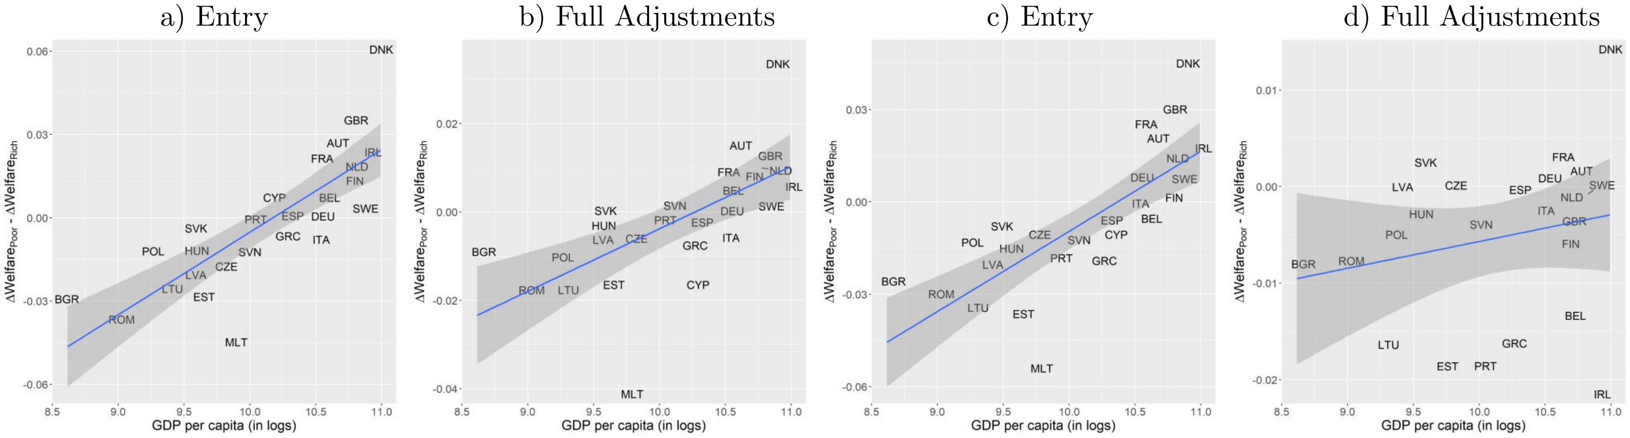

a) Entry b) Full Adjustments c) Entry d) Full Adjustments

	Fixed Productivity					Variable Productivity					
	Gap			Gap			Gap			Gap	
	Entry	Full		Entry	Full		Entry	Full		Entry	Full
FRA	2.389	0.990	AUT	2.949	1.667	FRA	2.279	0.198	AUT	1.864	0.090
NLD	2.165	0.997	MLT	-4.736	-3.959	NLD	1.660	-0.182	MLT	-5.170	-3.124
GER	0.323	0.190	EST	-2.580	-1.487	GER	0.570	0.161	EST	-3.406	-1.934
ITA	-0.527	-0.421	LVA	-1.908	-0.785	ITA	-0.277	-0.167	LVA	-1.802	0.072
UK	3.254	1.113	LTU	-2.810	-1.597	UK	2.777	-0.274	LTU	-3.193	-1.562
IRL	2.622	0.415	POL	-1.474	-0.861	IRL	1.494	-2.075	POL	-1.552	-0.575
DNK	5.812	3.522	CZE	-1.497	-0.436	DNK	4.727	1.349	CZE	-1.390	0.089
GRC	-0.919	-0.596	SVK	-0.636	-0.139	GRC	-1.675	-1.697	SVK	-1.025	0.174
PRT	-0.319	-0.347	HUN	-1.436	-0.470	PRT	-1.594	-1.932	HUN	-1.742	-0.356
ESP	-0.203	-0.070	ROM	-3.915	-1.932	ESP	-0.353	-0.107	ROM	-3.233	-0.847
BEL	0.989	0.416	BGR	-2.646	-1.059	BEL	-0.310	-1.411	BGR	-2.803	-0.874
LUX	0.887	-0.631	SVN	-0.977	-0.014	LUX	-2.587	-4.370	SVN	-1.460	-0.316
SWE	0.609	0.293	CYP	0.459	-1.810	SWE	0.518	-0.074	CYP	-1.248	-4.014
FIN	1.074	1.038				FIN	-0.095	-0.516			

The above figures and table show the dummies δ_m from regression (36) under inclusion of product fixed effects. The 2 figures on the left and the *Fixed Productivity* columns show the respective values when productivity of entering firms is held constant. The figures and results on the right hand side refer to the model that allows for variation in productivity across firms and between entering and incumbent domestic firms. Outlier Luxembourg is not shown in the plots.

C Derivations

C.1 Household Decision

C.1.1 Variety Choice

This part shows that when household utility for variety j is given by

$$u_{jk}^{(i)} = x_{ijk} e^{\frac{q_{jk}+\varepsilon_{ijk}}{\alpha(z_i)}},$$

the resulting probability that a household with service consumption z_i will buy variety j^* will be

$$\Pr(i \mapsto j^*) = \frac{\exp[q_{j^*k} - \alpha(z_i)\ln p_{j^*k}]}{\sum_{j\in J_k}\exp[q_{jk} - \alpha(z_i)\ln p_{jk}]}$$

when ε_{ijk} follows a type 1 extreme value distribution. The proof largely follows Handbury (2013). Generally, consumers will choose variety j^* if

$$\begin{aligned}
u_{j^*k}^{(i)} &\geq u_{jk}^{(i)} \\
x_{j^*k} e^{\frac{q_{j^*k}+\varepsilon_{ij^*k}}{\alpha(z_i)}} &\geq x_{jk} e^{\frac{q_{jk}+\varepsilon_{ijk}}{\alpha(z_i)}} \\
E_{ij^*k}\frac{e^{\frac{q_{j^*k}+\varepsilon_{ij^*k}}{\alpha(z_i)}}}{p_{j^*k}} &\geq E_{ijk}\frac{e^{\frac{q_{jk}+\varepsilon_{ijk}}{\alpha(z_i)}}}{p_{jk}}
\end{aligned}$$

where E_{ijk} denotes the expenditure which a consumer spends on variety j. The term $\exp[q_{jk} + \varepsilon_{ijk}/\alpha(z_i)]/p_{jk}$ represents the utility per dollar a household receives when it chooses variety j. Ultimately households will want to choose the variety which maximizes this utility per dollar and it will be unnecessary to explicitly need to keep track of expenditures E_{ijk}. In addition, given that utility over different product categories is of Cobb-Douglas form, the optimal expenditure on a product will in any case be independent of the optimal variety and hence $E_{ijk} = E_{ik}$, $\forall j \in J_k$. Hence, taking logs, consumers will optimally choose variety j^* if

$$\begin{aligned}
\frac{q_{j^*k}+\varepsilon_{ij^*k}}{\alpha(z_i)} - \ln p_{j^*k} \geq \max_{j\in J_k}\left(\frac{q_{jk}+\varepsilon_{ijk}}{\alpha(z_i)} - \ln p_{jk}\right) \\
\Leftrightarrow q_{j^*k}+\varepsilon_{ij^*k} - \alpha(z_i)\ln p_{j^*k} \geq \max_{j\in J_k}\left(q_{jk}+\varepsilon_{ijk} - \alpha(z_i)\ln p_{jk}\right).
\end{aligned}$$

In order to derive the optimal decision rule, first note that if ε_{ijk} follows a type 1 extreme value distribution, then

$$\varepsilon_{ijk} + q_{jk} - \alpha(z_i)\ln p_{jk}$$

follows the same distribution with location parameter $v_j := q_{jk} - \alpha(z_i)\ln p_{jk}$. Further, the maximum over N T1EV distributed variables u_j with location parameters v_j is again T1EV distributed, since

$$\Pr(\max\{u_j\} < x) = \prod_{j=1}^{N}\Pr(u_j < x)$$

$$= \prod_{j=1}^{N} e^{-e^{-(x-v_j)}} = e^{-\sum_{j=1}^{N} e^{-(x-v_j)}} = e^{-\sum_{j=1}^{N} e^{-x+v_j}}$$

$$
\begin{aligned}
&= e^{-e^{-x}\sum_{j=1}^{N} e^{v_j}} = e^{-e^{-x} e^{\left[\ln\left(\sum_{j=1}^{N} e^{v_j}\right)\right]}} \\
&= e^{-e^{-x}e^{v}} = e^{-e^{-(x-v)}}.
\end{aligned}
$$

where v is the new location parameter with

$$
v = \ln\left(\sum\nolimits_{j=1}^{N} e^{v_j}\right).
$$

The final useful property when deriving the choice probabilities is that the difference between two T1EV distributed random variables $X \sim T1EV(v_j)$ and $Y \sim T1EV(v)$ follows a logistic distribution with location parameter $\mu = v_j - v$. Using this property, we can then derive the household's choice probability as:

$$
\begin{aligned}
\Pr(i \mapsto j^*) &= \Pr\left(\varepsilon_{ij^*k} + q_{j^*k} - \alpha(z_i)\ln p_{j^*k} \geq \max_{j \neq j^*}\{\varepsilon_{ijk} + q_{jk} - \alpha(z_i)\ln p_{jk}\}\right) \\
&= \Pr\left(\max_{j \neq j^*}\{\varepsilon_{ijk} + q_{jk} - \alpha(z_i)\ln p_{jk}\} - (\varepsilon_{ij^*k} + q_{j^*k} - \alpha(z_i)\ln p_{j^*k}) \leq 0\right) \\
&= \frac{1}{1 + \exp[v - v_{j^*}]} \\
&= \frac{1}{1 + \exp\left[\ln\left(\sum_{j \neq j^*} e^{q_{jk} - \alpha(z_i)\ln p_{jk}}\right) - (q_{j^*k} - \alpha(z_i)\ln p_{j^*k})\right]} \\
&= \frac{\exp[q_{j^*k} - \alpha(z_i)\ln p_{j^*k}]}{\sum_{j \in J_k} \exp[q_{jk} - \alpha(z_i)\ln p_{jk}]}
\end{aligned}
$$

which is the familiar logit expression. Line 3 uses the cdf of the logistic distribution: $F(x; \tilde{\mu}, \tilde{\sigma}) = 1/(1 + \exp(-(x - \tilde{\mu})/\tilde{\sigma}))$

C.2 Firm Choices

C.2.1 Prices

Firm profits are given by

$$
\begin{aligned}
\pi_{jm} &= [p_{jm} - mc_{jm}(q_{jm})]\frac{s_{jm}}{p_{jm}}E_m - f \\
&= [p_{jm} - mc_{jm}(q_{jm})]\sum_{i \in I_m} \frac{E_{im}}{E_m} s_{jm}^{(i)} \frac{E_m}{p_{jm}} - f \\
&= [p_{jm} - mc_{jm}(q_{jm})]\sum_{i \in I_m} \frac{E_{im}}{E_m} \frac{\exp[q_{jm} - \alpha(z_i)\ln p_{jm}]}{\sum_{j' \in J_k} \exp[q_{j'm} - \alpha(z_i)\ln p_{j'm}]} \frac{E_m}{p_{jm}} - f \\
&= \frac{p_{jm} - mc_{jm}(q_{jm})}{p_{jm}} \sum_{i \in I_m} E_{im} \frac{\exp[q_{jm} - \alpha(z_i)\ln p_{jm}]}{\sum_{j' \in J_k} \exp[q_{j'm} - \alpha(z_i)\ln p_{j'm}]} - f.
\end{aligned}
$$

where line 3 uses the closed-form solution for the probability that household i chooses variety j which equals

$$
s_{jm}^{(i)} = \frac{\exp[q_{jm} - \alpha(z_i)\ln p_{jm}]}{\sum_{j' \in J_k} \exp[q_{j'm} - \alpha(z_i)\ln p_{j'm}]}.
$$

The partial derivative of profits with respect to price p_{jm} is then

$$\begin{aligned}
\frac{\partial \pi_{jm}}{\partial p_{jm}} &= \frac{mc_{jm}(q_{jm})}{p_{jm}^2} \sum_{i \in I_m} E_{im} \frac{\exp[q_{jm} - \alpha(z_i) \ln p_{jm}]}{\sum_{j' \in J_k} \exp[q_{j'm} - \alpha(z_i) \ln p_{j'm}]} \\
&+ \frac{p_{jm} - mc_{jm}(q_{jm})}{p_{jm}} \sum_{i \in I_m} E_{im} \frac{-\frac{\alpha(z_i)}{p_{jm}} \exp[q_{jm} - \alpha_i \ln p_{jm}] \sum_{j' \in J_k} \exp[q_{j'm} - \alpha_i \ln p_{j'm}] + \frac{\alpha_i}{p_{jm}} \exp[q_{jm} - \alpha_i \ln p_{jm}]^2}{\left(\sum_{j' \in J_k} \exp[q_{j'm} - \alpha_i \ln p_{j'm}]\right)^2} \\
&= \frac{mc_{jm}(q_{jm})}{p_{jm}^2} \sum_{i \in I_m} E_{im} s_{jm}^{(i)} \\
&- \frac{p_{jm} - mc_{jm}(q_{jm})}{p_{jm}} \sum_{i \in I_m} E_{im} \frac{\alpha(z_i)}{p_{jm}} \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right].
\end{aligned}$$

The first-order condition with respect to the price then implies for the percentage markup:

$$\begin{aligned}
&\frac{\partial \pi_{jm}}{\partial p_{jm}} = 0 \\
\Leftrightarrow\ & \frac{mc_{jm}(q_{jm})}{p_{jm}^2} \sum_{i \in I_m} E_{im} s_{jm}^{(i)} = \frac{p_{jm} - mc_{jm}(q_{jm})}{p_{jm}} \sum_{i \in I_m} E_{im} \frac{\alpha_i}{p_{jm}} \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right] \\
\Leftrightarrow\ & mc_{jm}(q_{jm}) \sum_{i \in I_m} E_{im} s_{jm}^{(i)} = (p_{jm} - mc_{jm}(q_{jm})) \sum_{i \in I_m} E_{im} \alpha_i \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right] \\
\Leftrightarrow\ & \frac{p_{jm} - mc_{jm}(q_{jm})}{mc_{jm}(q_{jm})} = \frac{\sum\limits_{i \in I_m} E_{im} s_{jm}^{(i)}}{\sum\limits_{i \in I_m} E_{im} \alpha_i \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right]} \\
\Leftrightarrow\ & \frac{p_{jm} - mc_{jm}(q_{jm})}{mc_{jm}(q_{jm})} = \frac{\sum\limits_{i \in I_m} E_{im} s_{jm}^{(i)}}{\sum\limits_{i \in I_m} E_{im} \alpha_i s_{jm}^{(i)} \left(1 - s_{jm}^{(i)} \right)}.
\end{aligned}$$

C.2.2 Quality

The first-order condition with respect to quality is

$$\begin{aligned}
\frac{\partial \pi_{jm}}{\partial q_{jm}} &= -\frac{1}{p_{jm}} \frac{\partial mc_{jm}(q_{jm})}{\partial q_{jm}} \sum_{i \in I_m} E_{im} \frac{\exp[q_{jm} - \alpha_i \ln p_{jm}]}{\sum_{j' \in J_k} \exp[q_{j'm} - \alpha_i \ln p_{j'm}]} \\
&+ \frac{p_{jm} - mc_{jm}(q_{jm})}{p_{jm}} \sum_{i \in I_m} E_{im} \frac{\exp[q_{jm} - \alpha_i \ln p_{jm}] \sum_{j' \in J_k} \exp[q_{j'm} - \alpha_i \ln p_{j'm}] - (\exp[q_{jm} - \alpha_i \ln p_{jm}])^2}{\left(\sum_{j' \in J_k} \exp[q_{j'm} - \alpha_i \ln p_{j'm}]\right)^2} \\
&= -\frac{1}{p_{jm}} \frac{\partial mc_{jm}(q_{jm})}{\partial q_{jm}} \sum_{i \in I_m} E_{im} s_{jm}^{(i)} \\
&+ \frac{p_{jm} - mc_{jm}(q_{jm})}{p_{jm}} \sum_{i \in I_m} E_{im} \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right] = 0
\end{aligned}$$

and so

$$\begin{aligned}
& \frac{\partial mc_{jm}(q_{jm})}{\partial q_{jm}} \sum_{i \in I_m} E_{im} s_{jm}^{(i)} = (p_{jm} - mc_{jm}(q_{jm})) \sum_{i \in I_m} E_{im} \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right] \\
\Leftrightarrow\ & \frac{\partial mc_{jm}(q_{jm})}{\partial q_{jm}} \sum_{i \in I_m} \frac{E_{im}}{E_m} s_{jm}^{(i)} = (p_{jm} - mc_{jm}(q_{jm})) \sum_{i \in I_m} \frac{E_{im}}{E_m} \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right]
\end{aligned}$$

$$\Leftrightarrow \quad \frac{\partial mc_{jm}(q_{jm})}{\partial q_{jm}} s_{jm} = (p_{jm} - mc_{jm}(q_{jm})) \sum_{i \in I_m} \frac{E_{im}}{E_m} \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right].$$

Intuitively, the left-hand side describes the additional cost of increasing quality and the right-hand side the additional profit through a greater number of units sold.

C.2.3 Solution

In summary, the firm's first-order conditions are

$$\frac{p_{jm} - mc_{jm}(q_{jm})}{mc_{jm}(q_{jm})} = \frac{\sum\limits_{i \in I_m} E_{im} s_{jm}^{(i)}}{\sum\limits_{i \in I_m} E_{im}(\alpha_i) s_{jm}^{(i)} \left(1 - s_{jm}^{(i)} \right)}$$

$$\frac{\partial mc_{jm}(q_{jm})}{\partial q_{jm}} s_{jm} = (p_{jm} - mc_{jm}(q_{jm})) \sum_{i \in I_m} \frac{E_{im}}{E_m} \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right]$$

where

$$s_{jm}^{(i)} = \frac{\exp[q_{jm} - \alpha(z_i) \ln p_{jm}]}{\sum_{j' \in J_k} \exp[q_{j'm} - \alpha(z_i) \ln p_{j'm}]}.$$

Now suppose each firm from country c that sells to country m makes the same choices in terms of p_{jm} and q_{jm}, i.e.

$$p_{jm} = p_{cm}, \quad \forall j \in c$$
$$q_{jm} = q_{cm}, \quad \forall j \in c$$

In that case, the expenditure share of country c in market m for consumer group i becomes

$$s_{cm}^{(i)} = \sum_{j \in c} s_{jm}^{(i)} = \sum_{j \in c} \frac{\exp[q_{jm} - \alpha(z_i) \ln p_{jm}]}{\sum_{j' \in J_k} \exp[q_{j'm} - \alpha(z_i) \ln p_{j'm}]}.$$
$$= \frac{N_{cm} \exp[q_{cm} - \alpha(z_i) \ln p_{cm}]}{\sum_{c' \in J_k} \exp[q_{c'm} - \alpha(z_i) \ln p_{c'm}]}$$

and so the share of an individual firm from c becomes

$$s_{c^*m}^{(i)} = \frac{s_{cm}^{(i)}}{N_{cm}} = \frac{\exp[q_{cm} - \alpha(z_i) \ln p_{cm}]}{\sum_{c' \in J_k} \exp[q_{c'm} - \alpha(z_i) \ln p_{c'm}]}.$$

The first-order conditions can hence be written as

$$\frac{p_{cm} - mc_{cm}(q_{cm})}{mc_{cm}(q_{cm})} = \frac{\sum\limits_{i \in I_m} E_{im} s_{cm}^{(i)} / N_{cm}}{\sum\limits_{i \in I_m} E_{im}(\alpha_i) s_{cm}^{(i)} / N_{cm} \left(1 - s_{cm}^{(i)} / N_{cm} \right)} \tag{37}$$

$$\frac{\partial mc_{cm}(q_{cm})}{\partial q_{cm}} s_{cm} = (p_{cm} - mc_{cm}(q_{cm})) \sum_{i \in I_m} \frac{E_{im}}{E_m} \left[s_{cm}^{(i)} / N_{cm} - \left(s_{cm}^{(i)} / N_{cm} \right)^2 \right]. \tag{38}$$

Now suppose that $\tilde{j} = 1, ..., N_{cm} - 1$ firms choose p_{cm} and q_{cm} but firm $j = N_{cm}$ potentially deviates.

Also in this case, the expenditure share of each individual firm that chooses p_{cm} and q_{cm} is

$$s_{\tilde{j}m}^{(i)} = \frac{1}{N_{cm}-1} \frac{(N_{cm}-1)\exp[q_{cm} - \alpha(z_i)\ln p_{cm}]}{\sum_{c' \in J_k} \exp[q_{c'm} - \alpha(z_i)\ln p_{c'm}]} = \frac{\exp[q_{cm} - \alpha(z_i)\ln p_{cm}]}{\sum_{c' \in J_k} \exp[q_{c'm} - \alpha(z_i)\ln p_{c'm}]}.$$

and so the first-order conditions are the same as those above (37) and (38). The first-order condition of firm j,

$$\frac{p_{jm} - mc_{jm}(q_{jm})}{mc_{jm}(q_{jm})} = \frac{\sum\limits_{i \in I_m} E_{im} s_{jm}^{(i)}}{\sum\limits_{i \in I_m} E_{im}(\alpha_i) s_{jm}^{(i)} \left(1 - s_{jm}^{(i)}\right)}$$

$$\frac{\partial mc_{jm}(q_{jm})}{\partial q_{jm}} s_{jm} = (p_{jm} - mc_{jm}(q_{jm})) \sum_{i \in I_m} \frac{E_{im}}{E_m} \left[s_{jm}^{(i)} - \left(s_{jm}^{(i)} \right)^2 \right].$$

with

$$s_{jm}^{(i)} = \frac{\exp[q_{jm} - \alpha(z_i)\ln p_{jm}]}{\sum_{j' \in J_k} \exp[q_{j'm} - \alpha(z_i)\ln p_{j'm}]}.$$

will then also be satisfied at p_{cm} and q_{cm}.

C.3 Price Index

This section derives price indexes for each consumer i with income y_i, living in country m. First, the indirect utility of a consumer i can be written as

$$\begin{aligned} V^{(i)} &= U_M^{(i)}(\mathbf{x}, z) + u(z^*) \\ &= \sum_k \omega_k \ln x_k^* e^{\frac{q_k + \varepsilon_{ik}}{\widetilde{\alpha}_0 + \widetilde{\alpha}_1 \ln z^*}} + u(z^*) \end{aligned}$$

where $U_M^{(i)}(\mathbf{x}, z)$ and $u(z^*)$ denote utility from manufacturing and services, respectively. x_k^* and z^* are the optimal choices of a household. Replacing x_k^* by (7) and using the normalization for p_z delivers an expression for the indirect utility as function of income y_i, service consumption z^*, prices p_{jk}, and characteristics, q_k and ε_{ik}:

$$V^{(i)} = \sum_k \omega_k \ln \omega_k \frac{y_i - z^*}{p_k} e^{\frac{q_k + \varepsilon_{ik}}{\widetilde{\alpha}_0 + \widetilde{\alpha}_1 \ln z^*}} + u(z^*).$$

I focus on the utility derived from manufacturing goods, $U_M^{(i)}$: To construct the price index, set $U_M^{(i)} = 1$ and solve for the corresponding necessary expenditure on manufacturing goods $y_i - z^*$:

$$\begin{aligned} 1 &= \sum_k \omega_k \ln \left(\omega_k \frac{y_i - z^*}{p_k} e^{\frac{q_k + \varepsilon_{ik}}{\widetilde{\alpha}_0 + \widetilde{\alpha}_1 \ln z^*}} \right) \\ \Leftrightarrow \ 1 &= \sum_k \omega_k \ln \frac{\omega_k}{p_k} + \sum_k \omega_k \ln(y_i - z^*) + \sum_k \omega_k \frac{q_k + \varepsilon_{ik}}{\widetilde{\alpha}_0 + \widetilde{\alpha}_1 \ln z^*} \\ \Leftrightarrow \ \ln(y_i - z^*) &= 1 - \sum_k \omega_k \ln \frac{\omega_k}{p_k} - \sum_k \omega_k \frac{q_k + \varepsilon_{ik}}{\widetilde{\alpha}_0 + \widetilde{\alpha}_1 \ln z^*} \end{aligned}$$

$$\Leftrightarrow \ \mathbb{P}_m(z^*) = \exp\left(\sum_k \omega_k \left[1 - \ln\frac{\omega_k}{p_k} - \frac{q_k + \varepsilon_{ik}}{\widetilde{\alpha}_0 + \widetilde{\alpha}_1 \ln z^*}\right]\right)$$

$$\Leftrightarrow \ \mathbb{P}_m(y_i) = \exp\left(\sum_k \omega_k \left[1 - \ln\frac{\omega_k}{p_k} - \frac{q_k + \varepsilon_{ik}}{\alpha_0 + \alpha_1 \ln y_i}\right]\right)$$

where the last step follows from the functional form assumptions described in section 4.2.1. Since $\alpha_0 + \alpha_1 \ln y_i$ will be positive (since the price coefficient will be negative), it can be easily shown that

$$\frac{\partial \mathbb{P}_m}{\partial p_k} > 0, \ \frac{\partial \mathbb{P}_m}{\partial q_k} < 0, \ \frac{\partial \mathbb{P}_m}{\partial \varepsilon_k} < 0.$$

The optimal price index is hence increasing in the price of the optimal variety and decreasing in their respective characteristics. Notice that the price index captures both changes along the intensive as well as the extensive margin: As new varieties become available, the chosen variety k^* may change and with it the respective p_{k^*}, q_{k^*}, and ε_{k^*}.

C.4 Hidden Varieties

This part shows that equation (ES) can also be derived through an approximation. Specifically, the expenditure share in market m of exporters from origin c can be approximated through first-order Taylor expansions as

$$\begin{aligned}
s_{cm} &= \sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\sum_{j\in J_c} \exp(q_{jm} - \alpha_i \ln p_{jm})}{\sum_{j'} \exp(q_{j'm} - \alpha_i \ln p_{j'm})} \\
&\approx \sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\exp(\overline{q}_{cm} - \alpha_i \ln \overline{p}_{cm}) \sum_{j\in J_c} \left[1 + (q_{jm} - \overline{q}_{cm}) - \frac{\alpha_i}{\overline{p}_{jm}}(p_{jm} - \overline{p}_{cm})\right]}{\sum_{c'} \exp(\overline{q}_{c'm} - \alpha_i \ln \overline{p}_{c'm}) \sum_{j'\in J_{c'}} \left[1 + (q_{j'm} - \overline{q}_{c'm}) - \frac{\alpha_i}{\overline{p}_{j'm}}(p_{j'm} - \overline{p}_{c'm})\right]} \\
&= \sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{N_{cm} \exp(\overline{q}_{cm} - \alpha_i \ln \overline{p}_{cm})}{\sum_{c'} N_{c'm} \exp(\overline{q}_{c'm} - \alpha_i \ln \overline{p}_{c'm})} \\
&= \sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\exp(\overline{q}_{cm} + N_{cm} - \alpha_i \ln \overline{p}_{cm})}{\sum_{c'} \exp(\overline{q}_{c'm} + N_{c'm} - \alpha_i \ln \overline{p}_{c'm})}.
\end{aligned}$$

A valid instrument for the average price $\overline{p}_{cm}$ is hence sufficient to identify α_i up to a first-order approximation.

C.5 Own-Price Elasticities

The elasticity of firm $j's$ sold quantity in market m, x_{jm}, with respect to its own price is

$$\begin{aligned}
\frac{\partial x_{jm}}{\partial p_{jm}} \frac{p_{jm}}{x_{jm}} &= \frac{\partial (s_{jm} E_m / p_{jm})}{\partial p_{jm}} \frac{p_{jm}}{s_{jm} E_m / p_{jm}} \\
&= \frac{\partial}{\partial p_{jm}} \left(\sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\exp(q_{jm} - \alpha_i \ln p_{jm})}{\sum_{j'} \exp(q_{j'm} - \alpha_i \ln p_{j'm})} \frac{E_m}{p_{jm}}\right) \frac{p_{jm}^2}{s_{jm} E_m} \\
&= \frac{\partial}{\partial p_{jm}} \left(\sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\exp(q_{jm} - (\alpha_i + 1) \ln p_{jm})}{\sum_{j'} \exp(q_{j'm} - \alpha_i \ln p_{j'm})}\right) \frac{p_{jm}^2}{s_{jm}}
\end{aligned}$$

$$
\begin{aligned}
&= -\sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\frac{\alpha_i+1}{p_{jm}} \exp(q_{jm} - (\alpha_i+1)\ln p_{jm}) \sum_{j'} \exp(q_{j'm} - \alpha_i \ln p_{j'm})}{\left(\sum_{j'} \exp(q_{j'm} - \alpha_i \ln p_{j'm})\right)^2} \frac{p_{jm}^2}{s_{jm}} \\
&\quad + \sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\frac{\alpha_i}{p_{jm}} \exp(q_{jm} - \alpha_i \ln p_{jm}) \exp(q_{jm} - (\alpha_i+1)\ln p_{jm})}{\left(\sum_{j'} \exp(q_{j'm} - \alpha_i \ln p_{j'm})\right)^2} \frac{p_{jm}^2}{s_{jm}} \\
&= -\sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\frac{\alpha_i+1}{p_{jm}^2} \exp(q_{jm} - \alpha_i \ln p_{jm}) \sum_{j'} \exp(q_{j'm} - \alpha_i \ln p_{j'm})}{\left(\sum_{j'} \exp(q_{j'm} - \alpha_i \ln p_{j'm})\right)^2} \frac{p_{jm}^2}{s_{jm}} \\
&\quad + \sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\frac{\alpha_i}{p_{jm}^2} \exp(q_{jm} - \alpha_i \ln p_{jm}) \exp(q_{jm} - \alpha_i \ln p_{jm})}{\left(\sum_{j'} \exp(q_{j'm} - \alpha_i \ln p_{j'm})\right)^2} \frac{p_{jm}^2}{s_{jm}} \\
&= -\sum_{i\in I_m} \frac{E_{im}}{E_m} \left[(\alpha_i+1) s_{jm}^{(i)} - \alpha_i \left(s_{jm}^{(i)}\right)^2\right] \frac{1}{s_{jm}} \\
&= -\left(1 + \sum_{i\in I_m} \frac{E_{im}}{E_m} \left[\alpha_i s_{jm}^{(i)} - \alpha_i \left(s_{jm}^{(i)}\right)^2\right] \frac{1}{s_{jm}}\right) \\
&= -\left(1 + \sum_{i\in I_m} \frac{E_{im}}{E_m} \frac{\alpha_i s_{jm}^{(i)} (1 - s_{jm}^{(i)})}{s_{jm}}\right)
\end{aligned}
$$

D Data and Reduced Form Evidence

D.1 Production Data

Domestic production data is taken from Eurostat, the statistical agency of the European Union, for 27 countries. Since 1995, Eurostat collects harmonized information on sold as well as physical production of its member states at a high level of disaggregation, the so-called prodcom-classification. There are close to 4,000 prodcom product categories, classified by 8-digit identifiers.

A major reason why the EU collects this information was to allow trade and production data being linked after the creation of a single market in 1992. For that reason, there is considerable overlap between 8-digit Prodcom and 8-digit Combined Nomenclature categories, which allows a match between the available trade and production data.[1] I use data on sold instead of total production since the latter consists of production used for sale as well as for inputs within the company and because sold production is reported both in terms of values (in Euros) as well as volume (weight, units, etc.). The availability of both allows me to construct unit values as proxies for prices, just as with trade data.

The availability of production data on such a disaggregated level comes at the cost of some data being marked as confidential by Eurostat. This is for example the case if there are only a few firms producing within an industry and a firm's production could in principle be inferred from the more aggregate data.[2] I impute data that is marked as confidential or missing using other information on a country's imports, exports, and production in the same industry, or closely related industries. An

---

[1] Eurostat states for example that "Prodcom codes normally relate to one or more Combined Nomenclature headings, thus enabling external trade data to be related to production data" (See Eurostat - Statistics on the production of manufactured goods (prom); Reference Metadata in Euro SDMX Metadata Structure (ESMS)).

[2] See Williams (2008) for a more detailed description of data collection and reporting.

alternative would obviously be to delete those observations and compute price indexes only based on product categories without confidential or missing data. I chose to impute this data instead mainly to make the final results as representative as possible, given the available data.

Table D1: Regression Results - Production Value

	Dependent Variable: Production Value (logs)				
	(1)		(2)		(3)
Export Value (logs)	0.378^{***}	(6.05)	0.558^{***}	(17.27)	
Import Value (logs)	0.067^{**}	(2.98)	0.112^{***}	(5.43)	
Production Quantity (logs)	0.338^{***}	(4.95)			
Country FEs	Yes		Yes		Yes
3-Digit Industry FEs	Yes		Yes		Yes
Observations	32,562		40,025		47,224
R^2	0.816		0.725		0.552

Standard errors are clustered by 3-digit industry. t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Specifically, to predict missing information on country $c's$ production values in 8-digit industry k, I run versions of the regression

$$\begin{aligned}\ln(\text{Prod. Value}_{c,k}) &= \alpha_0 \cdot \ln(\text{Export Value}_{c,k}) + \alpha_1 \cdot \ln(\text{Import Value}_{c,k}) + \alpha_2 \cdot \ln(\text{Prod. Quantity}_{c,k}) \\ &\quad + \beta_c \cdot \mathbb{I}\{\text{Country} = c\} + \gamma_k \mathbb{I}\{\text{Product k} \in K\}, \end{aligned} \tag{39}$$

on the non-confidential data and use the estimated coefficients to predict the non-reported values. I use production quantities on the right-hand side since they are correlated with production values by definition. Similarly, I expect export values to be positively correlated with production. K denotes the 3-digit industry which k belongs to and the inclusion of 3-digit industry fixed effects captures the idea that a country which is productive in producing e.g. compact cars will plausibly also have the set of skills and capital to produce middle-sized cars or other closely related vehicles. Therefore, one would expect a country's production in relatively similar industries to be positively correlated.

I replace confidential and missing values in 3 steps. First, I run regression (39) and replace missing observations with the predicted ones from this regression. In cases in which also information on production quantity is not available, I use regression (39) with α_2 set to 0 instead. Finally, in cases in which data on quantities as well as trade values is missing, I predict production values based on country and industry fixed effects alone.

Table D1 summarizes the results of the 3 regressions described above. As expected, export values as well as production quantities show a strong positive correlation with production values. Import values are also positively correlated but have comparably less explanatory power for production values. The R^2 of each regression is quite high with values of 0.82 and 0.73, when production values are explained by quantities and trade flows, and 0.55 if only fixed effects are included.
In order to infer missing quantity data, I analogously use the regression

$$\begin{aligned}\ln(\text{Prod. Quantity}_{c,k}) &= \alpha_0 \cdot \ln(\text{Export Value}_{c,k}) + \alpha_1 \cdot \ln(\text{Import Value}_{c,k}) \\ &\quad + \alpha_2 \cdot \ln(\text{Prod. Value}_{c,k}) + \beta_c \cdot \mathbb{I}\{\text{Country} = c\} + \gamma_k \mathbb{I}\{\text{Product k} \in K\},\end{aligned}$$

Table D2: Regression Results - Production Quantity

	Dependent Variable: Production Quantity (logs)				
	(1)		(2)		(3)
Export Quantity (logs)	0.329***	(6.80)	0.586***	(21.39)	
Import Quantity (logs)	0.218***	(4.69)	0.192***	(8.78)	
Production Value (logs)	0.667***	(11.70)			
Country FEs	Yes		Yes		Yes
3-Digit Industry FEs	Yes		Yes		Yes
Observations	25,288		25,742		37,989
R^2	0.893		0.828		0.654

Standard errors are clustered by 3-digit industry. t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

to infer confidential values. As summarized in Table D2, the results are largely consistent with those for production values: Production quantities are positively correlated with values and traded quantities. The R^2 is slightly higher in each of these regressions relative to their counterpart in Table D1. Figure D1 also summarizes the fit of the above regressions graphically.

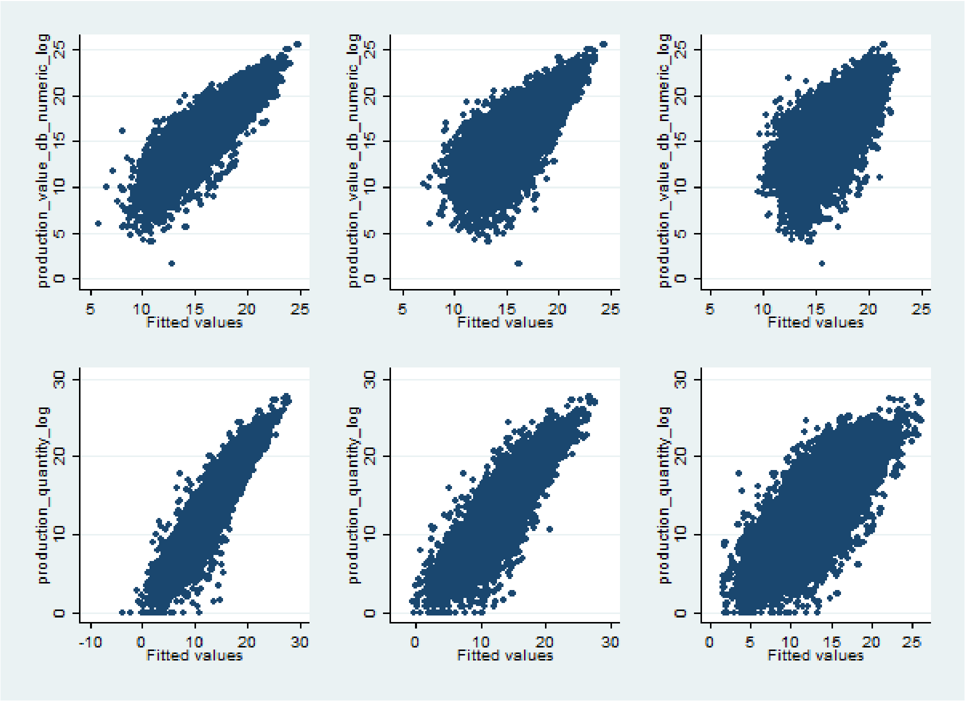

Figure D1: **Actual vs fitted values**: The top 3 figures plot fitted against actual production values (in logs) for the 3 regressions discussed above. The bottom 3 figures plot fitted against actual production quantities.

D.2 Matching Trade and Production Data

I match trade and production data using the concordance developed by van Beveren, Bernard, and Vandenbussche (2012). Such a concordance is necessary for 2 reasons: First, trade flows are reported using the so-called Combined Nomenclature (cn) classification while production data is classified using prodcom codes (pc). Second, these classifications are frequently revised over time, for example because of the introduction of new goods or some products becoming less importing over time or obsolete. The concordance constructs product categories that are consistent over time by matching categories that have been renamed or aggregating several categories into a single one that is robust to changes over time.

Since the prodcom dataset covers only the manufacturing sector, there are naturally some cn8 codes that do not match into pc8 codes, especially in the agricultural sector. I delete those products from the dataset. Additionally, some cn8 categories match into 2 or 3 pc8 categories (In 2003, for example, 490 out 9,532 cn8-categories matched into 2 pc8 categories and 5 into 3). In those cases, I distribute the respective trade flows equally onto the corresponding prodcom categories.

For the construction of unit values, it is also important to keep track of the units in which both trade and production data are reported. Eurostat reports the data on imported and exported quantities in tons as well as a supplementary unit (e.g. the number of items or liters) which varies by product. Production data on the other hand, while in most cases also being reported in weight, is frequently reported in other units only. This can create inconsistencies when constructing unit values if trade and production data are reported in different units.

In many cases, I am able to resolve these inconsistencies in a straightforward way, e.g. if trade data is reported in number of items and the production data in dozens. I delete observations in which this is not possible. A similar problem arises when multiple cn8 categories match into a pc8 category but have different units and I also delete those observations.

In some cases, a country either does not import a product at all or does not produce it. While the first case is unproblematic for the counterfactual of shutting down trade (the price index will be unaffected), the second one is more complicated since it is not obvious if and with which production technology the home country could produce the respective product in autarky. To get around this issue, I slightly aggregate the data up so that each country produces each product.

One way to assess the quality of the match is to check if the crosswalked trade data matches the one that was reported by companies in the prodcom files. In the prodcom surveys, in addition to the production data, companies report their total amount of imports and exports by prodcom category. While this does not provide information about the actual trade partner, it allows me to test if these trade flows by pc8 match those that I obtained after using the crosswalk from cn8 to pc8 using trade data.

All in all, it appears that the match is reasonable. Figure D2 shows that actual and predicted imports are close to the 45-degree line. Table D3 confirms this in a regression of the reported import values on the inferred ones (in logs). The slope is 0.91 and the R^2 of the regression is 0.986.

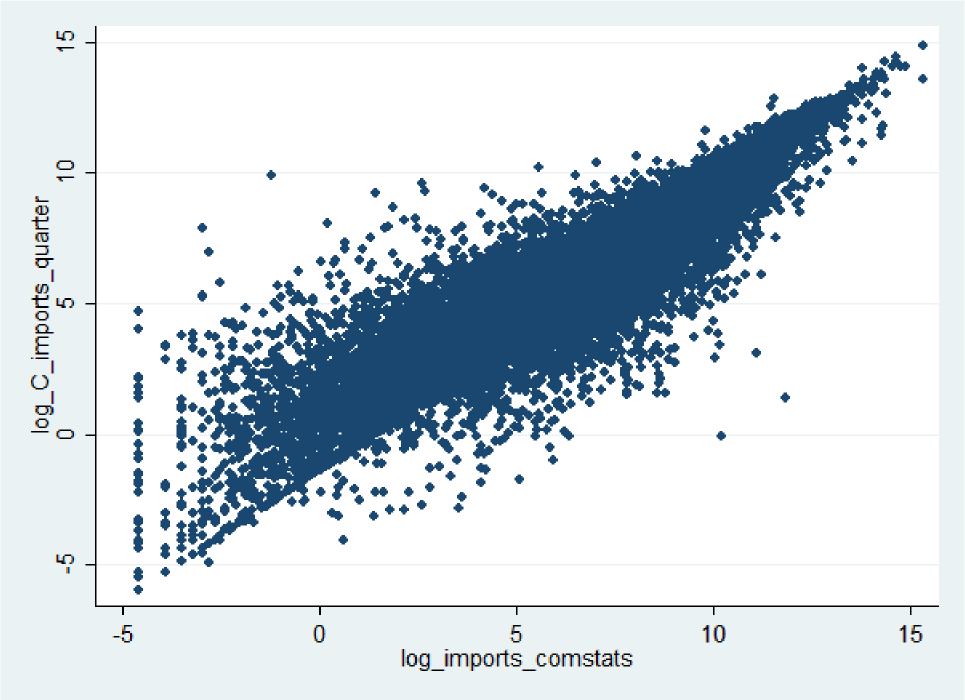

Figure D2: **Reported versus inferred import values**: The graph plots reported import values in the prodcom surveys against the ones inferred from the crosswalked trade data.

Table D3: Regression - Reported versus inferred imports by PC8

	Dependent Variable: Reported imports (in logs)
	(1)
Inferred imports (in logs)	0.910*** (43.01)
Observations	53,802
R^2	0.986

Regression includes importer fixed effects and standard errors are clustered by importer. t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Finally, I match the resulting trade and production data by pc8 codes into pc8plus codes that are consistent over time for the years 2003-2007. In almost 90% of cases a pc8 category did not change over time and no adjustment needed to take place. In the remaining cases, several pc8 categories were aggregated into a pc8plus category: In 6% of cases, 2-3 pc8 were aggregated into a pc8plus category. Only one pc8plus category consisted of more than 10 pc8 categories.

D.3 Outliers and Missing Observations

International trade data is known for being subject to measurement error and the dataset used for this paper is no exception. In order to limit the influence of outliers, I remove unit values that are 30 times bigger or smaller than the median price in a product category. I also remove observations for which a positive quantity is reported but a zero value and vice versa. In some cases, the export share is unrealistically large and I delete cases in which it is greater than 1. I also delete cases in which exports are greater than 0 but imports and production are 0.

For some countries, information on demographics is not available for all years. Population data is frequently not available for sparsely inhabited islands and non-specified territories for example and I deleted the respective observations. Similarly, data on GDP per capita could not be obtained for about 30 countries.[3] All in all, these countries account for only just about 0.05% of exports to the EU and it seems unlikely that they affect the results significantly.

In some cases, the Hausman-type instruments cannot be constructed because a country sells only in a single market. Since this concerns only a small volume of total EU imports (0.02%), I excluded these observations when estimating the model with these instruments. I did however keep them when using the alternative instruments. Since the results did not depend significantly on the instrument used, and given the small amount of trade affected, it appears unlikely that this deletion is driving the results in a major way.

D.4 Income Distributions

As stated in the main text, in order to approximate the income distribution in each country, I use information on disposable household income by decile, quartile and the five highest and lowest percentiles for each country and year. I fit these numbers using a log - normal distribution with country-time specific location and scale parameters and minimize the distance between actual and simulated quantiles, i.e.

$$\min_{\mu_{c,t},\sigma_{c,t}} \left(\sum_q |Q(Y_{c,t}^{(q)}) - \hat{Q}(Y_{c,t}^{(q)}(\mu_{c,t}, \sigma_{c,t}))| \right)$$

where Q is the quantile function and q denotes the respective quantile. Figure D3 plots the inferred log-normal distributions along with the quantiles given in the data for a selection of countries in 2006. All in all, the log-normal-assumption appears to result in a good fit, with the simulated quantiles matching the actual ones closely.
Table D4 shows the estimated parameters for each country in 2012. These estimates are largely consistent with two priors one might have regarding them. First, the richer countries in Northern and Western Europe tend to have higher values of μ, implying bigger average incomes. Second, the Scandinavian countries, which are known for being quite equal economies, exhibit among the lowest values for σ.

D.5 Instruments

D.5.1 Feenstra and Romalis (2014) - Instruments

As described in the main text, this instrument makes use of the observation that exporters report free-on-board (FOB) trade values while those reported by the importing country include cost, insurance, and freight (CIF) incurred by the seller. The difference between the corresponding FOB and CIF unit values hence provides a measure of shipping costs and can be used as an instrument

---

[3] These countries are Ceuta, Iraq, Melilla, North Korea, Gibraltar, Nauru, Holy See (Vatican), New Caledonia, Saint Helena, Wallis and Futuna, Somalia, Pitcairn, Mayotte, Northern Marinara Islands, Saint Pierre and Miquelon, French Polynesia, Anguilla, American Samoa, Turks and Caicos Islands, Guam, U.S. Virgin Islands, Cocos/ Keeling Islands, Cayman Islands, Christmas Islands, British Virgin Islands, Norfolk Island, Montserrat, Cook Islands, Netherlands Antilles, Niue, Falkland Islands, Tokelau.

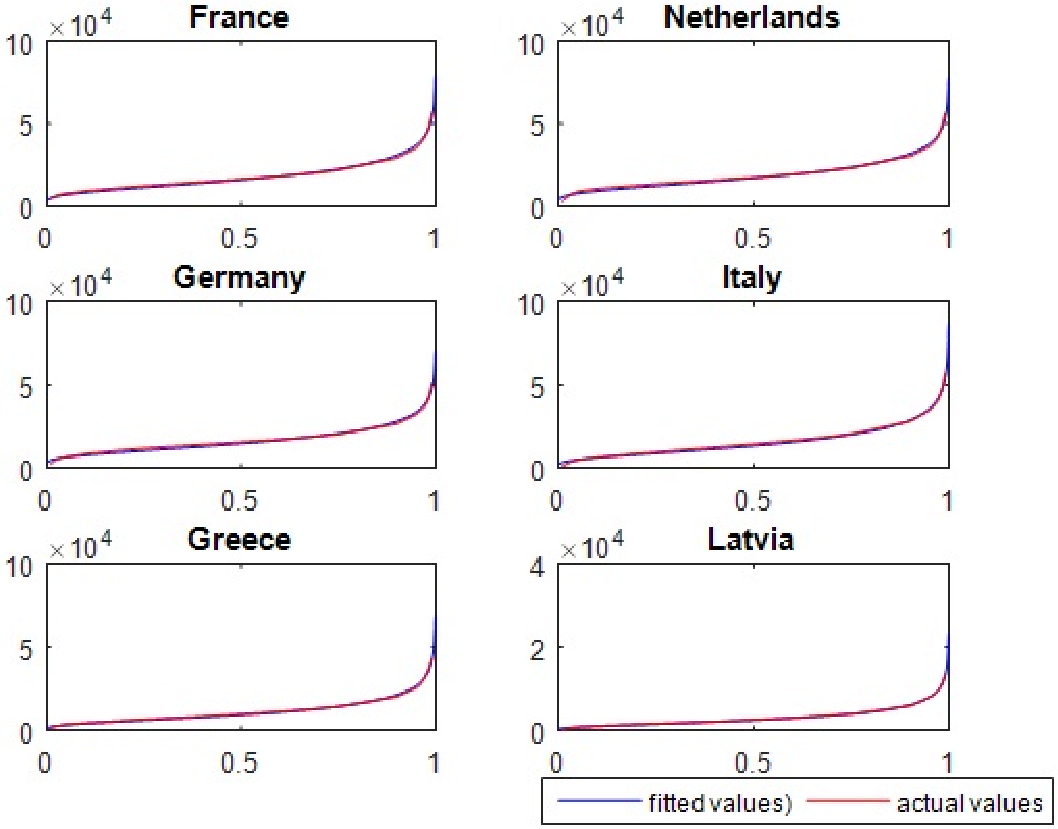

Figure D3: Fitted versus actual income distribution (2006)

for prices charged. In order to construct these instruments, I need to rely on an additional data source for trade flows since the trade flows in the Eurostat data are only reported by the importing country or if the exporting country is an EU member as well. This would not allow me to construct instruments for prices of non-EU exporters and hence for a sizable part of the sample.

Instead, I use data on imports and exports by 6-digit Harmonized System Code (HS6) from the UN Comtrade database. Trade data by HS6 and cn8 an easily be matched to each other since the first 6 digits of the cn classification coincide with those of the Harmonized System.

To construct the instrument, I compute the unit values of each transaction in which an EU country imports a product. Each transaction is reported twice in the Comtrade database: Once as import by the respective EU country and once as export by the trading partner. I define the instrument as the percentage difference between the unit value reported by the importer (uv^{cif}) and that by the exporter (uv^{fob}):

$$\mathbb{Z}^{(1)}_{cmkt} = \ln(uv^{\text{cif}}_{cmkt}) - \ln(uv^{\text{fob}}_{cmkt})$$

where uv_{cmkt} denotes the unit values of exports from country c to m of product k at time t.

Some trade flows are only reported by one trading partner, in which case this instrument cannot be constructed on the HS6 level. The same problem arises in cases in which a trade flow is reported in the Eurostat data but not the Comtrade data. In those cases, I first use unit values on a higher

Table D4: Income Distributions: Parameter Estimates 2012

	μ	σ		μ	σ
Luxembourg	10.382	0.4941	Slovenia	9.3564	0.4277
Denmark	10.081	0.5240	Malta	9.3155	0.4876
Sweden	10.038	0.4614	Portugal	8.9922	0.6724
Finland	9.9988	0.4790	Greece	8.9895	0.6625
Austria	9.9390	0.5273	Czech Republic	8.9524	0.4912
Netherlands	9.9181	0.4763	Slovakia	8.8086	0.4573
France	9.8768	0.6323	Estonia	8.6808	0.5883
Belgium	9.8573	0.4812	Poland	8.4881	0.5949
Germany	9.8495	0.5092	Hungary	8.4302	0.5241
Ireland	9.8270	0.5580	Latvia	8.3695	0.6697
United Kingdom	9.8266	0.5795	Lithuania	8.3618	0.5773
Cyprus	9.7163	0.5749	Bulgaria	7.8616	0.6332
Italy	9.5944	0.5989	Romania	7.6037	0.5594
Spain	9.4695	0.6091			

The graph shows estimates of the parameters of the log-normal distribution which approximates the income distribution in each country in 2012.

level of aggregation (e.g. the 5-digit instead of the 6-digit level) which allows me to resolve about 85% of cases of observations with missing instruments. The remaining cases are those in which a country generally does not report export data (typically smaller countries or dictatorships). In these situations I use the unit values of geographically close countries instead which are plausibly correlated because of similar and similarly long shipping routes, as well as cultural similarities.

To demonstrate that there is a correlation between the instrument and unit values, Table D5 shows the results of regressions of unit values on the instrument and various fixed effects. As expected, there is a positive and significant correlation between cost, freight and insurance (CIF) charges and unit values, consistent with shipping costs shifting up the eventual prices charged by firms. Depending on the specification, a one percentage point increase in shipping costs results in 0.14 to 0.17 percent higher unit values.

Table D5: Alternative Instrument 1: First Stage

	Dependent Variable: Price (logs)		
	(1)	(2)	(3)
CIF Charges (in %)	0.167^{**}	0.144^{***}	0.144^{***}
	(3.25)	(8.00)	(10.84)
Time FEs	Yes	Yes	Yes
Product FEs	No	Yes	Yes
Exporter FEs	No	No	Yes
Importer FEs	No	No	Yes
Observations	12,493,156	12,493,156	12,493,156

Standard Errors are clustered by period and exporter. t statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

D.5.2 Alternative Instrument based on U.S. shipping costs

As a second alternative I use U.S. data to construct a measure of shipping cost between the trade partners in the EU data. While Eurostat does not publish a measure of shipping cost, the U.S. Census Bureau provides data on cost, insurance and freight (cif) charges at the HS10 level of disaggregation for each export to the United States.[4] I use this information to proxy for the cost that countries incur when exporting to the EU, which I then use as an instrument for the prices they charge.

More specifically, I first compute the average charges by unit and kilometer (km) incurred by countries when selling a product k to the United States during period t. For that purpose, I add up all charges paid by exporters of k to the U.S. and divide this sum by the total quantities shipped and the total distances covered:

$$\widehat{\text{Cost per unit}} \text{ \& km}_{kt} = \frac{\sum_c \text{Charges}_{ckt}}{\sum_c \text{Quantity}_{ckt} \cdot \text{Distance}_c}.$$

The idea is that this measure captures the common component in the shipping cost per unit and km which all sellers face when selling abroad. The actual instrument is then constructed by multiplying the estimate of this cif costs by the distance between the respective trade partners c (the exporter) and m (the importer) in the Eurostat data, i.e.

$$\mathbb{Z}^{(2)}_{cmkt} = \widehat{\text{Cost per unit}} \text{ \& km}_{kt} \cdot \text{Distance}_{c,m}.$$

Since the U.S. data is reported on the 10-digit HS level, I construct the $\widehat{\text{Cost per unit}}$ & km_{kt} measure on the 6-digit HS level and make again use of the fact that the first 6 digits of the Harmonized System and the Combined Nomenclature Classifications coincide. If no quantity data is available, I again use the values on a higher level of aggregation (e.g. the 5-digit instead of the 6-digit level).

When matching the average unit shipping cost to the Eurostat data, some observations cannot directly be matched. This is for example the case when the EU imported a product while the U.S. did not. Also in those cases, I rely on values for a higher aggregation level until all observations are matched to a unit cost. Finally, when matching trade to production data, I use the simple average over the unit shipping costs of all cn8 that belong to a pc8plus category.

I also compute an alternative statistic based on trade values instead of quantities, which has the advantage of not suffering from missing values:

$$\widehat{\text{Cost per \$}} \text{ \& km}_{kt} = \frac{\sum_c \text{Charges}_{ckt}}{\sum_c \text{Trade Value}_{ckt} \cdot \text{Distance}_c}.$$

This measure captures the average cost of shipping a quantity of product k worth 1 US dollar over 1 km from country c to the United States. I found however little difference in the results when using an instrument based on this cost per value relative to the one based on quantities, $\mathbb{Z}^{(2)}_{cmkt}$.

Table D6 demonstrates the correlation of the instrument with unit values of EU imports. After controlling for GDP per capita or the inclusion of exporter fixed effects, there is a robust positive relationship between shipping costs and prices: A one percent increase in shipping costs imply a .045 to .075 increase in the price.

---

[4] This data, used and described in more detail in Schott (2008), is generously provided on Peter Schott's website.

Table D6: Alternative Instrument 2: First Stage

	Dependent Variable: Price (logs)			
	(1)	(2)	(3)	(4)
Charges (logs)	-0.00747	0.0454^{*}	0.0585^{***}	0.0746^{***}
	(-0.18)	(2.09)	(7.33)	(6.94)
GDP per capita$_{Exporter}$		0.0155^{***}		
		(7.21)		
Product FEs	Yes	Yes	No	No
Product-Exporter FEs	No	No	Yes	No
Product-Exporter-Time FEs	No	No	No	Yes
Observations	13,212,843	13,212,843	13,212,285	12,648,046

Standard Errors are clustered by product and exporter. *t* statistics in parentheses. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. GDP per capita is measured in 1,000 Euros.

D.5.3 Price Coefficients: OLS versus 2SLS

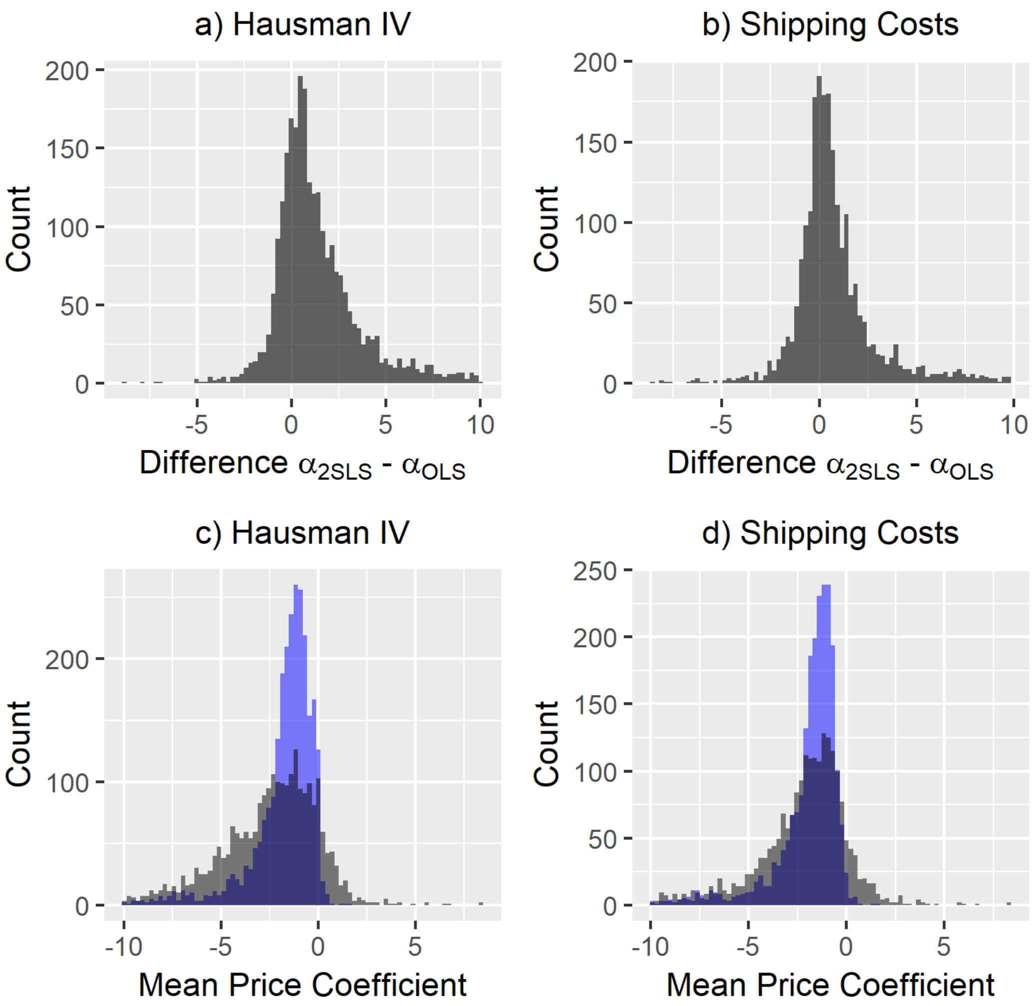

Figure D4: **Price Coefficients - OLS versus 2SLS**: Panels a) and b) plot the distribution of differences between the price coefficient estimates using OLS versus 2SLS. Panels c) and d) picture the distributions of the price coefficients in the OLS (blue) and the 2SLS case (gray). Outliers are excluded from the graphs, i.e. differences that are above 10 in absolute value in panels a) and b) and coefficients above 10 in absolute value in panels c) and d).

D.5.4 Cobb-Douglas Utility Weights

Table D7: Cobb-Douglas Utility Weights

Baseline weights - Summary Statistics:					
	N	Mean	St. Dev.	Min	Max
ω_k	3,103	0.0003	0.001	2.179e-09	0.032
Country-specific weights - Correlations:					
	Germany	France	UK	Italy	
Germany	1.0000	0.7050	0.8386	0.5870	
France	0.7050	1.0000	0.6400	0.6834	
UK	0.8386	0.6400	1.0000	0.5796	
Italy	0.5870	0.6834	0.5796	1.0000	
Year-specific weights - Correlations:					
	2003	2004	2005		
2003	1.0000	0.9866	0.9673		
2004	0.9866	1.0000	0.9854		
2005	0.9673	0.9854	1.0000		

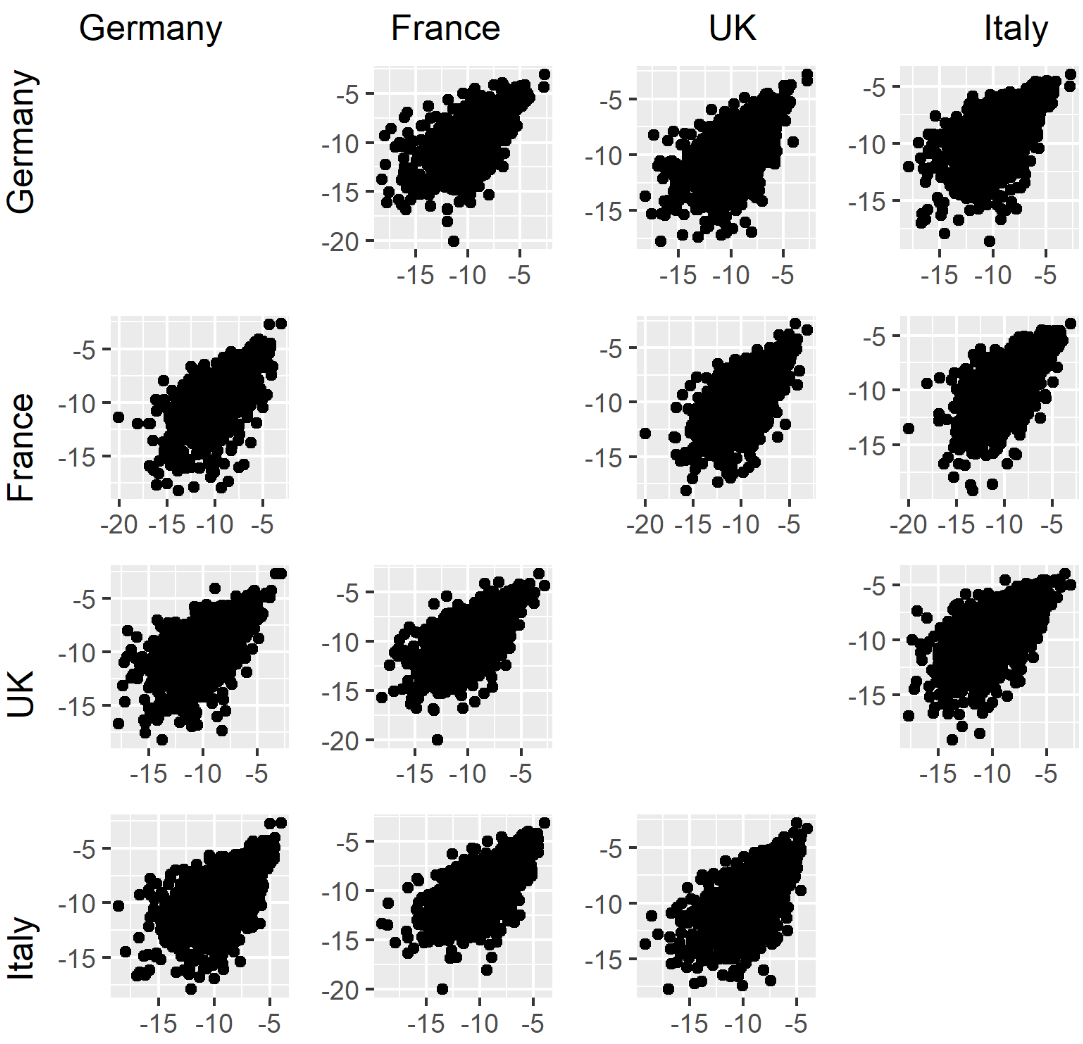

Figure D5: **Country-specific Cobb-Douglas Utility Weights**: The plot shows the inferred utility weights ω when the weights are allowed to differ by country. Each panel plots the country-specific values (in logs) against each other for a given country pair for the year 2005.

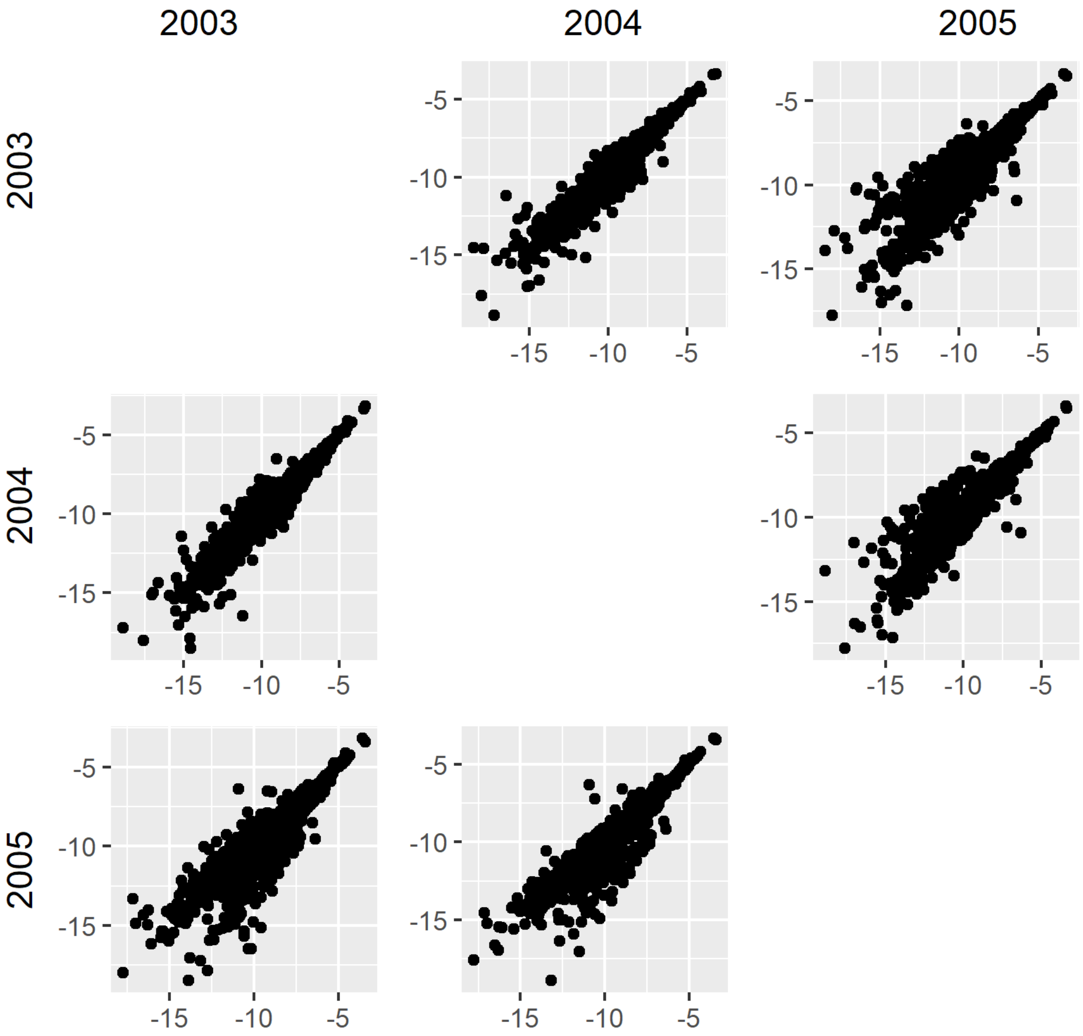

Figure D6: **Time-specific Cobb-Douglas Utility Weights**: The plot shows the inferred utility weights ω when the weights are allowed to differ by year. Each panel plots the year-specific values (in logs) against each other for each given pair of years.

References

[1] Amiti, Mary, and Amit Khandelwal (2013): "Import Competition and Quality Upgrading"; The Review of Economics and Statistics 95 (2), pp. 476–490.

[2] Anderson, Simon, André De Palma, and Jaques-François Thisse (1987): "The CES is a Discrete Choice Model?"; Economics Letters, 24, pp. 139-140.

[3] Autor, David, David Dorn, Gordon Hanson and Jae Song (2014): "Trade Adjustment: Worker-Level Evidence"; Quarterly Journal of Economics 129 (4), pp. 1799-1860.

[4] Berry, Steven (1994): "Estimating Discrete Choice Models of Product Differentiation"; RAND Journal of Economics, 25, 242-262.

[5] Berry, Steven, James Levinsohn, and Ariel Pakes (1995): "Automobile Prices in Market Equilibrium"; Econometrica, 63, 841-890.

[6] Beveren, Ilke Van, Andrew Bernard, and Hylke Vandenbussche (2012): "Concording EU Trade and Production Data over Time"; NBER Working Paper No. 18604.

[7] Bloom, Nicholas, Mirko Draca, and John Van Reenen (2016): "Trade-induced Technical Change? The Impact of Chinese Imports on Innovation, IT, and Productivity"; Review of Economic Studies 83, 87–117

[8] Broda, Christian, and John Romalis (2009): "The Welfare Implications of Rising Price Dispersion"; NBER Working Paper No. 18314.

[9] Broda, Christian, and David Weinstein (2006): "Globalization and the Gains from Variety"; Quarterly Journal of Economics, 121, 541-585.

[10] Chaney, Thomas (2008): "Distorted Gravity: The Intensive and Extensive Margins of International Trade"; American Economic Review, 98, pp. 1707–1721.

[11] Ciliberto, Federico, and Elie Tamer (2009): "Market Structure and Multiple Equilibria in Airline Markets"; Econometrica, 77, pp. 1791–1828.

[12] Cosar, Kerem, Paul Grieco, Shengyu Li, and Felix Tintelnot (2015): "What Drives Home Market Advantage?"; Working Paper.

[13] Deaton, Angus, and John Muellbauer (1980): "Economics and Consumer Behaviour"; Cambridge: Cambridge University Press.

[14] Dingel, Jonathan (2017): "The Determinants of Quality Specialization"; Review of Economic Studies 84, pp. 1551–1582

[15] Dubé, Jean-Pierre, Jeremy Fox and Che-Lin Su (2012): "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation"; Econometrica, 80, pp. 2231–2267.

[16] Eaton, Jonathan and Samuel Kortum (2002): "Technology, Geography and Trade"; *Econometrica* 70 (5) (September 2005), 1741-1779

[17] Eddelbuettel, Dirk and Romain Francois (2011): ”Rcpp: Seamless R and C++ Integration”; Journal of Statistical Software, 40, pp. 1-18.

[18] Faber, Benjamin (2014): ”Trade Liberalization, the Price of Quality, and Inequality: Evidence from Mexican Store Prices”; Working Paper.

[19] Faber, Benjamin and Fally, Thibault (2017) - ”Firm Heterogeneity in Consumption Baskets, Evidence from Home and Store Scanner Data”; NBER Working Paper 23101.

[20] Fajgelbaum, Pablo, Gene Grossman, and Elhanan Helpman (2011): ”Income Distribution, Product Quality, and International Trade”; Journal of Political Economy, 119, 721-765.

[21] Fajgelbaum, Pablo, and Amit Khandelwal (2016): ”Measuring the Unequal Gains from Trade”; The Quarterly Journal of Economics (2016), pp. 1113–1180.

[22] Feenstra, Robert, and John Romalis (2014): ”International Prices and Endogenous Quality”; The Quarterly Journal of Economics, 129 (2), pp 477 - 527.

[23] Fieler, Ana Cecília (2011): Nonhomotheticity and Bilateral Trade: Evidence and a Quantitative Explanation”; Econometrica, 79, pp. 1069–1101.

[24] Goldberg, Pinelopi and Nina Pavcnik (2007): ”Distributional Effects of Globalization in Developing Countries”; Journal of Economic Literature 45, pp. 39–82.

[25] Goldberg, Pinelopi (Forthcoming): ”Introduction”; In Trade And Inequality, edited by Pinelopi Goldberg. International Library of Critical Writings in Economics Series. Cheltenham, UK: Edward Elgar Publishing.

[26] Hallak, Juan Carlos (2006): ”Product Quality and the Direction of Trade”; Journal of International Economics, 68, pp 238-265.

[27] Hallak, Juan Carlos, and Peter K. Schott (2011): ”Estimating Cross-Country Differences in Product Quality”; Quarterly Journal of Economics, 126, pp. 417–474.

[28] Handbury, Jessie (2013): ”Are Poor Cities Cheap for Everyone? Non-Homotheticity and the Cost of Living Across U.S. Cities”; Working Paper.

[29] Hausman, Jerry (1996): ”Valuation of New Goods Under Perfect and Imperfect Competition”; in *The Economics of New Goods, Studies in Income and Wealth*, Vol. 58, ed. by T. Bresnahan and R. Gordon. Chicago: National Bureau of Economic Research.

[30] Helpman, Elhanan, Marc Melitz, and Yona Rubinstein (2008): ”Estimating Trade Flows: Trading Partners and Trading Volumes”; The Quarterly Journal of Economics, 123, pp. 441-487.

[31] Hummels, David, Rasmus Jørgensen, Jakob Munch, and Chong Xiang (2014): ”The Wage Effects of Offshoring: Evidence from Danish Matched Worker-Firm Data”; American Economic Review 104 (6), pp. 1597–1629.

[32] Hummels, David, and Peter Klenow (2005): ”The Variety and Quality of a Nation’s Exports”; American Economic Review, 95, pp. 704-723.

[33] Hummels, David, and Volodymyr Lugovskyy (2009): ”International Pricing in a Generalized Model of Ideal Variety”; Journal of Money, Credit, and Banking, 41, pp. 3-33.

[34] Khandelwal, Amit (2010): "The Long and Short (of) Quality Ladders"; Review of Economic Studies, 77, pp. 1450-1476.

[35] Knittel, Christopher and Konstantinos Metaxoglou (2008): "Estimation of Random Coefficient Demand Models: Challenges, Difficulties and Warnings"; NBER Working Paper No. 14080.

[36] Kugler, Maurice and Eric Verhoogen (2012): "Prices, Plant Size, and Product Quality"; Review of Economic Studies 79, 307–339.

[37] Melitz, Marc (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity"; Econometrica, 71, pp. 1695-1725.

[38] Nevo, Aviv (2001): "Measuring Market Power in the Ready-to-Eat Cereal Industry"; Econometrica 69, pp. 307-342.

[39] Ossa, Ralph (2015): "Why Trade Matters after all"; Journal of International Economics 97, pp. 266–277.

[40] Petrin, Amil (2002): "Quantifying the Benefits of New Products: The Case of the Minivan"; Journal of Political Economy 110 (4), pp. 705-729.

[41] Porto, Guido (2006): "Using survey data to assess the distributional effects of trade policy"; Journal of International Economics 70 (1), pp. 140-160.

[42] Schott, Peter (2004): "Across-Product versus Within-Product Specialization in International Trade"; Quarterly Journal of Economics, 119, 646-677.

[43] Schott, Peter (2008): "The Relative Sophistication of Chinese Exports"; Economic Policy, 23, 5-49.

[44] Soderbery, Anson (2015): "Estimating import supply and demand elasticities: Analysis and implications"; Journal of International Economics 96, pp. 1–17.

[45] Su, Che-Lin, and Kenneth Judd (2012): "Constrained Optimization Approaches to Estimation of Structural Models"; Econometrica, 89, pp. 2213–2230.

[46] Verhoogen, Eric (2008): "Trade, Quality Upgrading, and Wage Inequality in the Mexican Manufacturing Sector"; Quarterly Journal of Economics 123 (2), pp 489–530.

[47] Wächter, Andreas and Lorenz Biegler (2006): "On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming"; Mathematical Programming, 106, pp. 25-57.

[48] Williams, Brian (2008): "Europroms Prodcom User Guide"; 06.02.2008.