

Per capita income and the extensive margin of bilateral trade*

Christian Hepenstrick[†] Alexander Tarasov[‡]

November 12, 2012

Abstract

This paper quantitatively explores the role of the demand structure in explaining the relationship between an importer’s per capita income and the extensive margin of bilateral trade. The underlying mechanism is based on the fact that agents expand the set of goods they consume with income. This in turn affects the structure of a country’s import demand and therewith the extensive margin of trade. We formalize this intuition by incorporating preferences that allow for binding non-negativity constraints into an otherwise standard Ricardian multi-country model. We quantify the model using the data on US consumer expenditures and aggregate values of bilateral trade flows and find that the behavior of the model’s extensive margin of bilateral trade is consistent with the data (as opposed to the standard model). Two popular counterfactual experiments – lower trade costs and the rise of China and India – demonstrate that the mechanism outlined in this paper is indeed quantitatively important.

Keywords: Non-homothetic preferences, extensive margin, Ricardian trade

JEL Classification: F10, F11, F19

1 Introduction

Trade flows vary greatly in the number of different goods that are traded between countries. This “extensive margin of bilateral trade” can be thought of as the manifestation of an interplay

*We thank Ana Cecília Fielér, Kevin Staub, Gino Gancia, Andreas Kohler, Pedro Mira, Josef Zweimüller, Alessandro Barattieri, Marc Melitz, Pol Antràs, Tomasz Swiecki, Elhanan Helpman, Samuel Kortum and seminar participants at the University of Zurich, the Vienna Institute for International Economic Studies, Harvard University, and EconCon for helpful comments and discussions. Christian Hepenstrick gratefully acknowledges financial support by the Verein zur Förderung des Akademischen Nachwuchses, the Ecosciencia Foundation, and the Swiss National Science Foundation.

[†]Swiss National Bank, Börsenstrasse 15, CH-8022 Zurich, Switzerland, email: christian.hepenstrick@snb.ch.

[‡]Department of Economics, University of Munich, Ludwigstrasse 28, 80539 Munich, Germany, phone: +49 89 2180 57 54, email: alexander.tarasov@lrz.uni-muenchen.de.

of the exporter’s production technology, the importer’s demand structure, and bilateral trade costs - two countries are more likely to trade a given good if the exporter is particularly strong in producing this good, the importer has an especially high demand for that good, or bilateral trade barriers are low. Most existing analyses of the extensive margin of bilateral trade focus on the exporter’s technology and bilateral trade costs.¹ This paper provides a complementary perspective by emphasizing the role of the importing country’s demand structure.

Several studies have documented that richer countries both import and export more varieties (see e.g. Hummels and Klenow (2002), Baldwin and Harrigan (2011), or Sauré (2012)). Table 1 summarizes these findings by regressing the extensive margin of bilateral trade in consumption goods² (in logs) on the per capita incomes and population sizes (both in logs) of the trading partners controlling for bilateral resistance.

Table 1

The positive elasticity of exporter per capita income can be explained by a standard Ricardian mechanism: per capita income is high due to a country’s advanced technology. Being technologically advanced in turn implies that this country is competitive in many industries and, therefore, tends to export a broad set of varieties (which generates the positive correlation between exporter income and the extensive margin of trade). At the same time, however, it follows immediately that the number of varieties the country imports is relatively low since it is optimal to produce most varieties locally due to the advanced technology. As a result, a standard Ricardian framework (with CES preferences) implies a negative elasticity of the extensive margin of trade with respect to importer per capita income, which is inconsistent with the data. In this paper, we argue that the fact that consumers expand the set of goods they purchase – henceforth called the “extensive margin of consumption” – with income generates a countervailing force that can change the sign of the elasticity.³

To formalize this intuition, we adapt the Ricardian multi-country model by Eaton and Kortum (2002) (henceforth EK). The production technology in the EK framework is country-variety specific. Together with trade costs this determines the price at which a supplying country can offer a variety in a particular destination market. In any given destination market the producer country offering the lowest price will be the sole supplier of this variety. However, whether this trade flow then actually materializes depends not only on the supplier country’s good technology and low bilateral trade costs, but crucially also on the importing country’s

¹See, for instance, Chaney (2008).

²We limit the analysis to consumption goods, as the channel proposed in this paper affects the patterns of final demand only. Details on the data can be found in Appendix A.6.

³Jackson (1984) documents the empirical relevance of the “extensive margin of consumption” using US consumer expenditure data. Falkinger and Zweimüller (1996) provide evidence that richer countries consume a broader set of goods using the World Bank’s ICP-data. Other studies documenting a positive correlation between variety and income include Jekanowski and Binkley (2000), Moon, Florkowski, Beuchat, Resurreccion, Paraskova, Jordanov, and Chinnan (2002), and Thiele and Weiss (2003) for food consumption and Gronau and Hamermesh (2008) documenting similar effects in time use data.

demand structure - particularly if there is actually a positive demand for this variety at the price offered. With traditional CES preferences demand is always positive, as marginal utility is unbounded. In the present paper, we relax this assumption and allow agents to adjust their extensive margin of consumption with income. In choosing their optimal consumption bundle agents order the varieties in accordance with their prices and decide up to which price consuming positive quantities is optimal. It is therefore possible that the lowest price offered by the supplier countries for a given variety is still too high, so that agents in the importing country find it optimal not to consume (and, therefore, not to import) this variety. As a result, the model is in principle able to explain the empirical behavior of the extensive margin of trade.

In order to assess whether this demand side mechanism is actually strong enough to dominate the negative effect coming from the supply side we quantify the model. In particular, we use aggregate values of bilateral trade flows to estimate the model's technologies and trade costs and microdata on US consumer expenditure to pin down the preference parameters. The advantage of using data on US consumer expenditure is that this data set is independent from the trade data, implying that the estimates of the preference parameters are independent from the trade data as well. Therefore, when we then simulate the calibrated model and compare the behavior of the model's extensive margin of bilateral trade to the data, we compare the model's predictions to an aspect of the data that has not been used in quantifying the model. We find that the demand side forces are sufficiently strong to dominate the negative effect of technology on the extensive margin of imports. Comparing the simulated income elasticities of the extensive margin of bilateral trade to the data, we find that they are of the right sign and close to the empirical elasticities.

What is the quantitative importance of this new demand side channel? To answer this question, we use the calibrated model to perform two classical counterfactual experiments - lowering trade costs and the rise of China and India. A traditional model neglecting the demand side predicts that lower trade costs lead to a higher extensive margin of trade since trade becomes worthwhile for more varieties. The demand side reinforces this effect: lower trade costs lead to a higher real income and, thereby, induce agents to expand their extensive margin of consumption. This in turn increases the number of traded varieties. Quantitatively, we find that the predicted changes in the extensive margin of bilateral trade flows are at least twice as high when allowing for the extensive margin of consumption to adjust. In the second experiment, we explore the effect of technological progress in China and India. Better local technologies imply that these countries will tend to produce more varieties locally since they become more competitive. However, a rise in income will lead consumers to expand the measure of varieties that they consume. This effect dominates: the model with non-homothetic preferences therefore predicts that the extensive margin of imports will rise, while the model neglecting demand side effects would actually predict falling extensive margins (that is, the two countries would import fewer varieties as they become richer).

This paper contributes to two strands of the trade literature. First, by highlighting the

role of the importer's demand structure, it contributes to the broad literature investigating the extensive margin of trade. Second, on a more theoretical side, it contributes to the growing literature that recognizes the potential importance of non-homothetic consumer behavior in understanding different aspects of trade patterns. The focus of previous contributions was on how non-homothetic consumer behavior helps us to understand variations in the aggregate values of trade flows and differences in unit prices⁴, whereas this study focuses on the extensive margin of trade. As for the aggregate values of trade flows, Fielser (2011) argues that non-homothetic preferences help to explain the higher trade share of rich countries. She extends the EK model to two industries with differing demand elasticities. Richer countries then relatively concentrate their expenditures in the high-elasticity industries. If the variability in productivities across countries is relatively high in these industries, their share of traded output will be high. Together with the demand pattern, this implies that rich countries' trade shares are higher. However, agents in her model do not adjust their extensive margin of consumption. The predictions with respect to the extensive margin of trade are therefore the same as in EK: i.e., a counterfactual negative correlation of per capita income and the extensive margin of imports. Another aspect of trade patterns where non-homotheticity is potentially relevant is the variation in unit-prices across importing countries. Simonovska (2010) argues that differences in unit prices reflect differences in markups arising due to demand elasticities that change with income. Choi, Hummels, and Xiang (2009) and Fajgelbaum, Grossman, and Helpman (2011), on the other hand, understand differences in unit-prices as reflecting quality differences due to an increasing taste for quality with rising income.

Concerning the extensive margin of trade, Sauré (2012) argues that richer countries have more trading partners (the country level extensive margin) due to non-homothetic preferences and Foellmi, Hopenstrick, and Zweimüller (2010) show that non-homothetic preferences can generate incentives for parallel trade and influence the extensive margin of trade via this channel. Matsuyama (2000) shows how a hierarchy in consumption can interact with comparative advantages and, therefore, give a rise to product cycles and effects reminiscent of the Prebisch-Singer hypothesis. The contribution of the present paper is a quantitative multi-country model of trade where consumers in an importing country decide on their extensive margin of consumption. This decision then determines, together with the exporters' technologies and the structure of trade costs, the extensive margins of bilateral trade flows.

The rest of the paper is organized as follows. Section 2 develops the model. Section 3 then discusses the quantification and its results. In this paper, we use the EK framework since it is parsimonious in the context of our objective - it allows us to develop our argument in a very intuitive way, while providing us at the same time with a model, which is general enough to be quantified directly. However, the message of the demand side being an important determinant of the extensive margin of trade is more general. Therefore, in Section 4, we discuss

⁴For a more complete overview of the literature on non-homothetic preferences and trade see Markusen (2010).

how demand side effects would play in a model of monopolistic competition and heterogeneous firms and what additional effects and complications may arise. Section 5 discusses a number of extensions. First, we show that accounting for trade in intermediates does not significantly change the quantitative predictions of the model. Second, we consider the implications of allowing for within-country inequality. Finally, we show that the results are not driven by a particular functional form of the utility function. Section 6 concludes.

2 The Model

The world economy consists of N asymmetric countries. Country i 's population is denoted by L_i . Each agent is endowed with one unit of labor, which is inelastically supplied on the domestic market. There is one industry producing differentiated consumption goods. The measure of varieties is exogenous and normalized to unity.

2.1 Consumer behavior

Agents maximize the same symmetric additively separable utility function

$$U = \int_0^1 v(x(j)) dj$$

subject to their budget constraints $E \geq \int_0^1 p(j)x(j) dj$ and the non-negativity constraints $x(j) \geq 0 \forall j$. E denotes an agent's income. The sub-utility function $v(x)$ is concave, $v'(x) > 0$ and $v''(x) < 0$, and the marginal utility is bounded from above, $v'(0) < \infty$. With bounded marginal utility, the non-negativity constraints are potentially binding and the corresponding first order conditions for some variety j are

$$\begin{aligned} v'(x(j)) &= \lambda p(j) & \text{for } x(j) > 0 \\ v'(0) &< \lambda p(j) & \text{for } x(j) = 0 \end{aligned} \tag{2.1}$$

where λ is the Lagrange multiplier. Intuitively, for every available variety j , an agent compares the marginal utility from starting to consume this variety $v'(0)$ with her utility costs $\lambda p(j)$ associated with buying a marginal unit of this variety. If the marginal costs are higher than the marginal utility, the optimal quantity is zero - the non-negativity constraint binds. The optimal quantities are positive for all other varieties and are chosen such that the marginal rates of substitution equal relative prices. Figure 1 depicts the demand function for a particular

variety j . Note that there is a finite price $v'(0)/\lambda$ above which the optimal quantity is zero.

Figure 1

As the varieties enter the utility function symmetrically, agents simply order the varieties according to their prices (think of reindexing the varieties such that the prices rise in the index j) and then choose up to which price they still want to consume positive quantities. We denote the index of this marginal variety by M . Its price follows from rearranging the first order condition (2.1) when the non-negativity constraint just binds, $p(M) = v'(0)/\lambda$. As the indices are increasing in prices, M also denotes the measure of varieties consumed in positive quantities and thus represents the *extensive margin of consumption* in this model. Because the goods spectrum is normalized to one, the extensive margin of consumption simultaneously represents the share of available varieties consumed in positive quantities.

Assuming that the price distribution can be represented by a continuous cdf $G(p)$ - this will be the case in the general equilibrium - the share of consumed varieties (and therefore also the extensive margin of consumption) is

$$M = G(v'(0)/\lambda). \quad (2.2)$$

For varieties $j < M$, the Marshallian demand follows from rearranging the first order condition, $x(\lambda p) = v'^{-1}(\lambda p)$. Inserting the Marshallian demand into the budget constraint and making the change of variable $p = G^{-1}(j)$ yields

$$E = \int_0^{v'(0)/\lambda} px(\lambda p) dG(p) \quad (2.3)$$

implicitly determining the marginal utility of income λ for a given price distribution $G(p)$ and income E . We will know the price distribution and the income in the equilibrium and thus be able to solve for the multiplier λ . Using (2.2), one then solves for the extensive margin of consumption and (2.1) determines the optimal quantities for varieties $j < M$.

The potentially binding non-negativity constraints contrast with previous EK models (see for example EK and Fieler, 2011) that assume CES preferences.⁵ The unbounded marginal utility of CES preferences implies that agents in this class of models always consume all varieties, $M = 1$, independent of the price distribution and income.

⁵Fieler (2011) models two industries with different elasticities of substitution. Agents therefore adjust their relative expenditures across industries with income and thus exhibit a non-homothetic consumption pattern. However, since the marginal utility is unbounded, the agents do not adjust the sets from which they consume, i.e. all agents always consume all varieties in both industries.

2.2 Production technology and market structure

The supply side of the model is similar to the basic EK framework: The production technology exhibits constant returns to scale and uses one input⁶, which we call labor. Labor is assumed to be perfectly mobile within countries, but immobile across countries, so that in equilibrium there will be one wage rate w_i per country. $z_i(j)$ denotes country i 's productivity in producing variety j . Assuming perfect competition and iceberg trade costs⁷ - $d_{ni} \geq 1$ units need to be shipped in i for one unit to arrive at destination n - implies that the price at which country i offers variety j in country n is

$$p_{ni}(j) = \frac{w_i d_{ni}}{z_i(j)}. \quad (2.4)$$

The country-variety specific productivity $z_i(j)$ is the realization of a Fréchet distributed random variable $Z_i(j)$:

$$\Pr [Z_i(j) \leq z] = \exp \{-T_i z^{-\theta}\},$$

where T_i is country specific and governs the expected productivity draw. We will therefore refer to T_i as country i 's technology (higher T_i implies a higher expected productivity and, therefore, represents a more advanced technology in country i). θ is common to all countries and controls the variation in the productivities (the lower θ , the more variation there is in productivity draws). We will show later that θ also governs the elasticity of trade volumes with respect to trade costs, which is why we sometimes refer to θ as the trade elasticity.

2.3 Equilibrium

All countries i are in principle able to produce each variety j . However, consumers will purchase each variety from only one producing country - the country offering the lowest price. International trade thus emerges if the country with the lowest price is a foreign country. In Appendix A.1, we show that the lowest prices available in country n can be represented by the following cdf:

$$G_n(p) = 1 - \exp \{-\Phi_n p^\theta\}, \quad (2.5)$$

where

$$\Phi_n = \sum_{i=1}^N T_i (w_i d_{ni})^{-\theta}.$$

I.e., $G_n(p)$ is the share of varieties in country n with a price (weakly) below p . As all agents are endowed with one unit of labor, the income of an agent in country n is simply the wage

⁶For parsimony, we abstain from modeling multiple inputs. This implies that, in the calibration exercise, differences in non-tradable endowments (e.g. human and physical capital) and price indices for tradable intermediates are absorbed into the calibrated technology. This is admissible for non-tradable endowments in the context of the counterfactuals. For tradable intermediate inputs, we show in Section 5 how the results change when allowing for them.

⁷We normalize trade costs within a country to unity: $d_{nn} = 1$ for all n ; and assume that the triangle inequality, $d_{ni} \leq d_{ki} d_{nk}$, holds for all i, k , and n .

rate, w_n (to be endogenously determined). For a given wage rate, the budget constraint (2.3) together with the price distribution therefore determines the marginal utility in country n , λ_n . The extensive margin of consumption in n follows immediately,

$$M_n = 1 - \exp \left\{ -\Phi_n (v'(0) / \lambda_n)^\theta \right\}. \quad (2.6)$$

In Appendix A.2, we derive the probability, π_{ni} , that country i is the cheapest supplier of a certain variety in country n ,

$$\pi_{ni} = \frac{T_i (w_i d_{ni})^{-\theta}}{\Phi_n}.$$

Note that, since there is a continuum of goods, π_{ni} is also the share of varieties for which country i is the cheapest supplier in n . Moreover, because the probability is the same for all goods, π_{ni} is not only the share of the total goods spectrum, but also the share for any sub-spectrum; in particular also for an importing country's extensive margin of consumption. This in turn implies that the *extensive margin* of the *bilateral trade flow* from i to n , which we denote by m_{ni} , is simply the importing country's extensive margin of consumption, M_n , multiplied by the share of varieties for which the exporter i is the cheapest producer, π_{ni} ,

$$m_{ni} = \pi_{ni} M_n. \quad (2.7)$$

The source of the remarkable simplicity of this result is worth discussing. The distributional assumption of the EK framework implies that conditional on entering market n , prices have the same distribution across supplier countries.⁸ Therefore, the prices of the goods that the importer n actually buys bear no information about the likely source of these goods, so that π_{ni} is the share of varieties of any subinterval of the variety space for which the supplier i offers the cheapest price in n - in particular, for the subinterval M_n representing the M_n percent cheapest varieties as well (in Appendix A.3, we provide the corresponding derivations).

Since the distribution of prices of goods that are actually sold in n is the same across supplier countries i , average sales (per product) do not vary by source. In the model, average sales in n are given by the total expenditures, $w_n L_n$, divided by the extensive margin of consumption, $w_n L_n / M_n$. The aggregate value of the bilateral trade flow from i to n , X_{ni} , is given by the average sales multiplied by the measure of varieties for which i is the cheapest supplier in n and which are actually consumed in positive quantities - the extensive margin of trade, m_{ni} - which yields

$$X_{ni} = m_{ni} (w_n L_n / M_n) = \pi_{ni} w_n L_n. \quad (2.8)$$

Note that, taking the derivative of the log of volumes with respect to the log of trade costs,

⁸Eaton and Kortum (2005) call this feature "neutrality". Note that neutrality is not unique to the Ricardian framework. In the monopolistic competition model with heterogeneous firms and market entry costs, neutrality follows if market entry costs are only destination specific and productivities are drawn from a Pareto distribution. Costinot and Komunjer (2007) provide a discussion of general productivity distributions in the Ricardian multi-country framework.

yields θ (neglecting general equilibrium effects on the wage rates), which is why θ is sometimes called the trade elasticity.

To close the model and determine the equilibrium wage vector, we use the labor market clearing conditions⁹

$$w_i L_i = \sum_{n=1}^N \frac{T_i (w_i d_{ni})^{-\theta}}{\sum_{k=1}^N T_k (w_k d_{nk})^{-\theta}} w_n L_n \quad \text{for } i = 1, \dots, N. \quad (2.9)$$

In summary, the structure of the global economy is characterized by the countries' technologies, T_i , and populations, L_i , the matrix of bilateral trade costs, d_{ni} , the trade elasticity, θ , and the shape of the utility function, $v(\cdot)$. In the equilibrium, producers price according to (2.4) and consumers choose their optimal quantities and extensive margins (2.6) as implied by (2.1). Market clearing (2.9) pins down the set of equilibrium wage rates and bilateral trade patterns are characterized by their aggregate value (2.8) and their extensive margin (2.7).¹⁰

2.4 The role of per capita income

In this section, we discuss how per capita income affects trade patterns and contrast the results with the standard model with homothetic preferences.¹¹ In the context of this paper, the role of non-homothetic preferences is to endogenize the extensive margin of consumption. Figure 2 illustrates this by depicting equation (2.6).

Figure 2

The price of the marginal variety is $v'(0)/\lambda_n$. Using the country specific price distribution $G_n(p)$, one gets the share of varieties with prices lower than this marginal price and thus the share (and measure) of varieties consumed in positive quantities - the extensive margin of consumption. An increase in per capita income lowers the marginal utility of income λ_n and thus increases the extensive margin of consumption - richer countries consume a broader set of varieties. A first order stochastic dominance shift in the price distribution increases

⁹The labor market clearing condition follows from imposing balanced trade, $\sum_{k \neq i} X_{ik} = \sum_{n \neq i} X_{ni}$, and adding the domestically sourced consumption, X_{ii} , on both sides to get $\sum_{k=1}^N X_{ik} = \sum_{n=1}^N X_{ni}$. Total expenditures in i are $w_i L_i = \sum_{k=1}^N X_{ik}$. Substituting for X_{ni} and π_{ni} on the right hand side then yields the labor market clearing condition as stated in the text.

¹⁰The general equilibrium exists and is unique. To see this, note that the labor market clearing conditions can be rewritten as excess demand for labor. It is straightforward to show that the resulting system of excess demand functions satisfies the sufficient properties for existence and uniqueness (see for example Propositions 17.B.2 and 17.F.3 in Mas-Colell, Whinston, and Green (1995)). The extensive margins and the optimal quantities are unique, as they follow from maximizing a concave object over a convex constraint. As wages and extensive margins fully summarize the general equilibrium, this implies that the general equilibrium is unique.

¹¹With homothetic preferences, expenditure shares are constant. One can show that a linear transformation of the widely used CES-preferences is indeed the most general form of additively separable homothetic preferences. Note that homotheticity (for additive preferences) requires $v'(0) = \infty$, as sufficiently poor agents will otherwise not buy an expensive variety, i.e. their expenditure share is zero, whereas the expenditure share of sufficiently rich agents is positive, which contradicts homotheticity.

the extensive margin of consumption as well - countries that are better integrated and that, therefore, have lower prices consume a broader set of varieties. Note finally that for $v'(0) \rightarrow \infty$, the extensive margin approaches one and agents would not adjust their extensive margins with income.

In the Ricardian framework, a country's per capita income depends on that country's technology - the better technology, the higher the wage rate and thus per capita income. Thus, in the present model, there are two channels of how an importing country n 's technology affects the extensive margin of trade. The first channel is the standard supply side channel. A better technology implies better productivity draws for more varieties, so that the share of varieties π_{nn} for which local producers offer the best prices rises. But this means that the share of varieties that are imported, $(1 - \pi_{nn})$, and thus the extensive margin of bilateral trade tends to fall. On the other hand, the non-homothetic model exhibits the second channel. Better technology leads to higher wages and agents, therefore, extend their extensive margin of consumption, which tends to increase the extensive margin of trade. In the calibrated version of the present model, it turns out that the latter effect dominates, so that the extensive margin of bilateral trade is positively correlated with per capita income (which is consistent with the data). In the homothetic model, on the other hand, only the first effect is present (all countries' extensive margins of consumption are equal to one) and the correlation between the extensive margin of trade and an importer income is unambiguously negative. Note also that the decomposition of aggregate GDP into population size and per capita income matters in the non-homothetic model - a rich but small country has a high extensive margin of consumption and thus also tends to import a broad set of varieties, whereas the opposite is true for a poor but large country, although the two countries may have the same aggregate GDP. This is also consistent with the data (see Table 1).

While the non-homothetic and the homothetic model can have opposing predictions for the extensive margin of bilateral trade, they exhibit the same pattern for the aggregate volumes. This feature is very useful when calibrating the model and comparing its performance to the homothetic model. Note also that with respect to aggregate volumes, neither model has a separate role for differences in per capita incomes induced by different technologies.¹²

3 Quantifying the model

In this section, we quantify the model to assess whether the above theory is able to explain the behavior of the extensive margin of trade. We calibrate the parameters in the model using data on aggregate trade volumes and US consumer behavior. We then simulate the calibrated

¹²In the data, richer countries tend to trade more. One potential explanation is brought forward by Waugh (2010) and Tarasov (2012a), who argue that richer countries have systematically lower trade costs (specifically, Waugh (2010) focuses on cross-country variation in variable trade costs, while Tarasov (2012a) explores the role of fixed costs of trade in explaining the evidence). We capture this in the calibration by following Waugh (2010)'s approach to modeling variable trade costs.

model and compare the behavior of its extensive margin of bilateral trade with the data. The data are described in detail in Appendix A.6.

3.1 Parametrization of the utility function

We have shown above that the effects discussed emerge for a broad class of sub-utility functions $v(x)$. The central property is a bounded marginal utility, $v'(x) < \infty$. To quantify the model, we need to choose a certain parametric form for $v(x)$. We will use the Stone-Geary form

$$v(x) = \log(\bar{x} + x),$$

where $\bar{x} \geq 0$, as the thus parametrized model nests the standard homothetic model with $\bar{x} = 0$.¹³ The preference parameter \bar{x} represents the degree of non-homotheticity. In the context of this paper, it is particularly important that it governs the marginal utility of starting to consume an additional variety, $v'(0) = 1/\bar{x}$. The closer the non-homotheticity parameter \bar{x} gets to zero, the larger the marginal utility of consuming new varieties and thus the weaker the demand side effects on the extensive margin of trade. For $\bar{x} = 0$, the marginal utility approaches infinity and agents find consuming all available varieties optimal, no matter how expensive they are.

In the robustness section, we show that the results remain unchanged for alternative preferences specifications with bounded marginal utility such as quadratic preferences and CARA preferences.

3.2 Calibration strategy

The theoretical model's bilateral trade patterns are characterized by two moments - the extensive margin of trade and the aggregate trade volume. Volumes are governed by

$$X_{ni} = \pi_{ni} w_n L_n,$$

and the extensive margins are

$$m_{ni} = \pi_{ni} M_n.$$

The labor market clearing condition, $w_i L_i = \sum_{n=1}^N X_{ni}$, determines the equilibrium wage rates and the equilibrium extensive margins of consumption (that follow from the budget restriction (2.3)). In Appendix A.4, we show that with the Stone-Geary preferences, the budget restriction can be written as

$$w_n = \bar{x} (\Phi_n)^{-\frac{1}{\theta}} \left(M_n (-\log(1 - M_n))^{\frac{1}{\theta}} - \gamma \left(\frac{1}{\theta} + 1; -\log(1 - M_n) \right) \right), \quad (3.1)$$

¹³For $\bar{x} = 0$, the preferences become CES preferences with an elasticity of substitution of one. Since the quantitative behavior of the homothetic model is independent of the elasticity of substitution (see Alvarez and Lucas (2007)), the results represent the general CES model.

where $\gamma(z, \bar{t}) = \int_0^{\bar{t}} t^{z-1} e^{-t} dt$ is the incomplete Gamma function.

The model parameters are the countries' technologies T_i and populations L_i , the bilateral trade costs d_{ni} , the non-homotheticity parameter \bar{x} , and the trade elasticity θ . While the populations can be taken from the data, we need to calibrate the remaining parameters. In the following, we describe how we calibrate these parameters. The data used for the calibration is discussed en passant with a more complete description in Appendix A.6. We start with the calibration strategy for the non-homotheticity parameter, as this is the most novel part of the model. We then discuss the trade elasticity, trade costs, and technologies in turn.

3.2.1 Non-homotheticity parameter

To calibrate the non-homotheticity parameter, we use data from the US Consumer Expenditure Survey (CEX) for the year 2003. The advantage of using this source is its independence from the trade data. Based on this database, we construct the expenditures of approximately 3000 households for 107 different categories of tradable goods such as “Encyclopedia and other sets of reference books”, “Wigs, hairpieces, or toupees”, and “Winter sports equipment”. Details can be found in Appendix A.6. Counting the categories with positive expenditure gives us a measure for the extensive margin in consumption of a household. Table 2 reports the elasticity of this measure of the extensive margin with respect to total expenditures controlling for demographic variables such as household size, age of the reference person, and geography, i.e. rural/urban and region (Northeast, Midwest, South, West).

Table 2

Table 2 also reports the elasticities for a more liberal expenditure classification that comprises 186 categories.¹⁴ Clearly, the positive elasticity is robust across specifications lying between 0.4 and 0.5. For the calibration, we target the elasticity associated with the conservative classification and the specification including all the controls, 0.41.

To understand the details of the calibration, first note that only the relative technologies matter for trade volumes, which implies that we cannot identify the absolute level of the technologies using aggregate trade volumes. In the budget constraint (3.1), on the other hand, the absolute level of technologies is relevant since it determines the level of Φ_n . Moreover, note that scaling all technologies by a constant has the same effect as scaling the non-homotheticity parameter. We can therefore normalize either the level of technologies or the preference parameter. It is convenient to normalize the level of the technologies such that the US aggregator,

¹⁴The main difference lies in the treatment of housing- and gender-related categories. The conservative classification excludes these categories, whereas the liberal classification aggregates over categories that represent the same item but are differentiated by renter/owner or women/men/girls/boys.

Φ_{US} , equals one. Then, the budget constraint of a US-agent h with income $w_{US,h}$ is

$$w_{US,h} = \bar{x} \left(M_{US,h} (-\log(1 - M_{US,h}))^{\frac{1}{\theta}} - \gamma \left(\frac{1}{\theta} + 1; -\log(1 - M_{US,h}) \right) \right).$$

We set $\theta = 4.5$ (see the next section). For a given non-homotheticity parameter \bar{x} , we can feed all the CEX households' expenditures into the budget constraint and calculate the corresponding extensive margins of consumption. We then choose the preference parameter such that the resulting elasticity of the extensive margin of consumption matches the empirical CEX elasticity.^{15,16}

3.2.2 Trade elasticity

For the trade elasticity, we take the value estimated by Simonovska and Waugh (2012), $\theta = 4.5$. In general, one cannot identify the trade elasticity and the level of trade costs separately by estimating a gravity equation - the trade elasticity may be high and trade costs low or, conversely, the elasticity low and trade costs high. To tackle this problem, EK argue that one can use disaggregated price data from the World Bank's International Comparison Project (ICP) and take the maximal (or second highest) within good price difference as an estimate for bilateral trade costs. As the resulting value for trade costs has been obtained independently from trade volumes, one then can solve for the trade elasticity, which is implied by a gravity type regression. Simonovska and Waugh (2012) extend this approach. They provide a more elaborate estimation strategy that controls for an aggregation bias arising from the fact that it is very unlikely that the highest price difference represents actually the trade costs due to the small number of goods categories in the ICP (around 80). They also use a broader set of countries (123) and are thus able to check if the trade elasticity systematically varies with development level, which they find not to be the case. The structural framework for their estimation is the EK model. Since the present model behaves identically to the EK model with respect to volumes and prices, we can directly adapt the Simonovska and Waugh (2012) baseline estimate of $\theta = 4.5$.

¹⁵The resulting value is $\bar{x} = 3.14$. In contrast to the elasticity of substitution of CES preferences, for example, this parameter does not have a standalone interpretation (beyond the fact that it is not zero), since measures such as the demand elasticity or the elasticity of substitution change with income and prices. Therefore, \bar{x} is only meaningful when income and price distribution are known or, as in the case here, when a model giving rise to income and prices is parametrically specified.

¹⁶At first sight, the fact that we use within country inequality to calibrate the model where countries are populated by representative agents may seem surprising. We address this objection in the robustness section, where we consider the potential role of within-country inequality. We also refer to the well-established practice in the macro literature using micro-economically estimated elasticities to calibrate macro models populated by representative agents.

3.2.3 Trade costs

We calibrate the trade costs using aggregate bilateral trade volumes of the year 2003. In particular, we follow Waugh (2010) in modeling unobserved trade costs as a function of observed proxies and an exporter fixed effect:

$$d_{ni} = \exp \{ \delta_k + b + l + ex_i + \varepsilon_{ni} \},$$

where we suppressed the associated dummy variables for expositional simplicity. δ_k ($k = 1, \dots, 6$) is the effect of the bilateral distance between countries i and n lying in the k th distance interval. The intervals are (in miles): $[0, 375)$, $[375, 750)$, $[750, 1500)$, $[1500, 3000)$, $[3000, 6000)$, and $[6000, \infty)$. b is the effect of sharing a border, and l the effect of having the same language. ex_i is an exporter fixed effect that allows for asymmetry in bilateral trade costs and ε_{ni} captures all other trade barriers and is assumed to be orthogonal to the exporter fixed effects, distance, border, and language.

Normalizing the volume of the bilateral trade flow from i to n (equation (2.8)) with the importer's home sales X_{nn} yields

$$\frac{X_{ni}}{X_{nn}} = (d_{ni})^{-\theta} \frac{S_i}{S_n}, \quad (3.2)$$

where $S_i = T_i(w_i)^{-\theta}$ is a country fixed effect. The value of the bilateral trade flow X_{ni} is observed, while - in the context of the model - X_{nn} is simply a country's aggregate GDP less its total imports.¹⁷ Imposing the above trade cost function yields a Gravity-type equation, which we estimate using the Poisson pseudo-maximum likelihood estimator proposed by Silva and Tenreyro (2006).¹⁸ Table 3 reports the resulting estimates for the trade cost parameters together with the implied effects on trade costs (in percentage, relative to a country-pair with a bilateral distance between 0 and 375 miles without a common border and without a shared language) and some summary statistics on the estimated trade costs.

Table 3

The average and median trade costs among OECD countries are 2.01 and 1.89 respectively, which is slightly higher than the often cited 1.7 suggested by Anderson and Van Wincoop (2004) and very much in line with Waugh (2010). Consistent with the findings in the literature, trade costs among non-OECD countries are considerably higher.

¹⁷It is important to mention two potential caveats related to the way the model is mapped to the data. First, we use aggregate trade volumes, which include also non-consumption goods such as intermediates. Second, whereas trade is measured in gross values, GDP is measured as value added, which again in the context of intermediates may be of importance. We address both caveats in the robustness section where we extend the model by allowing for intermediates. It turns out that the main results remain basically the same. The reason for this may be that the share of consumption goods in trade flows is unrelated to per capita GDP and aggregate GDP, so that no systematic bias emerges.

¹⁸The reason we use the PPML estimator is that there are a number of zero bilateral trade flows in the data.

3.2.4 Technologies

The most straightforward approach to recovering technologies is to follow Fielser (2011) and take per capita incomes¹⁹ as a proxy for wages - indeed, wages and per capita incomes coincide in the model. Using the estimates for the country fixed effects, \hat{S}_i , one then can directly solve for the implied technologies, $\hat{T}_i = \hat{S}_i (w_i)^\theta$. An alternative approach disregards the country fixed effects and uses the market clearing conditions (2.9) to recover the technologies. Specifically, plugging the per capita incomes, the estimated trade costs, and the trade elasticity into the market clearing conditions allows us to solve for the unique set of technologies for which all markets clear. Figure 3 plots the expected productivity draw in a country, $E_i [z] = T_i^{1/\theta}$, against its per capita income.

Figure 3

Clearly, the two approaches yield very similar technologies. Moreover, as to be expected, there is a high correlation between estimated technology and observed incomes. Note, however, that the correlation is not perfect, as different geographic locations imply that countries with the same technology face different levels of demand and thus have different equilibrium incomes. For the remainder of the paper, we use the technologies calibrated using the market clearing conditions.²⁰

3.3 Calibration results

Given the calibrated parameters, we can now simulate the model and compare the behavior of its extensive margin to the data. Remember that with respect to volumes, the non-homothetic model behaves identically to the homothetic model. As we used volumes to calibrate technologies and trade costs, these calibrated values also apply to the homothetic model. That is, we obtain the homothetic model's predictions simply by setting the non-homotheticity parameter to zero, $\bar{x} = 0$.

We highlighted the positive correlation between the extensive margin of bilateral trade and the per capita incomes of the trading partners in the introduction (see Table 1). Table 4 reports the income elasticities that follow from repeating the regressions cited in the introduction using the data generated by the model.

Table 4

Both models yield the same elasticity with respect to the exporter income. Moreover, the elasticity is reasonably close to what is observed in the data. For the importer income elasticity, on the other hand, only the non-homothetic model's sign is consistent with the data. The non-homothetic model's elasticity is with 0.63 somewhat higher than the data's income elasticity

¹⁹In the context of the model, the use of nominal incomes is appropriate as deviations from PPP are endogenous in the EK framework.

²⁰The results for the alternative technologies are very similar, with differences in the outcomes typically lying within 1-2%.

of around 0.5, but still reasonably close. The homothetic model, on the other hand, predicts a negative importer income elasticity of -0.12. The reason for this counterfactual prediction is its negligence of the extensive margin of consumption meaning that only the negative effect of a technologically advanced country producing more varieties locally is present. In the non-homothetic model, this negative effect is dominated by a positive effect coming from the demand side - the expanding extensive margin of consumption. Figure 4 plots the calibrated extensive margin of consumption against per capita income.

Figure 4

Note that the relation is not perfect due to differences in the remoteness. Spain and New Zealand, for example, had the same nominal per capita income in 2003, but Spain's calibrated extensive margin of consumption is almost 10% higher than that of New Zealand as Spain's geographic location is much more favorable. This means that prices tend to be lower in Spain, so that its real income is higher and agents find consuming a broader set of varieties optimal.

3.4 The quantitative importance of the demand side

In this section, we perform two counterfactual experiments. The goal of these exercises is to demonstrate that accounting for demand side effects is quantitatively important when examining the reaction of the extensive margins of trade to changes in the economic environment. In each case, we start with the world economy as calibrated in the previous section and compare the counterfactual outcome to the initial situation.

3.4.1 The rise of China and India

One of the most important trends in the global economy is the rise of China and India. These two countries experienced spectacular growth rates in the recent years - according to the World Development Indicators (World Bank (2010)), China's per capita income almost doubled (95%) relative to the world per capita income between 1993 and 2003, and India's per capita income grew by 31% relative to world per capita income. In this experiment, we consider the calibrated 2003 world economy and ask how trade patterns would change if China's and India's technologies were to improve such that their incomes rose again by the same magnitudes relative to the world income.

The homothetic model predicts that, as China's and India's technologies improve, these two countries become competitive in more varieties. These countries will therefore export a broader set of varieties to their trading partners (the median increase is 64% for China and 19% for India). At the same time, the better technology implies also that these countries will import a narrower set of varieties - China is predicted to source 44% more varieties locally, whereas India's home share extends by 5%. This translates into a median decrease in the extensive

margin of bilateral import flows of 19% and 14% for China and India, respectively.

These predictions change when one allows for non-homothetic consumer behavior. In the non-homothetic model, China and India extend their extensive margin of consumption as their incomes increase (83% in China and 34% in India), so that the measures of imported varieties actually increase by 47% and 29%. This translates into a median increase in the extensive margins of bilateral import flows of 49% for China and 15% for India.

3.4.2 Changes in trade costs

In the second experiment, we consider the effect of reduction in trade costs. Lower trade costs affect the extensive margin of trade through two channels. The first channel is the standard supply side channel - lower trade costs imply that trade becomes worthwhile for more varieties, i.e. the extensive margin of trade will expand. The second channel operates through the demand side - lower trade costs lead to lower prices,²¹ which increases real incomes. With higher incomes, agents will find consuming a broader set of varieties optimal, which in turn will tend to increase the extensive margins of trade. In order to assess the relative importance of these two channels, we take again the calibrated 2003 world economy and uniformly decrease trade costs by 10, 25, and 50 percent.

From Table 5 reporting summary statistics on the percentage changes in the extensive margins of bilateral trade, the starkly differing predictions of the homothetic and the non-homothetic model become apparent.

Table 5

The homothetic model, for example, predicts that more than half of country-pairs actually experience a reduction in the bilateral extensive margins in case of a 25%-reduction in trade costs. This is because low-tech supplier countries were competitive in nearby markets in the initial situation with high trade costs, but lower trade costs imply that they are now dominated by countries with better technologies such that their export margins tend to fall. In the non-homothetic model, this supply-side effect is attenuated by the demand-side effect of a rising extensive margin of consumption, the absolute and relative changes in which are plotted in Figure 5 against per capita income.

Figure 5

While poor countries experience the highest relative changes, middle income countries have the highest absolute changes. Relative to the supply side effect, these changes are so large, for example, that the number of country pairs experiencing decreasing extensive margins falls to 15% and the predicted median change is an increase of 22% instead of a decrease of 1% in the homothetic model. These contrasting predictions demonstrate that accounting for demand side

²¹In the above experiments, we choose the US wage as the numéraire. Therefore, we describe the adjustment via lower prices. Alternatively, with some goods price as the numéraire, the corresponding adjustment would be described by lower trade costs increasing productivity and hence wages.

effects is indeed quantitatively important when thinking about the extensive margin of bilateral trade.

4 How general is the proposed channel?

We have presented a model of international trade where agents adjust their extensive margin of consumption with income, which has quantitatively important effects for the extensive margin of bilateral trade. In order to highlight this novel demand side channel, we kept the supply side very simple by adapting a perfectly competitive Ricardian framework. We found that, when allowing for non-homothetic consumer behavior, the EK model is not only able to capture the pattern of aggregate trade volumes, but also the behavior of the extensive margin of trade. The more general message is that the extensive margin of trade may be strongly driven through differences in demand pattern across countries. In the following, we discuss how this channel generalizes to richer models and why accounting for the demand side is potentially important.

A richer framework of international trade would model the firm explicitly by allowing for market power as in Krugman (1980) and Melitz (2003). Fixed market entry costs imply that not all firms find it optimal to enter all markets, which gives rise to the extensive margin of trade. In particular, for a given level of entry costs, entering bigger markets is more attractive, since the contribution margin in these markets is larger. With homothetic preferences, the notion of a “big market” is purely driven by aggregate GDP - it does not matter if we have a large and poor population or a small and rich population due to constant expenditure shares. With non-homothetic preferences, however, the decomposition of aggregate GDP becomes relevant as poor agents adjust their expenditure shares with income. This is particularly apparent in the model developed above when the expenditure share goes from zero (when the non-negativity constraint is binding) to some positive share - the emerging extensive margin of consumption then drives the extensive margin of trade. In a more general setting with market entry costs, not only the bounded marginal utility is relevant, but also the fact that poor agents concentrate their expenditures on relatively cheap goods. Thus, a small but rich market may be sufficiently “big”, whereas demand in a poor but populous market (with the same aggregate GDP) may be too low for the operating profits to cover market entry costs - accordingly, the small but rich market’s extensive margin of bilateral trade will tend to be higher than that in the large but poor market. Clearly, not only average income, but also the entire shape of the income distribution is relevant for a firm’s entry decision in such a framework. Ignoring non-homothetic demand thus leads a researcher to attribute differences in the extensive margins of import flows entirely to differences in market entry costs, while a considerable part of the differences may be driven by asymmetries in average income and income distribution (which is in fact supported by the above findings).

5 Robustness and extensions

5.1 Extending the model to trade in intermediates

The model we developed above is one of consumption goods only. We chose to abstract from intermediate goods to keep the model as simple as possible. In this section, we outline a model with intermediate goods and final goods - both tradable. The purpose of this extension is twofold: First, we use the extended model to assess if abstracting from intermediate goods introduces a significant bias in the quantification. Second, we used trade and GDP data to compute the left hand side of equation (3.2). However, many papers using the EK framework are models of trade in intermediates, in the context of which gross manufacturing output is used instead of GDP data. With the intermediate extension, we need to use gross manufacturing output, which allows us to assess the extent to which the results were driven by these different ways of computing the normalized trade flows.

There are two industries in the extended model, $\iota = I, F$. I produces tradable intermediate goods and F produces tradable final goods.²² Both industries bundle labor and a CES-aggregate of intermediates using a Cobb-Douglas production technology with labor share β . Cost minimization implies that the price at which country i can offer a industry ι -variety j_ι in market n is

$$p_{ni}(j_\iota) = \frac{d_{ni}}{z_i(j_\iota)} w_i^\beta P_i^{1-\beta},$$

where $P_i = \left(\int_0^1 p_i(j_I)^{1-\sigma} dj_I \right)^{1/(1-\sigma)}$ is the CES price index ($p_i(j_I)$ is the price of variety j_I in country i).²³ Assuming Fréchet distributed productivities with the same parameters across the industries yields a gravity-like expression that looks in its reduced form similar to that derived above

$$\frac{X_{ni}}{X_{nn}} = (d_{ni})^{-\theta} \frac{S_i}{S_n}.$$

However, there are two crucial differences. First, the total expenditures, X_n , are now the total intermediate absorption, X_n^I , plus the total expenditures on final goods, X_n^F . Consequently, the home supply must now be imputed by subtracting country's total manufacturing exports from its gross manufacturing output, which we mostly get from UNIDO (2003) (details in Appendix A.6). Data constraints reduce the sample to 71 countries. Second, the country fixed effects now include the countries' intermediate price indices: $S_i = T_i \left(w_i^\beta P_i^{1-\beta} \right)^{-\theta}$. Note that in the simple model without intermediates the price indices P_i were absorbed into the calibrated technologies and wages entered with an exponent of one. This implies that the model tended to overstate the dispersion in technologies (the standard deviation in log of the calibrated technologies is 2.7 with intermediates vs. 7.5 without intermediates). Explicitly considering intermediates may particularly matter for counterfactual experiments changing the trade costs, as this implies

²²Note that it is usually assumed that final goods are non-traded in EK and follow up papers.

²³We assume that the trade costs and labor share are the same across the two industries.

potentially large changes in the intermediate price indices.

Using the same procedure as in the main text (but a different measure for X_n), we can estimate the trade costs and the country fixed effects. There are again two approaches to recover technologies. The more direct approach uses the estimated country fixed effects, \hat{S}_i and trade costs, \hat{d}_{ni} , to compute the implied price indices, $\hat{P}_n = \left(\sum_{i=1}^N T_i \left(w_i^\beta P_i^{1-\beta} d_{ni} \right)^{-\theta} \right)^{-1/\theta} = \left(\sum_{i=1}^N \hat{S}_i \left(\hat{d}_{ni} \right)^{-\theta} \right)^{-1/\theta}$, and then uses these price indices together with the per capita incomes, w_i , and the calibrated values for β and θ to solve for the implied technologies, $\hat{T}_i = \hat{S}_i \left(w_i^\beta \hat{P}_i^{1-\beta} \right)^\theta$.²⁴ Alternatively, one can combine the estimated trade costs and the per capita incomes and solve directly for the unique set of technologies for which all markets clear.²⁵ The correlation (in logs) between the technologies thus calibrated is very high at 0.96. In the counterfactual experiments, we use the technologies based on imposing market clearing.

The price distribution of final goods in country n is

$$G_n(p) = 1 - \exp \left\{ -p^\theta \sum_{i=1}^N T_i \left(w_i^\beta P_i^{1-\beta} d_{ni} \right)^{-\theta} \right\}.$$

Normalizing technologies such that $\sum_{i=1}^N T_i \left(w_i^\beta P_i^{1-\beta} d_{USi} \right)^{-\theta} = 1$, we can use the same value for the non-homotheticity parameter as above. Simulating the model thus calibrated yields the importer income elasticity of the extensive margin of bilateral trade of 0.58 (the corresponding elasticity in the model without intermediates is 0.57; the empirical elasticity is 0.50 in the restricted sample of 71 countries). Considering the counterfactual experiments, the model with intermediates generally features even stronger differences between the predictions of the homothetic and the non-homothetic model. This is because the intermediate price index reacts to changes in trade costs and technologies amplifying the reaction of the price distribution of final goods, and agents thus tend to adjust their extensive margins of consumption more strongly.

5.2 Inequality

Up to now, we abstracted from within-country inequality and had each country populated by representative agents. Although most of the global inequality is indeed explained by differences in average incomes, within-country inequality is a potentially important determinant for the extensive margin of trade. Consider, for instance, two countries with the same average income, but one with a wider range of the income distribution. The model predicts that the country

²⁴In the calibration, β is set to 1/3. This value is derived by calculating the average labor share in UNIDO (2003).

²⁵Specifically, we take an initial guess for the technology vector, compute the implied price indices, and use these together with the initial guess to compute π_{ni} and then the countries' balances of payments. We adjust the technology guess using a tâtonnement-like algorithm until all markets clear.

with richer agents has a broader extensive margin of consumption and thus tends to import more varieties.²⁶

Taking the present model literally implies that if each country features one very rich agent, then all countries' extensive margins of consumption were one and all the demand side effects disappeared. That is, the model's extensive margin is very sensitive to the upper tail of the income distribution. The main explanation of why this may not be the case is based on the presence of fixed market entry costs (beachhead costs). However, incorporating these costs into the present framework would require departure from the competitive setting to allow for positive markups that can be used to cover the beachhead costs. This in turn would disproportionately complicate the model.²⁷

Nonetheless, in order to get a feeling for the potential importance of the within-country income distribution, we propose a simple exercise that allows us to stay within the Ricardian framework. In particular, we choose to use the average income in the top quintile to compute the extensive margin of consumption. Remember that conditional on entering market n , the price distribution is the same across supplier countries. This in turn implies that the share of consumer income spent on products from country i does not depend on income and equals to π_{ni} .²⁸ Consequently, the aggregate value of the flow from i to n is still $X_{ni} = \pi_{ni}w_nL_n$, i.e. aggregate volumes do not depend on the income distribution. As a result, we can use the same calibration strategy as above.

The non-homotheticity parameter \bar{x} calibrated using the CEX data remains unchanged. However, we acknowledge the presence of within-country inequality by using the average income among the top quintile in the budget constraint (3.1) instead of per capita income. Taking the model literally, this amounts to allowing for a general income distribution, which is bounded by the average income in the top quintile. In a more general sense, we learn from this exercise how the results change when trying to account for differences at the top of countries' income distributions.

²⁶Indeed, when repeating the regressions cited in the introduction including the importer's top quintile of the income distribution, we get positive coefficients for the top quintile. The elasticity considering consumption goods only is 0.23 and significant at the 1% level, whereas the elasticity for all types of goods is lower at 0.15 and significant only at the 10% level.

²⁷Markups become endogenous with non-homothetic preferences. With a representative agent, the model still preserves some tractability (see Simonovska (2010)), as there is one cutoff productivity per market above which firms enter this market and below which firms abstain from entering. However, if one introduces within-country inequality, firms not only decide whether to enter a market or not, but also whom to supply in this market. The equilibrium in this case is characterized by a correspondence between a firm productivity and the income of the agent, who consumes this firm's variety at the optimal quantity of zero. Unfortunately, there is no simple expression for this correspondence in general (see Tarasov (2012b) for a special case with two countries and two income classes).

²⁸Consider consumer h with income w_n^h in country n . Define her extensive margin of consumption by M_n^h . Then, it is straightforward to show that the number of products produced in country i the consumer buys is $\pi_{ni}M_n^h$ (the argument is exactly the same as in the case with a representative agent). As, conditional on entering market n , the price distribution is the same across supplier countries, the expenditures per product of consumer h are w_n^h/M_n^h . As a result, the expenditures on products produced by country i are $\pi_{ni}M_n^h$ multiplied by w_n^h/M_n^h . This implies that, irrespective of income, consumers spend equal shares on products produced by different countries.

We use quintile data from UNU-WIDER (2008) that are described in detail in the Appendix A.1. As the quintiles are not available for the full sample, we consider a smaller sample of 112 countries. We then reestimate the model parameters for this smaller sample of countries.²⁹ We find that accounting for within-country inequality by using average incomes in the top quintiles in the budget constraint yields a lower importer income elasticity of 0.52 (vs. 0.46 in the data).

5.3 Alternative utility functions

In the theory part, we worked with a general utility function with the crucial feature of a bounded marginal utility. We then had to assume some particular functional form for the utility function (Stone-Geary) to quantify the model. This section considers two alternative utility functions with bounded marginal utility and shows that the calibration results are robust to changes in the functional form of the utility function. As our calibration strategy for the demand side targets only one moment, we focus on one-parameter utility functions.

In particular, we consider a quadratic utility

$$v^{quadr}(x) = x - \frac{1}{2}a^{quadr}x^2,$$

which is popular for its linear demand function and a constant absolute risk aversion utility (CARA)

$$v^{cara}(x) = -\exp\{-a^{cara}x\}.$$

A reader might note that often these utility functions are written with three parameters.³⁰ However, an aggregate utility function has a purely ordinal purpose in the context of a static trade model, so that all monotonic transformations of the utility function, $U = \int v(x(j)) dj$, yield the same economic behavior.³¹ The one-parameter versions above are simply linear transformations of the three parameter versions often seen.

As aggregate trade volumes do not depend on the particular functional form of the utility function, the supply side parameters calibrated above (trade elasticity, trade costs, technologies) still apply. Thus, we only need to recalibrate the demand side parameter. In Appendix A.5, we present the analogues to equation (3.1) governing the extensive margin of consumption for a given income and price distribution. Using these equations, we calibrate the new utility parameters by targeting US consumers' income elasticity of the extensive margin consumption.³²

We then simulate the calibrated models and calculate the income elasticities of the extensive margin. The resulting exporter income elasticities are the same as above, as they do not depend on the demand side of the model. The importer income elasticities, on the other hand, crucially

²⁹For this smaller sample of countries, the calibrated importer income elasticity in the representative agent model barely changes (0.62 instead of 0.63).

³⁰ $v^{cara}(x) = B^{cara} - C^{cara} \exp\{-a^{cara}x\}$ and $v^{quadr}(x) = B^{quadr} + C^{quadr}x - 1/2a^{quadr}x^2$.

³¹It is important to note that the transformation is applied to the aggregate utility function, U , and not directly to the sub-utility function.

³²The resulting parameters are $a^{quadr} = 0.28$ and $a^{cara} = 0.33$.

depend on the demand side as demonstrated above when comparing the homothetic model with the non-homothetic model. However, when considering the alternative non-homothetic utility functions, the importer income elasticities change only very little (0.618 for CARA and 0.623 for quadratic preferences instead of 0.63 for Stone-Geary). Similarly, the quantitative effects in the counterfactual experiments do not change significantly. These results demonstrate that the quantitative behavior of the model does not seem to depend very much on the particular functional form of the sub-utility function.

6 Conclusion

In this paper, we discuss an importing country's demand structure as a determinant of the extensive margin of bilateral trade. We draw on the evidence of microeconomic studies that show that richer agents consume more varieties. Allowing for such an extensive margin of consumption in an otherwise standard Ricardian trade model offers an explanation for the positive correlation between the extensive margin of bilateral trade and importer's per capita income. We then quantify the model using data on aggregate trade volumes and US consumer behavior. We find that the extensive margin of trade generated by the model behaves similarly to what we observe in the data. Two counterfactual experiments demonstrate that this novel demand side channel is quantitatively important.

We mention in the introduction that other authors have used non-homothetic preferences to discuss different aspects of the pattern of international trade such as aggregate trade volumes and quality. A potentially fruitful avenue for future research is a model where these two aspects and the extensive margin of trade could be analyzed simultaneously. On the demand side, such a framework would feature agents who adjust their consumption decision at the intensive, the extensive, and the quality margin. On the supply side, variations in countries' abilities to produce quality goods would introduce comparative advantages. Trade pattern – volumes, extensive margin, and quality – could then be analyzed as the result of interactions of the exporter country's production structure and the importer country's demand pattern.

References

- ALVAREZ, F., AND R. LUCAS (2007): “General equilibrium analysis of the Eaton-Kortum model of international trade,” *Journal of Monetary Economics*, 54(6), 1726–1768.
- ANDERSON, J., AND E. VAN WINCOOP (2004): “Trade costs,” *Journal of Economic Literature*, 42(3), 691–751.
- BALDWIN, R., AND J. HARRIGAN (2011): “Zeros, quality and space: trade theory and trade evidence,” *American Economic Journal: Microeconomics*, 3(2), 60–88.
- BLS (2003): “2003 consumer expenditure interview survey, public use microdata, documentation,” U.S. Department of Labor, Bureau of Labor Statistics, Division of Consumer Expenditure Surveys.
- BRODA, C., AND D. WEINSTEIN (2006): “Globalization and the gains from variety,” *Quarterly Journal of Economics*, 121(2), 541–585.
- CHANEY, T. (2008): “Distorted gravity: the intensive and extensive margins of international trade,” *American Economic Review*, 98(4), 1707–1721.
- CHOI, Y., D. HUMMELS, AND C. XIANG (2009): “Explaining import quality: the role of the income distribution,” *Journal of International Economics*, 78(2), 293–303.
- COSTINOT, A., AND I. KOMUNJER (2007): “What goods do countries trade? New Ricardian predictions,” NBER Working Paper No. 13691.
- DOLLAR, D., AND A. KRAAY (2002): “Growth is good for the poor,” *Journal of economic growth*, 7(3), 195–225.
- EATON, J., AND S. KORTUM (2002): “Technology, geography, and trade,” *Econometrica*, 70(5), 1741–1779.
- EATON, J., AND S. KORTUM (2005): “Technology in the global economy: A framework for quantitative analysis,” Book manuscript, New York University and University of Chicago.
- EATON, J., S. KORTUM, AND F. KRAMARZ (2004): “Dissecting trade: firms, industries, and export destinations,” *American Economic Review*, 94(2), 150–154.
- FAJGELBAUM, P., G. GROSSMAN, AND E. HELPMAN (2011): “Income distribution, product quality, and international trade,” *Journal of Political Economy*, 119(4), 721–765.
- FALKINGER, J., AND J. ZWEIMÜLLER (1996): “The cross-country Engel curve for product diversification,” *Structural Change and Economic Dynamics*, 7(1), 79–97.

- FEENSTRA, R. (1994): “New product varieties and the measurement of international prices,” *American Economic Review*, 84(1), 157–177.
- FIELER, A. C. (2011): “Non-homotheticity and bilateral trade: evidence and a quantitative explanation,” *Econometrica*, 79(4), 1069–1101.
- FOELLM, R., C. HEPENSTRICK, AND J. ZWEIMÜLLER (2010): “Non-homothetic preferences, parallel imports and the extensive margin of international trade,” CEPR Discussion Paper No. 7939.
- GAULIER, G., S. ZIGNAGO, D. SONDJIO, A. SISSOKO, AND R. PAILLACAR (2010): “BACI: a world database of international trade at the product-level, 1995-2007 version,” Centre d’Etudes Prospectives et d’Informations Internationales Working Paper No. 2010-23.
- GRONAU, R., AND D. HAMERMESH (2008): “The demand for variety: a household production perspective,” *Review of Economics and Statistics*, 90(3), 562–572.
- HUMMELS, D., AND P. KLENOW (2002): “The variety and quality of a nation’s trade,” NBER Working Paper No. 8712.
- JACKSON, L. (1984): “Hierarchic demand and the Engel curve for variety,” *Review of Economics and Statistics*, pp. 8–15.
- JEKANOWSKI, M., AND J. BINKLEY (2000): “Food purchase diversity across US markets,” *Agribusiness*, 16(4), 417–433.
- KRUGMAN, P. (1980): “Scale economies, product differentiation, and the pattern of trade,” *American Economic Review*, 70(5), 950–959.
- MARKUSEN, J. (2010): “Putting per-capita income back into trade theory,” NBER Working Paper No. 15903.
- MAS-COLELL, A., M. WHINSTON, AND J. GREEN (1995): *Microeconomic theory*. Oxford University Press, New York.
- MATSUYAMA, K. (2000): “A Ricardian model with a continuum of goods under nonhomothetic preferences: demand complementarities, income distribution, and North-South trade,” *Journal of Political Economy*, 108(6), 1093–1120.
- MELITZ, M. (2003): “The impact of trade on intra-industry reallocations and aggregate industry productivity,” *Econometrica*, 71(6), 1695–1725.
- MOON, W., W. FLORKOWSKI, L. BEUCHAT, A. RESURRECCION, P. PARASKOVA, J. JORDANOV, AND M. CHINNAN (2002): “Demand for food variety in an emerging market economy,” *Applied Economics*, 34(5), 573–581.

- SAURÉ, P. (2012): “Bounded love of variety and patterns of trade,” *Open Economies Review*, 23, 645–674.
- SILVA, J., AND S. TENREYRO (2006): “The log of gravity,” *Review of Economics and Statistics*, 88(4), 641–658.
- SIMONOVSKA, I. (2010): “Income differences and prices of tradables,” NBER working paper No. 16233.
- SIMONOVSKA, I., AND M. WAUGH (2012): “The elasticity of trade: Estimates and evidence,” UC Davis Working Paper No. 11-2.
- TARASOV, A. (2012a): “Per capita income, market access costs, and trade volumes,” *Journal of International Economics*, 86(2), 284–294.
- (2012b): “Trade Liberalization and Welfare Inequality: A Demand-Based Approach,” *Scandinavian Journal of Economics*, forthcoming.
- THIELE, S., AND C. WEISS (2003): “Consumer demand for food diversity: evidence for Germany,” *Food Policy*, 28(2), 99–115.
- UNIDO (2003): *Industrial statistics database*. United Nations Industrial Development Organization.
- UNU-WIDER (2008): “World income inequality database, version 2.0c,” World Institute for Development Economics Research of the United Nations University.
- WAUGH, M. (2010): “International trade and income differences,” *American Economic Review*, 100(5), 2093–2124.
- WORLD BANK (2010): *World development indicators*. World Bank.

A Appendix

A.1 Derivation of the Country Specific Price Distribution $G_n(p)$

Using the productivity distribution and the pricing equation (2.4), the probability of country i supplying a particular variety j at a price lower than p in market n can be written as

$$\begin{aligned} G_{ni}(p) &= \Pr [P_{ni} \leq p] = \Pr \left[\frac{w_i d_{ni}}{Z_i(j)} \leq p \right] = \Pr \left[\frac{w_i d_{ni}}{p} \leq Z_i(j) \right] \\ &= 1 - \exp \left\{ -T_i (w_i d_{ni})^{-\theta} p^\theta \right\}. \end{aligned}$$

The probability that the lowest price on offer in market n is below p is the complement to the probability that all offered prices lie above p

$$G_n(p) = \Pr \left[\min \{P_{ni}(j)\}_{i=1}^N \leq p \right] = 1 - \Pr \left[\min \{P_{ni}(j)\}_{i=1}^N > p \right].$$

As the productivity draws are iid across countries, this probability is simply the product of the individual probabilities, which yields the price distribution from the main text:

$$G_n(p) = 1 - \prod_{i=1}^N \Pr [P_{ni}(j) > p] = 1 - \exp \left\{ -p^\theta \sum_{i=1}^N T_i (w_i d_{ni})^{-\theta} \right\}.$$

A.2 Derivation of the Trade Share π_{ni}

The probability that country i is the cheapest supplier for variety j in market n is given by

$$\pi_{ni}(j) = \Pr \left[P_{ni}(j) < \min \{P_{nk}(j)\}_{k \neq i} \right] = \int_0^\infty \Pr \left[p < \min \{P_{nk}(j)\}_{k \neq i} \right] dG_{ni}(p).$$

Again, one can write the distribution of the minimum price as the product of the individual distributions

$$\begin{aligned} \int_0^\infty \Pr \left[p < \min \{P_{nk}(j)\}_{k \neq i} \right] dG_{ni}(p) &= \int_0^\infty \prod_{k \neq i} \Pr [p < P_{nk}(j)] dG_{ni}(p) \\ &= \int_0^\infty \prod_{k \neq i} [1 - G_{nk}(p)] dG_{ni}(p). \end{aligned}$$

Inserting for the price distributions yields

$$\pi_{ni}(j) = \int_0^\infty \prod_{k \neq i} [1 - G_{nk}(p)] dG_{ni}(p) = \frac{T_i (w_i d_{ni})^{-\theta}}{\sum_{k=1}^N T_k (w_k d_{nk})^{-\theta}}.$$

Note that this probability does not depend on the index j , so that it also represents the

share of varieties for which country i is the cheapest supplier in n

$$\pi_{ni}(j) = \pi_{ni} = \frac{T_i (w_i d_{ni})^{-\theta}}{\Phi_n}.$$

Specifically, the share of varieties for which country i is the cheapest supplier in n is given by $\int_0^1 I\left(P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right) dj$, where $I(\cdot)$ is the indicator function. Using the law of large numbers,

$$\begin{aligned} \int_0^1 I\left(P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right) dj &= E\left\{I\left(P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right)\right\} \\ &= \Pr\left\{P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right\} = \pi_{ni}. \end{aligned}$$

A.3 Conditional on Entry, Price Distributions are the Same Across Sources

The distribution of prices from country i in market n conditional on being the cheapest supplier is

$$\Pr\left[P_{ni}(j) \leq p \mid P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right] = \frac{\Pr\left[P_{ni}(j) \leq p, P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right]}{\Pr\left[P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right]}.$$

The denominator is π_{ni} . The nominator can be written as

$$\Pr\left[P_{ni}(j) \leq p, P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right] = \int_0^p \Pr\left[x < \min\{P_{nk}(j)\}_{k \neq i}\right] dG_{ni}(x).$$

Recall that

$$\Pr\left[x < \min\{P_{nk}(j)\}_{k \neq i}\right] = \prod_{k \neq i} [1 - G_{nk}(p)].$$

Thus,

$$\Pr\left[P_{ni}(j) \leq p, P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right] = \int_0^p \prod_{k \neq i} [1 - G_{nk}(x)] dG_{ni}(x) = \pi_{ni} G_n(p).$$

Reinserting this into the initial expression completes the proof:

$$\Pr\left[P_{ni}(j) \leq p \mid P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}\right] = G_n(p) = \Pr\left[\min\{P_{nk}(j)\}_{k=1}^N \leq p\right].$$

Note that the number of varieties exported by country i to country n is given by

$$\int_0^1 I\left(P_{ni}(j) < \min\{P_{nk}(j)\}_{k \neq i}, P_{ni}(j) < v'(0)/\lambda_n\right) dj.$$

Using the law of large numbers, the latter can be rewritten as follows:

$$\Pr \left[P_{ni}(j) < \min \{P_{nk}(j)\}_{k \neq i}, P_{ni}(j) < v'(0)/\lambda_n \right].$$

Using the previous analysis, one can derive that

$$\Pr \left[P_{ni}(j) < \min \{P_{nk}(j)\}_{k \neq i}, P_{ni}(j) < v'(0)/\lambda_n \right] = \pi_{ni} G_n(v'(0)/\lambda_n) = \pi_{ni} M_n.$$

A.4 Derivation of the Budget Constraint for Stone-Geary Preferences

First note that with the Stone-Geary utility, the first order conditions (2.1) become

$$\begin{aligned} \frac{1}{x(j) + \bar{x}} &= \lambda p(j) \quad \text{for } x(j) > 0 \\ \frac{1}{\bar{x}} &< \lambda p(j) \quad \text{for } x(j) = 0. \end{aligned}$$

Using these first order conditions, we can solve for the price of the marginal variety

$$p(M) = \frac{v'(0)}{\lambda} = \frac{1}{\bar{x}\lambda}$$

and for the inverse of the marginal utility of income respectively

$$\frac{1}{\lambda} = \bar{x}p(M).$$

Optimal expenditures for varieties $j < M$ are

$$p(j)x(j) = \frac{1}{\lambda} - \bar{x}p(j) = \bar{x}(p(M) - p(j)).$$

Inserting this into a country n agent's budget restriction (2.3) yields

$$w_n = \bar{x} \left(p(M_n) M_n - \int_0^{p(M_n)} p dG_n(p) \right).$$

Using the country specific price distribution $G_n(p)$, one can write

$$p(M_n) = G_n^{-1}(M_n) = \left(-\frac{\log(1 - M_n)}{\Phi_n} \right)^{\frac{1}{\theta}}$$

and

$$dG_n(p) = \theta p^{\theta-1} \Phi_n \exp \{ -p^\theta \Phi_n \} dp.$$

Substituting this into the budget constraint yields

$$w_n = \bar{x} \left(\left(-\frac{\log(1 - M_n)}{\Phi_n} \right)^{\frac{1}{\theta}} M_n - \int_0^{p(M_n)} p \theta p^{\theta-1} \Phi_n \exp \{-p^\theta \Phi_n\} dp \right).$$

Changing variables in the integral, $t = p^\theta \Phi_n$, we derive that

$$w_n = \bar{x} \left(\left(-\frac{\log(1 - M_n)}{\Phi_n} \right)^{\frac{1}{\theta}} M_n - \int_0^{p(M_n)^\theta \Phi_n} \left(\frac{t}{\Phi_n} \right)^{\frac{1}{\theta}} \exp \{-t\} dt \right),$$

where the integral equals the incomplete Gamma function so that we can write

$$w_n = \bar{x} \left(\left(-\frac{\log(1 - M_n)}{\Phi_n} \right)^{\frac{1}{\theta}} M_n - \left(\frac{1}{\Phi_n} \right)^{\frac{1}{\theta}} \gamma \left(\frac{1}{\theta} + 1, p(M_n)^\theta \Phi_n \right) \right).$$

Substituting for the price of the marginal variety and rearranging finally lead to the expression of the main text.

A.5 The Budget Constraints for Alternative Utility Functions

The steps to derive the budget constraints are very similar to those outlined in Appendix A.4. Therefore, we omit their derivations. For a given income, E , and price distribution, $G(p) = 1 - \exp \{-\Phi p^\theta\}$, the budget constraint with quadratic utility is

$$E = \frac{(\Phi)^{-\frac{1}{\theta}}}{a^{quadr}} \gamma \left(\frac{1}{\theta} + 1; -\log(1 - M) \right) - \frac{(\Phi)^{-\frac{1}{\theta}}}{a^{quadr}} (-\log(1 - M))^{-\frac{1}{\theta}} \gamma \left(\frac{2}{\theta} + 1; -\log(1 - M) \right).$$

With CARA preferences, the budget constraint spells

$$E = -\frac{\theta \Phi}{a^{cara}} \int_0^{(-\log(1-M)/\Phi)^{\frac{1}{\theta}}} \log \left(\left(-\frac{\log(1 - M)}{\Phi} \right)^{-\frac{1}{\theta}} p \right) p^\theta \exp(-\Phi p^\theta) dp.$$

A.6 Data

We use data for the year 2003. In the baseline specification, the sample consists of 164 countries, which corresponds to $26732 = 164 * 163$ bilateral trade relations. In the following, we describe the sources of the data used in the quantification.

A.6.1 Aggregate value of bilateral trade

We use the COMTRADE trade data of the year 2003 as provided by CEPII (Gaulier, Zignago, Sondjo, Sissoko, and Paillacar, 2010). This data set provides the dollar values of the bilateral trade flows between 239 economic entities (mostly countries) on the HS6 level of aggregation $X_{ni}(j)$, which corresponds to 5111 goods categories. Summing over all HS6 categories, we get the aggregate value of a bilateral trade flow from exporting country i to importing country n , $X_{ni} = \sum_{j=1}^{5111} X_{ni}(j)$.

A.6.2 Extensive margin of bilateral trade

We use a simple and intuitive measure for the extensive margin of bilateral trade, which counts the number of HS6 categories with positive volumes

$$m_{ni} = \sum_j I(X_{ni}(j) > 0),$$

where $I(X_{ni}(j) > 0)$ is an indicator function taking the value of one if the bilateral trade flow from i to n in the HS6-category j is positive. A potential drawback of this measure is the fact that the HS6 categories are defined for customs purposes. As a result, heavily regulated goods tend to have more categories. The associated measurement error is absorbed into the error term and the estimated elasticities are unbiased, if the coarseness of the traded HS6-categories is orthogonal to the regressors.

An alternative measure for the extensive margin is brought forward by Broda and Weinstein (2006)

$$m_{ni}^{BW} = \frac{\sum_j X_n(j) I(X_{ni}(j) > 0)}{\sum_j X_n(j)},$$

where $X_n(j) = \sum_{k \neq n} X_{nk}(j)$ is the value of country n 's total imports in category j . One of the main advantages of this measure is that the categories are weighted, which may alleviate measurement errors due to differences in the coarseness of the categorization. However, in the present context, this measure of the extensive margin may be inappropriate for two reasons. First, this measure is derived using a CES demand system (Feenstra, 1994), whereas the central assumption in our model is that preferences are non-CES. Second and more important, in the context of the present model, the numerator represents the trade flows from i to n , X_{ni} , while the denominator is the n 's total imports in all categories, $\sum_{k \neq n} X_{nk}$. Using the corresponding expressions from the model, one can see that $m_{ni}^{BW} = \pi_{ni}/(1 - \pi_{nn})$. That is, the central element of the theory - the extensive margin of consumption - cancels, implying that the Broda and Weinstein (2006) measure is unlikely to reflect the income effects coming from the non-homothetic consumer behavior.

A.6.3 Per capita incomes and population sizes

The per capita incomes and the population sizes are taken from the World Bank’s World Development Indicator. The per capita incomes are measured in current (year 2003) US-dollars. Following EK, we deliberately abstain from using purchasing power adjusted incomes as deviations from PPP arise endogenously in the EK framework.

A.6.4 Bilateral distances, shared border, and common language

All transportation cost proxies are from the database provided by CEPII. The bilateral distance is measured as the distance between two countries’ most populous cities. The common language indicator takes the value one if two countries have the same official language³³ and the common border dummy takes the value one if two countries share a common land border.

A.6.5 CEX

The US consumer expenditure survey (CEX) is a rotating panel collected by the Bureau of Labor Statistics (BLS). Its “interview survey” part provides detailed information on household characteristics and expenditures. One of the main purposes of the survey is its use in determining and revising the baskets that are used for the computation of the consumer price index. We obtain the CEX data from the website of the Inter-University Consortium for Political and Social Research (ICPSR). A detailed documentation of the data can be found in BLS (2003). In the following, we briefly discuss the raw data and how we processed the raw data.

The unit of observation in the CEX is a “consumer unit”, CU, which basically comprises all members of a household using their income to make joint expenditures.³⁴ Each CU is in the panel for 5 consecutive quarters with one interview per quarter. The initial interview only collects demographic characteristics, while the following four interviews collect expenditures from the previous three months. Expenditures are collected for around 600 categories (represented by “universal classification codes”, UCC). Of these 600 UCCs, we select the UCCs that correspond to tradable manufactures. This clearly involves some *ad hoc* decisions. Specifically, we develop two classifications - a conservative classification, where we disregard all uncertain UCCs, and a liberal classification, which includes more UCCs.

Housing related items are particularly difficult, since they are often separated by renter, owned home, and owned vacation home. We disregard these UCCs altogether in the conservative classification. In the liberal classification, we lump the renter and owner categories for the same expense together. For instance, we combine the categories “Installed and non-installed replacement wall to wall carpeting for owned homes” and “Installed and non-installed original wall to wall carpeting for rental homes” into one category, and ignore UCCs that are only

³³The results remain basically unchanged when using major languages instead of official languages.

³⁴Under this definition, a family constitutes a CU, while a boarder living with the family would constitute his own CU as he is financially independent.

available for either renter or owner such as “Installed and non-installed original wall to wall carpeting for owned homes”.

Another difficult class of UCCs is related to clothing - a CU buying “men’s footwear” depends very much whether this CU comprises an adult male. We therefore lump UCCs together across gender and age in the liberal classification, e.g. we collapse “men’s footwear”, “women’s footwear”, “Boys’ footwear”, and “Girls’ footwear” into one category. The conservative classification disregards these categories. In the end, the liberal classification consists of 186 distinct expenditure categories, and the conservative classification comprises 107 categories.

A.6.6 Manufacturing absorption

We use data from the United Nations Industrial Development Organization (UNIDO, 2003) on gross manufacturing output. For the year 2003, this database provides the gross manufacturing output for 74 countries. Unfortunately, the database does not include gross output for several large countries, most notably China. We therefore choose to impute the gross manufacturing output for countries that belonged to the 20 largest economies in 2003 and for which we have no gross manufacturing output. We do this by following Eaton, Kortum, and Kramarz (2004) and scaling value added in the manufacturing sector by the average ratio of gross output and value added across countries. These countries are China, Switzerland, Canada, and Mexico. Finally, for seven countries, the difference between the gross manufacturing output and the total exports, X_{nm} , appears to be negative, so that we exclude these countries from the sample. As a result, the sample consists of 71 countries.

A.6.7 Top quintiles of income distributions

We get data on the top quintiles of the income distributions from UNU-WIDER (2008). A well-known problem of inequality data is that the measure, which the inequality data refers to, varies across countries. In particular, some quintiles refer to expenditures and the others to income. Moreover, income may be measured in gross or net terms. To correct for this, we follow Dollar and Kraay (2002) and regress the observed quintiles on dummies for the underlying measure. We then use the resulting coefficients to estimate the net income quintiles.

Table 1: Dependent variable - extensive margin of bilateral trade

	variable	coefficient
per capita income	exporter	0.66***
	importer	0.47***
population size	exporter	0.65***
	importer	0.30***
bilateral distance	[375, 750)	-0.78***
	[750, 1500)	-1.49***
	[1500, 3000)	-2.26***
	[3000, 6000)	-2.50***
	[6000, ∞)	-2.88***
additional controls	shared border	0.45***
	same language	0.74***

$N = 16053$, $R^2 = 0.63$, *** implies significance at the 1%-level

Table 2: Dependent variable - extensive margin of consumption

specification	I		II		III		IV	
	cons.	liberal	cons.	liberal	cons.	liberal	cons.	liberal
expenditures	0.48***	0.52***	0.43***	0.48***	0.42***	0.46***	0.41***	0.46***
age dummies	no	no	yes	yes	yes	yes	yes	yes
HH members	no	no	no	no	yes	yes	yes	yes
geography	no	no	no	no	no	no	yes	yes
N	2981	2981	2981	2981	2846	2846	2846	2846
R ²	0.33	0.39	0.35	0.42	0.36	0.43	0.38	0.44

Table 3: Estimated trade costs

estimated coefficients		
variable	coefficient	%-effect
[375, 750)	-0.53***	13%
[750, 1500)	-1.52***	40%
[1500, 3000)	-1.97***	55%
[3000, 6000)	-2.84***	88%
[6000, ∞)	-3.33***	110%
shared border	0.77***	-16%
same language	0.82***	-17%

estimated trade costs (d_{ni})		
countries	mean	median
OECD	2.01	1.89
non-OECD	4.27	3.53
all countries	3.91	3.14

Table 4: : Income elasticities of the extensive margin of bilateral trade

	data	elasticities in models	
		non-homothetic	homothetic
exporter income	0.66	0.86	0.86
importer income	0.47	0.63	-0.12

Table 5: : Summary statistics for changes in trade costs

	10%-reduction		25%-reduction		50%-reduction	
	new	EK	new	EK	new	EK
mean	9%	2%	27%	5%	103%	37%
median	8%	0%	22%	-1%	85%	18%
top10%	22%	16%	65%	43%	207%	118%
bottom10%	-2%	-11%	-3%	-25%	16%	-31%
% negative	15%	49%	15%	51%	6%	35%

Figure 1: Demand function




Figure 2: Extensive margin of consumption




Figure 3: Calibrated technologies vs. observed incomes




Figure 4: Calibrated extensive margins of consumption vs. observed incomes




Figure 5: Absolute and relative changes in the extensive margin of consumption

