CGEB WORKING PAPER SERIES

Center for Global Economy and Business
NYU Stern School of Business

THE ELASTICITY OF TRADE: ESTIMATES AND EVIDENCE

Ina Simonovska
Princeton University, University of California - Davis, and NBER

Michael E. Waugh
New York University

Working Paper No. 04
March 2012

Center for Global Economy and Business
NYU Stern School of Business
44 West 4th Street, Suite 7-190
New York, NY 10012
http://www.stern.nyu.edu/cgeb

The Elasticity of Trade: Estimates and Evidence

Ina Simonovska
Princeton University, University of California - Davis, and NBER

Michael E. Waugh
New York University

First Version: April 2009
This Version: November 2011

ABSTRACT

Quantitative results from a large class of structural gravity models of international trade depend critically on the elasticity of trade with respect to trade frictions. We develop a new simulated method of moments estimator to estimate this elasticity from disaggregate price and trade-flow data and we use it within Eaton and Kortum's (2002) Ricardian model. We apply our estimator to new disaggregate price and trade-flow data for 123 countries in the year 2004. Our method yields a trade elasticity of roughly four, nearly fifty percent lower than Eaton and Kortum's (2002) approach. Moreover, robustness exercises result in trade elasticity estimates that are both lower and fall within a narrower range relative to the existing literature. This difference doubles the welfare gains from international trade.

JEL Classification: F10, F11, F14, F17
Keywords: elasticity of trade, bilateral, gravity, price dispersion, indirect inference

Email: inasimonovska@ucdavis.edu, mwaugh@stern.nyu.edu.
We are grateful to the World Bank for providing us with the price data from the 2005 ICP round. We thank George Alessandria, Alexander Aue, Robert Feenstra, Timothy Kehoe, Matthias Lux, B. Ravikumar, seminar participants at CUHK, Tsinghua, UC San Diego, Syracuse, ETH/KOF, Princeton, Uppsala, Oslo, San Francisco Fed, UC Berkeley, NYU and participants at the CESifo Area Conference on Global Economy 2011, NBER ITI Program Meeting Winter 2010, 2010 International Trade Workshop at the Philadelphia Fed, Conference on Microeconomic Sources of Real Exchange Rate Behavior at Vanderbilt, Midwest International Trade Meeting Fall 2010, SED 2010, WEAI 2010, Conference on Trade Costs and International Trade Integration: Past, Present and Future in Venice, International Comparisons Conference at Oxford, North American Winter Meeting of the Econometric Society 2010, AEA 2010 for their feedback.

1. Introduction

Quantitative results from a large class of structural gravity models of international trade depend critically on the elasticity of trade with respect to trade frictions.[1] To illustrate how important this parameter is, consider three examples: Anderson and van Wincoop (2003) find that the estimate of the tariff equivalent of the U.S.-Canada border varies between 48 and 19 percent, depending on the assumed elasticity of trade with respect to trade frictions. Yi (2003) points out that observed reductions in tariffs can explain almost all or none of the growth in world trade, depending on this elasticity. Arkolakis, Costinot, and Rodriguez-Clare (2011) argue that the trade elasticity is one of only two statistics needed to measure the welfare cost of autarky in a large and important class of trade models. Therefore, this elasticity is key to understanding the size of the frictions to trade, the response of trade to changes in tariffs, and the welfare gains or losses from trade.

Estimating this parameter is difficult because quantitative trade models can rationalize small trade flows with either large trade frictions and small elasticities, or small trade frictions and large elasticities. Thus, one needs satisfactory measures of trade frictions *independent* of trade flows to estimate this elasticity. Using their Ricardian model of trade, Eaton and Kortum (2002) (henceforth EK) provide an innovative and simple solution to this problem by arguing that, with product-level price data, one could use the maximum price difference across goods between countries as a proxy for bilateral trade frictions. The maximum price difference between two countries is meaningful because it is bounded by the trade friction between the two countries via simple no-arbitrage arguments.

We develop a new simulated method of moments estimator for the elasticity of trade incorporating EK's intuition. Our argument for a new estimator is that EK's method understates the true trade friction and results in estimates of the trade elasticity that are biased upward by economically significant magnitudes. Thus, we propose a new methodology, which is subject to the same data requirements as EK's approach, and we use it within the Ricardian model in order to correct the bias and arrive at a new estimate for the elasticity of trade.

We apply our estimator to disaggregate price and trade-flow data for the year 2004, which spans 123 countries that account for 98 percent of world output. Our benchmark estimate for the elasticity of trade is 4.12, rather than approximately eight, as EK's estimation strategy suggests. This difference doubles the measured welfare gains from trade across various models.

Since the elasticity of trade plays a key role in quantifying the welfare gains from trade, it is important to understand why our estimates of the parameter differ substantially from EK's.

[1] The class of models includes Krugman (1980), Anderson and van Wincoop (2003), Eaton and Kortum (2002), and Melitz (2003) as articulated in Chaney (2008), which all generate log-linear relationships between bilateral trade flows and trade frictions.

We show that the reason behind the difference is that their estimator is biased in finite samples of price data. The bias arises because the model's equilibrium no-arbitrage conditions imply that the maximum operator over a finite sample of prices underestimates the trade cost with positive probability and overestimates the trade cost with zero probability. Consequently, the maximum price difference lies strictly below the true trade cost, in expectation. This implies that EK's estimator delivers an elasticity of trade that lies strictly above the true parameter, in expectation. As the sample size grows to infinity, EK's estimator can uncover the true elasticity of trade, which necessarily implies that the bias in the estimates of the parameter is eliminated.

Quantitatively, the bias is substantial. To illustrate its severity, we discretize EK's model, simulate trade flows and product-level prices under an assumed elasticity of trade, and apply their estimating approach on artificial data. Assuming a trade elasticity of 8.28—EK's preferred estimate for 19 OECD countries in 1990—EK's procedure yields an elasticity estimate of 12.5, which is nearly 50-percent higher than originally postulated. Moreover, in practice, the true parameter can be recovered when 50,000 goods are sampled across the 19 economies, which constitutes an extreme data requirement to produce unbiased estimates of the elasticity of trade.

Based on these arguments, we propose an estimator that is applicable when the sample size of prices is small. Our approach builds on our insight that one can use observed bilateral trade flows to recover all sufficient parameters to simulate EK's model and to obtain trade flows and prices as functions of the parameter of interest. This insight then suggests a simulated method of moments estimator that minimizes the distance between the moments obtained by applying EK's approach on real and artificial data. We explore the properties of this estimator numerically using simulated data and we show that it can uncover the true elasticity of trade.

Applying our estimator to alternative data sets and conducting several robustness exercises allows us to establish a range for the elasticity of trade between 2.47 and 5.51. In contrast, EK's approach would have found a range of 4.17 to 9.6. Thus, our method finds elasticities that are roughly half the size of EK's approach. Because the inverse of this elasticity linearly controls changes in real income necessary to compensate a representative consumer for going to autarky, our estimates double the measured welfare gains from trade relative to previous findings.

The contribution of this paper is twofold. First, we develop a trade-elasticity estimation approach that is applicable to various existing models of international trade. The methodology and the moments that we use to estimate the trade elasticity within the Ricardian framework can be derived for other trade models with micro-level heterogeneity. In Simonovska and Waugh (2011), we illustrate this point by showing how the estimation strategy introduced in this paper applies to models with variable mark-ups, such as Bernard, Eaton, Jensen, and Kortum (2003), and models that build on the monopolistic-competition structure of Melitz (2003), as articulated in Chaney (2008). Thus, while we focus on the particulars of EK's Ricardian model and our method's relationship with EK's approach, our methodology contributes to the estimation of trade elasticities above and beyond a particular model.

Second, our findings suggest both a lower and narrower range for the trade elasticity relative to the existing literature. In particular, EK establish a range for the elasticity estimate between 3.6 and 12.8, with their benchmark strategy yielding values in the middle of this range. Our critique applies to the strategy that EK follow in order to obtain the estimate of 12.8 as well. When we apply our estimator to EK's alternative approach, we obtain an estimate of 4.4, which is nearly the same as our benchmark finding. Thus, we provide a lower and narrower range of 2.47 to 5.51, relative to EK's estimates.

In addition, the range of estimates that we provide is narrower and lower relative to the range that Anderson and van Wincoop (2004) establish. The authors survey the literature that estimates the trade elasticity using various approaches and they establish a range between five and ten. One set of estimates that Anderson and van Wincoop (2004) document is obtained using data on changes in bilateral trade flows and tariffs during trade liberalization episodes, as in Head and Ries (2001), Romalis (2007), and Caliendo and Parro (2011). The approach typically associates the entire change in trade flows during a trade liberalization with changes in tariffs. This necessarily results in high estimates of the trade elasticity, since changes in non-tariff barriers that occur during trade liberalizations are not accounted for in the estimation. Moreover, this approach is subject to a large data requirement, so it typically focuses on a particular episode that involves a handful of countries. In contrast, using our methodology, we provide estimates for the trade elasticity from data that spans as many as 123 countries, which account for 98 percent of world output.

Another set of estimates that Anderson and van Wincoop (2004) report is obtained using Feenstra's (1994) methodology. However, in heterogeneous frameworks that rely on constant-elasticity-of-substitution (CES) preferences, such as EK's Ricardian model, Feenstra's (1994) method cannot recover the elasticity of trade. To demonstrate this fact, we apply Feenstra's (1994) method to data generated from the Ricardian model. We find that the method recovers the preference parameter that controls the elasticity of substitution across goods. This parameter plays no role in determining aggregate trade flows and welfare gains from trade in the Ricardian model.

To summarize, in this paper, we develop a new methodology to estimate the elasticity of trade that is applicable across gravity-based models of international trade that feature micro-level heterogeneity. We apply our estimator to novel disaggregate price and trade-flow data for the year 2004, which spans 123 countries that account for 98 percent of world output. Across numerous exercises, we obtain estimates of the trade elasticity that are both lower and fall within a narrower range relative to the existing literature. Our findings imply that the measured welfare gains from international trade are twice as high as previously documented.

2. Model

We outline the environment of the multi-country Ricardian model of trade introduced by EK. We consider a world with N countries, where each country has a tradable final-goods sector. There is a continuum of tradable goods indexed by $j \in [0, 1]$.

Within each country i, there is a measure of consumers L_i. Each consumer has one unit of time supplied inelastically in the domestic labor market and enjoys the consumption of a CES bundle of final tradable goods with elasticity of substitution $\rho > 1$:

$$U_i = \left[\int_0^1 x_i(j)^{\frac{\rho-1}{\rho}} dj \right]^{\frac{\rho}{\rho-1}}.$$

To produce quantity $x_i(j)$ in country i, a firm employs labor using a linear production function with productivity $z_i(j)$. Country i's productivity is, in turn, the realization of a random variable (drawn independently for each j) from its country-specific Fréchet probability distribution:

$$F_i(z_i) = \exp(-T_i z_i^{-\theta}). \tag{1}$$

The country-specific parameter $T_i > 0$ governs the location of the distribution; higher values of it imply that a high productivity draw for any good j is more likely. The parameter $\theta > 1$ is common across countries and, if higher, it generates less variability in productivity across goods.

Having drawn a particular productivity level, a perfectly competitive firm from country i incurs a marginal cost to produce good j of $w_i/z_i(j)$, where w_i is the wage rate in the economy. Shipping the good to a destination n further requires a per-unit iceberg trade cost of $\tau_{ni} > 1$ for $n \neq i$, with $\tau_{ii} = 1$. We assume that cross-border arbitrage forces effective geographic barriers to obey the triangle inequality: For any three countries i, k, n, $\tau_{ni} \leq \tau_{nk}\tau_{ki}$.

Below, we describe equilibrium prices, trade flows, and welfare.

Perfect competition forces the price of good j from country i to destination n to be equal to the marginal cost of production and delivery:

$$p_{ni}(j) = \frac{\tau_{ni} w_i}{z_i(j)}.$$

So, consumers in destination n would pay $p_{ni}(j)$, should they decide to buy good j from i.

Consumers purchase good j from the low-cost supplier; thus, the actual price consumers in n

pay for good j is the minimum price across all sources k:

$$p_n(j) = \min_{k=1,\ldots,N} \Big\{ p_{nk}(j) \Big\}. \quad (2)$$

The pricing rule and the productivity distribution allow us to obtain the following CES exact price index for each destination n:

$$P_n = \gamma \Phi_n^{-\frac{1}{\theta}} \quad \text{where} \quad \Phi_n = \left[\sum_{k=1}^{N} T_k (\tau_{nk} w_k)^{-\theta} \right]. \quad (3)$$

In the above equation, $\gamma = \left[\Gamma\left(\frac{\theta + 1 - \rho}{\theta} \right) \right]^{\frac{1}{1-\rho}}$ is the Gamma function, and parameters are restricted such that $\theta > \rho - 1$.

To calculate trade flows between countries, let X_n be country n's expenditure on final goods, of which X_{ni} is spent on goods from country i. Since there is a continuum of goods, computing the fraction of income spent on imports from i, X_{ni}/X_n, can be shown to be equivalent to finding the probability that country i is the low-cost supplier to country n given the joint distribution of efficiency levels, prices, and trade costs for any good j. The expression for the share of expenditures that n spends on goods from i or, as we will call it, the trade share is:

$$\frac{X_{ni}}{X_n} = \frac{T_i (\tau_{ni} w_i)^{-\theta}}{\sum_{k=1}^{N} T_k (\tau_{nk} w_k)^{-\theta}}. \quad (4)$$

Expressions (3) and (4) allow us to relate trade shares to trade costs and the price indices of each trading partner via the following equation:

$$\frac{X_{ni}/X_n}{X_{ii}/X_i} = \frac{\Phi_i}{\Phi_n} \tau_{ni}^{-\theta} = \left(\frac{P_i \tau_{ni}}{P_n} \right)^{-\theta}, \quad (5)$$

where $\frac{X_{ii}}{X_i}$ is country i's expenditure share on goods from country i, or its home trade share.

In this model, it is easy to show that the welfare gains from trade are essentially captured by changes in the CES price index that a representative consumer faces. Because of the tight link between prices and trade shares, this model generates the following relationship between changes in price indices and changes in home trade shares, as well as, the elasticity parameter:

$$\frac{P_n'}{P_n} - 1 = 1 - \left(\frac{X_{nn}'/X_n'}{X_{nn}/X_n} \right)^{\frac{1}{\theta}}, \quad (6)$$

where the left-hand side can be interpreted as the percentage compensation a representative consumer in country n requires to move between two trading equilibria.

Expression (5) is not particular to EK's model. Several popular models of international trade relate trade shares, prices and trade costs in the same exact manner. These models include the Armington framework of Anderson and van Wincoop (2003) and the monopolistic competition framework of Krugman (1980). More importantly for the context of this paper, the heterogeneous Ricardian framework of Bernard, Eaton, Jensen, and Kortum (2003) and the model of firm heterogeneity by Melitz (2003), when parametrized as in Chaney (2008), also generate this relationship. Arkolakis, Costinot, and Rodriguez-Clare (2011) show how equation (6) arises in all of these models.

2.1. The Elasticity of Trade

The key parameter determining trade flows (equation (5)) and welfare (equation (6)) is θ. To see the parameter's importance for trade flows, take logs of equation (5) yielding:

$$\log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right) = -\theta\left[\log\left(\tau_{ni}\right) - \log(P_i) + \log(P_n)\right]. \tag{7}$$

As this expression makes clear, θ controls how a change in the bilateral trade costs, τ_{ni}, will change bilateral trade between two countries. This elasticity is important because if one wants to understand how a bilateral trade agreement will impact aggregate trade or to simply understand the magnitude of the trade friction between two countries, then a stand on this elasticity is necessary. This is what we mean by the elasticity of trade.

To see the parameter's importance for welfare, it is easy to demonstrate that (6) implies that θ represents the inverse of the elasticity of welfare with respect to domestic expenditure shares:

$$\log(P_n) = -\frac{1}{\theta}\log\left(\frac{X_{nn}}{X_n}\right) \tag{8}$$

Hence, decreasing the domestic expenditure share by one percent generates a $(1/\theta)/100$-percent increase in consumer welfare. Thus, in order to measure the impact of trade policy on welfare, it is sufficient to obtain data on realized domestic expenditures and an estimate of the elasticity of trade.

Given θ's impact on trade flows and welfare, this elasticity is absolutely critical in any quantitative study of international trade.

3. Estimating θ: EK's Approach

Equation (5) suggests that one could easily estimate θ if one had data on trade shares, aggregate prices, and trade costs. The key issue is that trade costs are not observed. In this section, we discuss how EK approximate trade costs and estimate θ. Then, we characterize the statistical properties of EK's estimator. The key result is Proposition 1, which states that their estimator is biased and overestimates the elasticity of trade with a finite sample of prices. The second result is Proposition 2, which states that EK's estimator is a consistent and an asymptotically unbiased estimator of the elasticity of trade.

3.1. Approximating Trade Costs

The main problem with estimating θ is that one must disentangle θ from trade costs, which are not observed. EK propose approximating trade costs using *disaggregate* price information across countries. In particular, the maximum price difference across goods between two countries bounds the bilateral trade cost, which solves the indeterminacy issue.

To illustrate this argument, suppose that we observe the price of good ℓ across locations, but we do not know its country of origin.[2] We know that the price of good ℓ in country n relative to country i must satisfy the following inequality:

$$\frac{p_n(\ell)}{p_i(\ell)} \leq \tau_{ni}. \tag{9}$$

That is, the relative price of good ℓ must be less than or equal to the trade friction. This inequality must hold because if it does not, then $p_n(\ell) > \tau_{ni} p_i(\ell)$ and an agent could import ℓ at a lower price. Thus, the inequality in (9) places a lower bound on the trade friction.

Improvements on this bound are possible if we observe a sample of L goods across locations. This follows by noting that the *maximum* relative price must satisfy the same inequality:

$$\max_{\ell \in L} \left\{ \frac{p_n(\ell)}{p_i(\ell)} \right\} \leq \tau_{ni}. \tag{10}$$

This suggests a way to exploit *disaggregate* price information across countries and to arrive at an estimate of τ_{ni} by taking the maximum of relative prices over goods. Thus, EK approximate τ_{ni}, in logs, by

$$\log \hat{\tau}_{ni}(L) = \max_{\ell \in L} \left\{ \log\left(p_n(\ell)\right) - \log\left(p_i(\ell)\right) \right\}, \tag{11}$$

where the "hat" denotes the approximated value of τ_{ni} and (L) indexes its dependence on the sample size of prices.

[2]This is the most common case, though Donaldson (2009) exploits a case where he knows the place of origin for one particular good, salt. He argues convincingly that in India, salt was produced in only a few locations and exported everywhere; thus, the relative price of salt across locations identifies the trade friction.

3.2. Estimating the Elasticity

Given the approximation of trade costs, EK derive an econometric model that corresponds to (7). For a sample of L goods, they estimate a parameter, β, using a method of moments estimator, which takes the ratio of the average of the left-hand side of (7) to the average of the term in the square bracket of the right-hand side of (7), where the averages are computed across all country pairs.[3] Mathematically, their estimator is:

$$\hat{\beta} = -\frac{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}{\sum_n \sum_i \left(\log \hat{\tau}_{ni}(L) + \log \hat{P}_i - \log \hat{P}_n\right)}, \tag{12}$$

$$\text{where} \quad \log \hat{\tau}_{ni}(L) = \max_{\ell \in L} \left\{ \log p_n(\ell) - \log p_i(\ell) \right\},$$

$$\text{and} \quad \log \hat{P}_i = \frac{1}{L} \sum_{\ell=1}^{L} \log(p_i(\ell)).$$

The value of β is EK's preferred estimate of the elasticity θ.[4] Throughout, we will denote by $\hat{\beta}$ the estimator defined in equation (12) to distinguish it from the value θ. As discussed, the second line of expression (12) approximates the trade cost. The third line approximates the aggregate price indices. The top line represents a rule that combines these statistics, together with observed trade flows, in an attempt to estimate the elasticity of trade.

3.3. Properties of EK's Estimator

Before describing the properties of the estimator $\hat{\beta}$, we want to be clear about the sources of error in equation (12). Bilateral trade flows are observable statistics. Trade barriers and price indices are approximated from price data using the last two equations in (12). Hence, they are potentially measured with error because of the approximation. In the model, prices are realizations of random variables, thus we treat the micro-level prices as being randomly sampled from the equilibrium distribution of prices. This allows us to theoretically characterize the properties of the approximation error and in turn to derive the properties of the estimator $\hat{\beta}$ in (12).[5]

Given our assumption that the prices are randomly sampled from the equilibrium distribution, we define the following objects.

[3] They also propose two other estimators. One uses the approximation in (11) and the gravity equation in (21). We show in Appendix C that our arguments are applicable to this approach as well. The other approach does not use disaggregate price data and we discuss it later.

[4] To alleviate measurement error, EK use the second-order statistic over price differences rather than the first-order statistic. Our estimation approach is robust to either specification.

[5] In practice, there may be other sources of measurement error in the data which are outside of the model. We discuss these issues in Section 7.2.

Definition 1 *Define the following objects:*

1. *Let $\epsilon_{ni} = \theta[\log p_n - \log p_i]$ be the log price difference of a good between country n and country i, multiplied by θ.*

2. *Let the vector $\boldsymbol{S} = \{\log(T_1 w_1^{-\theta}), \ldots, \log(T_N w_N^{-\theta})\}$.*

3. *Let the vector $\tilde{\boldsymbol{\tau}}_{\boldsymbol{i}} = \{\theta \log(\tau_{i1}), \ldots, \theta \log(\tau_{iN})\}$ and let $\tilde{\boldsymbol{\tau}}$ be a matrix with typical row, $\tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}$.*

4. *Let $g(p_i; \boldsymbol{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}})$ be the pdf of prices of individual goods in country i, $p_i \in (0, \infty)$.*

5. *Let $f_{\max}(\epsilon_{ni}; L, \boldsymbol{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{n}})$ be the pdf of $\max(\epsilon_{ni})$, given prices of a sample $L \geq 1$ of goods.*

6. *Let $\mathbb{X}$ denote the normalized trade share matrix, with typical (n, i) element, $\log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)$.*

The first item is simply the scaled log price difference. As we show in Appendix 2.1, this happens to be convenient to work with, as the second line in (12) can be restated in terms of scaled log price differences across locations. The second item is a vector in which each element is a function of a country's technology parameter and wage rate. The third item is a matrix of log bilateral trade costs, scaled by θ, with a typical vector row containing the trade costs that country i's trading partners incur to sell there. The fourth item specifies the probability distribution of prices in each country. The fifth item specifies the probability distribution over the maximum scaled log price difference and its dependence on the sample size of prices of L goods. We derive this distribution in Appendix 2.1. Finally, the sixth item summarizes trade data, which we view as observable statistics.

3.4. $\hat{\beta}$ is a Biased Estimator of θ

Given these definitions, we establish two intermediate results and then state Proposition 1, which characterizes the expectation of $\hat{\beta}$, shows that the estimator is biased and discusses the reason why the bias arises. The proof of Proposition 1 can be found in Appendix 2.1.

The first intermediate result is the following:

Lemma 1 *Consider an economy of N countries with a sample of L goods' prices observed. The expected value of the maximal difference of logged prices for a pair of countries is strictly less than the true trade cost,*

$$\Psi_{ni}(L; \boldsymbol{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{n}}) \equiv \frac{1}{\theta} \int_{-\theta \log(\tau_{in})}^{\theta \log(\tau_{ni})} \epsilon_{ni} f_{\max}(\epsilon_{ni}; L, \boldsymbol{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{n}}) d\epsilon_{ni} \; < \; \log(\tau_{ni}). \tag{13}$$

The difference in the expected values of logged prices for a pair of countries equals the difference in the price parameters, Φ, of the two countries,

$$\Omega_{ni}(\boldsymbol{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{n}}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}) \equiv \int_0^{\infty} \log(p_n) g(p_n; \boldsymbol{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{n}}) dp_n - \int_0^{\infty} \log(p_i) g(p_i; \boldsymbol{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}) dp_i = \frac{1}{\theta} \left(\log \Phi_i - \log \Phi_n \right), \quad (14)$$

with Φ_n defined in equation (3).

The key result in Lemma 1 is the strict inequality in (13). It says that Ψ_{ni}, the expected maximal log price difference, is less than the true log trade cost. Two forces drive this result. First, with a finite sample L of prices, there is positive probability that the maximal log price difference will be less than the true log trade cost. In other words, there is always a chance that the weak inequality in (11) does not bind. Second, there is zero probability that the maximal log price difference can be larger than the true log trade cost. This comes from optimality and the definition of equilibrium. These two forces imply that the expected maximal log price difference lies strictly below the true log trade cost.

The second result in Lemma 1 is that the difference in the expected log prices in expression (14) equals the difference in the aggregate price parameters defined in equation (3). This result is important because it implies that any source of bias in the estimator $\hat{\beta}$ does not arise because of systematic errors in approximating the price parameter Φ.

The next intermediate step computes the expected value of $1/\hat{\beta}$. This step is convenient because the inverse of $\hat{\beta}$ is linear in the random variables that Lemma 1 characterizes.

Lemma 2 *Consider an economy of N countries with a sample of L goods' prices observed. The expected value of $1/\hat{\beta}$ equals:*

$$E\left(\frac{1}{\hat{\beta}}\right) = \frac{1}{\theta} \left\{ -\frac{\sum_n \sum_i (\theta \Psi_{ni}(L) - (\log \Phi_i - \log \Phi_n))}{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)} \right\} < \frac{1}{\theta}, \quad (15)$$

with

$$1 > \left\{ -\frac{\sum_n \sum_i (\theta \Psi_{ni}(L) - (\log \Phi_i - \log \Phi_n))}{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)} \right\} > 0. \quad (16)$$

This results says that the expected value of the inverse of $\hat{\beta}$ equals the inverse of the elasticity multiplied by the bracketed term of (16). The bracketed term is the expected maximal log price difference minus the difference in expected log prices, both scaled by theta, and divided by trade data. This term is strictly less than one because Ψ_{ni} does not equal the log trade cost, as established in Lemma 1. If Ψ_{ni} did equal the log trade cost, then the bracketed term would equal one, and the expected value of the inverse of $\hat{\beta}$ would be equal to the inverse of θ. This can be seen by examining the relation between Φ's and aggregate prices P's in (3), and by substituting expression (7) into (16).

Inverting (15) and then applying Jensen's inequality establishes the main result: EK's estimator is biased above the true value of θ.

Proposition 1 *Consider an economy of N countries with a sample of L goods' prices observed. The expected value of $\hat{\beta}$ is*

$$E\left(\hat{\beta}\right) \;\geq\; \theta \times \left\{ -\frac{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}{\sum_n \sum_i (\theta \Psi_{ni}(L) - (\log \Phi_i - \log \Phi_n))} \right\} \;>\; \theta. \tag{17}$$

The proposition establishes that the estimator $\hat{\beta}$ provides estimates that exceed the true value of the elasticity θ. The weak inequality in (17) comes from applying Jensen's inequality to the strictly convex function of $\hat{\beta}$, $1/\hat{\beta}$. The strict inequality follows from Lemma 1, which argued that the expected maximal logged price difference is strictly less than the true trade cost. Thus, the bracketed term in expression (17) is always greater than one and the elasticity of trade is always overestimated.

3.5. Consistency and Asymptotic Bias

While the estimator $\hat{\beta}$ is biased in a finite sample, the asymptotic properties of EK's estimator are worth understanding. Proposition 2 summarizes the result. The proof to Proposition 2 can be found in Appendix 2.2.

Proposition 2 *Consider an economy of N countries. The maximal log price difference is a consistent estimator of the trade cost,*

$$\underset{L\to\infty}{\text{plim}} \max_{\ell=1,\ldots,L} (\log p_n(\ell) - \log p_i(\ell)) = \log \tau_{ni}. \tag{18}$$

The estimator $\hat{\beta}$ is a consistent estimator of θ,

$$\underset{L\to\infty}{\text{plim}} \; \hat{\beta}\left(L; \boldsymbol{S}, \tilde{\boldsymbol{\tau}}, \mathbb{X}\right) = \theta, \tag{19}$$

and the asymptotic bias of $\hat{\beta}$ is zero,

$$\lim_{L\to\infty} E\left[\hat{\beta}\left(L; \boldsymbol{S}, \tilde{\boldsymbol{\tau}}, \mathbb{X}\right)\right] - \theta = 0. \tag{20}$$

There are three elements to Proposition 2, each building on the previous one. The first statement says that the probability limit of the maximal log price difference equals the true log trade cost between two countries. Intuitively, this says that as the sample size becomes large, the probability that the weak inequality in (10) does not bind becomes vanishingly small.

The second statement says that the estimator $\hat{\beta}$ converges in probability to the elasticity of trade—i.e., $\hat{\beta}$ is a consistent estimator of θ. The reasons are the following. Because the maximal log price difference converges in probability to the true log trade cost, and the difference in averages of log prices converges in probability to the difference in log price parameters, $1/\hat{\beta}$ converges in probability to $1/\theta$. Since $1/\hat{\beta}$ is a continuous function of $\hat{\beta}$ (with $\hat{\beta} > 0$), $\hat{\beta}$ must converge in probability to θ because of the preservation of convergence for continuous functions (see Hayashi (2000)).

The third statement says that, in the limit, the bias is eliminated. This follows immediately from the argument that $\hat{\beta}$ is a consistent estimator of θ (see Hayashi (2000)).

The results in Proposition 2 are important for two reasons. First, they suggest that with enough data, EK's estimator provides informative estimates of the elasticity of trade. However, as we will show in the next section, Monte Carlo exercises suggest that the data requirements are extreme. Second, because EK's estimator has desirable asymptotic properties, it underlies the simulation-based estimator that we develop in Section 5.

4. How Large is the Bias? How Much Data is Needed?

Proposition 1 shows that EK's estimator is biased in a finite sample. Many estimators have this property, which raises the question: How large is the bias? Furthermore, even if the magnitude of the bias is large, perhaps moderate increases in the sample size are sufficient to eliminate the bias (in practical terms). The natural question is: How much data are needed to achieve that?

To answer these questions, we perform Monte Carlo experiments in which we simulate trade flows and samples of micro-level prices under a known θ. Then, we apply EK's estimator (and other estimators) to the artificial data. To simulate trade flows that mimic the data, we use the simulation procedure that is described in Steps 1-3 in Section 5.2 below. We estimate all the parameters necessary to simulate the model (except for θ) using the trade data from EK. We set the true value of θ equal to 8.28, which is EK's estimate when employing the approach described above. We then randomly sample prices from the simulated data and we apply EK's estimation to the simulated trade flows and prices. The sample size of prices is set to $L = 50$, which is the number of prices EK had access to in their data set.

Table 1 summarizes our findings. The columns of Table 1 present the mean and median estimates of β over 100 simulations. The rows present two different estimation approaches: method

Table 1: Monte Carlo Results, True θ = 8.28

Approach	Mean Estimate of θ (S.E.)	Median Estimate of θ
EK's Estimator	12.5 (0.06)	12.5
Least Squares	11.8 (0.06)	11.8
True Mean τ = 1.79	Estimated Mean τ = 1.48	

Note: S.E. is the standard error of the mean. In each simulation there are 19 countries and 500,000 goods. Only 50 realized prices are randomly sampled and used to estimate θ. 100 simulations performed.

of moments and least squares with suppressed constant. Also reported are the true average trade cost and the estimated average trade cost using maximal log price differences.

The first row in Table 1 shows that the estimates using EK's approach are larger than the true θ of 8.28, which is consistent with Proposition 1. The key source of bias in Proposition 1 was that the estimates of the trade costs were biased downward, as Lemma 1 argued. The final row in Table 1 illustrates that the estimated trade costs are below the true trade costs, where the latter correspond to an economy characterized by a true elasticity of trade among 19 OECD countries of 8.28.

The second row in Table 1 reports results using a least squares estimator with the constant suppressed rather than the method of moments estimator.[6] Similar to the method of moments estimates, the least squares estimates are substantially larger than the true value of θ. This is important because it suggests that the key problem with EK's approach is not the method of moment estimator per se, but, instead, the poor approximation of the trade costs.

The final point to note is that the magnitude of the bias is substantial. The underlying θ was set equal to 8.28, and the estimates in the simulation are between 11.8 and 12.5. Equation (8) can be used to formulate the welfare cost of the bias. It suggests that the welfare gains from trade will be underestimated by 50 percent as a result of the bias.

While Table 1 reflects the results from a particular calibration of the model to trade flow data, one would like to know how these results depend on the particulars of the economy like trade costs. Inspection of (15) and the integral in (13) shows that the bias will depend on trade flows and the level of trade costs in the economy. For example, as all trade costs approach one, the

[6] We have found that including a constant in least squares results in slope coefficients that either underestimate or overestimate the elasticity depending on the level of trade costs in the simulation. Hence, including a constant term does not resolve the bias.

bias will disappear holding fixed the sample size of prices. The reason is that as trade costs approach one, all goods become traded and hence the maximal price difference—even in a small sample—will likely reflect the true trade friction.

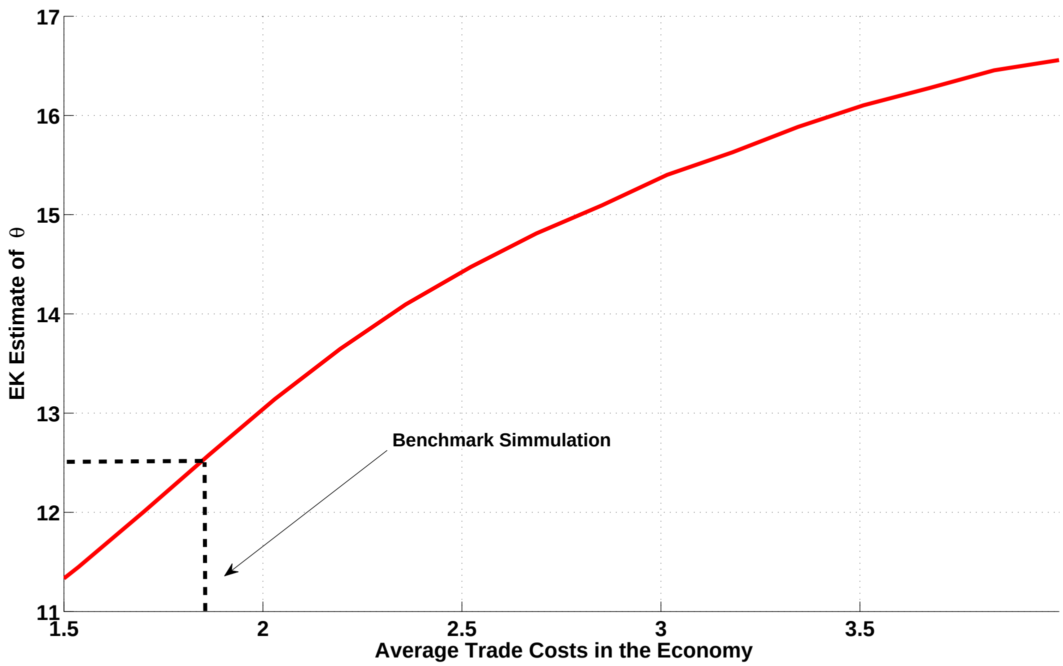

Figure 1: EK's Estimator and the Level of Trade Costs, True $\theta = 8.28$

Figure 1 shows how the bias behaves when trade costs are increasing away from one and the economy approaches autarky. To generate this figure, we keep the true θ equal to 8.28 and we uniformly scale the trade costs from the baseline simulation up or down. We then apply EK's estimation approach to the simulated data (now indexed by the level of trade costs) with the sample size of prices set equal to 50. The x-axis reports the average trade cost across all the countries and the y-axis reports the associated estimate of θ.

Figure 1 shows that, as trade costs increase, EK's estimate of θ increases and hence the bias increases. For example, when the average trade cost equals about three, EK's estimate of θ is 16—almost two times larger than the true θ of 8.28. In contrast, in the baseline simulation when average trade costs are about 1.8, EK's estimate is only fifty percent larger at 12.5. The intuition for this outcome is straightforward. As trade costs increase, more goods are likely to become non-traded and hence it is more likely that many of the prices in the sample are not informative about trade costs.

How much data is needed to eliminate the bias? Table 2 provides a quantitative answer. It performs the same Monte Carlo experiments described above, as the sample size of micro-level prices varies.

Table 2: Increasing the Sample of Prices Reduces the Bias, True θ = 8.28

Sample Size of Prices	Mean θ (S.E.)	Median θ	Mean τ
50	12.51 (0.06)	12.50	1.48
500	9.34 (0.02)	9.32	1.68
5,000	8.43 (0.01)	8.43	1.77
50,000	8.30 (0.002)	8.30	1.78

Note: S.E. is the standard error of the mean. In each simulation, there are 19 countries and 500,000 goods. The results reported use least squares with the constant suppressed. 100 simulations performed. True Mean τ = 1.79.

Table 2 shows that, as the sample size becomes larger, the estimate of θ becomes less biased and begins to approach the true value of θ. The final column shows how the reduction in the bias coincides with the estimates of the trade costs becoming less biased. This is consistent with the arguments of Proposition 2, which describes the asymptotic properties of this estimator.

We should note that the rate of convergence is extremely slow; even with a sample size of 5,000, the estimate of β is meaningfully larger than the value generating the data. Only when 50,000 prices are sampled does the estimate approach the true value. The exercise allows us to conclude that the data requirements to minimize the bias in estimates of the elasticity of trade (in practice) are extreme. This motivates our alternative estimation strategy in the next section.

5. A New Approach To Estimating θ

In this section, we develop a new approach to estimating θ and we discuss its performance on simulated data.

5.1. The Idea

Our idea is to exploit the structure of the model as follows. First, in Section 5.2, we show how to recover all the parameters that are needed to simulate the model up to the unknown scalar θ from trade data only. These parameters are the vector $\mathbf{S}$ and the scaled trade costs in matrix $\tilde{\boldsymbol{\tau}}$. Given these values, we can simulate moments from the model as functions of θ.

Second, Lemma 1 and Lemma 2 actually suggest which moments are informative. Inspection of the integral (13) and the density f_{max} in (b.28) leads to the observation that the expected maximal log price difference monotonically varies with θ and linearly with $1/\theta$. This follows because of the previous point—the vector $\mathbf{S}$ and scaled log trade costs $\tilde{\boldsymbol{\tau}}$ are pinned down by trade data, and these values completely determine all parameters in the integral (13), except the value $1/\theta$ lying outside the integral. Similarly, the integral (14) is completely determined by these values and scaled in the same way by $1/\theta$ as (13) is.

These observations have the following implication. While the maximum log price difference is biased below the true trade cost, if θ is large, then the value of the maximum log price difference will be small. Similarly, if θ is small, then the value of the maximum log price difference will be large. A large or small maximum log price difference will result in a small or large estimate of β. This suggests that the estimator $\hat{\beta}$ will vary monotonically with the true value of θ. Furthermore, this suggests that β is an informative moment with regard to θ.[7]

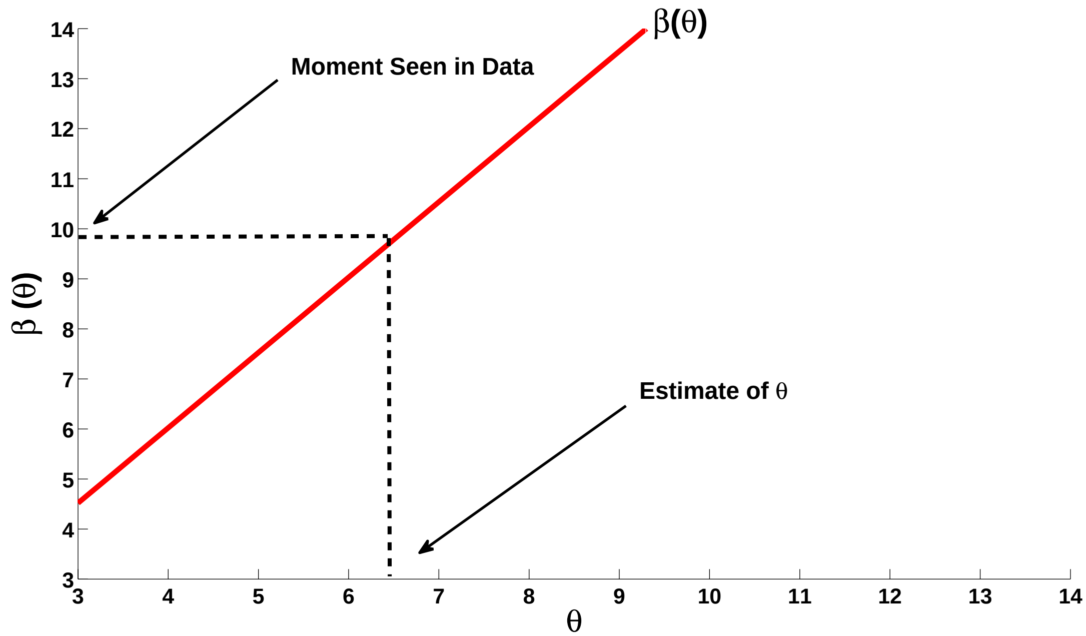

Figure 2: Schematic of Estimation Approach

Figure 2 quantitatively illustrates this intuition by plotting $\beta(\theta)$ from simulations as we varied θ. It is clear that β is a biased estimator because these values do not lie on the 45^o line. However, β varies near linearly with θ. These observations suggest an estimation procedure that matches the data moment β to the moment $\beta(\theta)$ implied by the simulated model under a known θ.[8] Because of the monotonicity implied by our arguments, the known θ must be the unique value that satisfies the moment condition specified.

5.2. Simulation Approach

In this subsection, we show how to recover all parameters of interest up to the unknown scalar θ from trade data only, and then we describe our simulation approach. This provides the foundation for the simulated method of moments estimator that we propose.

[7]Lemma 1 established that the expected value of $1/\hat{\beta}$ is proportional to $1/\theta$. Hence, modulo effects from Jensen's inequality, this suggests that $\hat{\beta}$ is roughly proportional to θ. Figure 2 confirms this.

[8]Another reason for using the moment β is that $\hat{\beta}$ is a consistent estimator of θ, as argued in Proposition 2.

Step 1.—We estimate the parameters for the country-specific productivity distributions and trade costs from bilateral trade-flow data. We follow closely the methodologies proposed by EK and Waugh (2010b). First, we derive the gravity equation from expression (4) by dividing the bilateral trade share by the importing country's home trade share,

$$\log\left(\frac{X_{ni}/X_n}{X_{nn}/X_n}\right) = S_i - S_n - \theta \log \tau_{ni}, \tag{21}$$

where S_i is defined as $\log\left[T_i w_i^{-\theta}\right]$ and is the same value in the parameter vector **S** in Definition 1. Note that (21) is a different equation than expression (5), which is derived by dividing the bilateral trade share by the exporting country's home trade share, and is used to estimate θ. S_i's are recovered as the coefficients on country-specific dummy variables given the restrictions on how trade costs can covary across countries. Following the arguments of Waugh (2010b), trade costs take the following functional form:

$$\log(\tau_{ni}) = d_k + b_{ni} + ex_i + \nu_{ni}. \tag{22}$$

Here, trade costs are a logarithmic function of distance, where d_k with $k = 1, 2, ..., 6$ is the effect of distance between country i and n lying in the k-th distance interval.[9] b_{ni} is the effect of a shared border in which $b_{ni} = 1$ if country i and n share a border and zero otherwise. The term ex_i is an exporter fixed effect and allows for the trade-cost level to vary depending upon the exporter. We assume that ν_{ni} reflects other factors and is orthogonal to the regressors and normally distributed with mean zero and standard deviation σ_ν. We use least squares to estimate equations (21) and (22).

Step 2.—The parameter estimates obtained from the first-stage gravity regression are sufficient to simulate trade flows and micro-level prices up to a constant, θ.

The relationship is obvious in the estimation of trade barriers since $\log(\tau_{ni})$ is scaled by θ in (21). To see that we can simulate micro-level prices as a function of θ only, notice that for any good j, $p_{ni}(j) = \tau_{ni} w_i / z_i(j)$. Thus, rather than simulating productivities, it is sufficient to simulate the inverse of marginal costs of production $u_i(j) = z_i(j)/w_i$. In Appendix 2.3, we show that u_i is distributed according to:

$$M_i(u_i) = \exp\left(-\tilde{S}_i u_i^{-\theta}\right), \quad \text{with} \quad \tilde{S}_i = \exp(S_i) = T_i w_i^{-\theta}. \tag{23}$$

Thus, having obtained the coefficients S_i from the first-stage gravity regression, we can simulate the inverse of marginal costs and prices.

[9] Intervals are in miles: $[0, 375)$; $[375, 750)$; $[750, 1500)$; $[1500, 3000)$; $[3000, 6000)$; and $[6000, \text{maximum}]$. An alternative to specifying a trade-cost function is to recover scaled trade costs as a residual using equation (5), trade data, and measures of aggregate prices as in Waugh (2010a).

To simulate the model, we assume that there are a large number (150,000) of potentially tradable goods. In Section 7.1, we discuss how we made this choice and the motivation behind it. For each country, the inverse marginal costs are drawn from the country-specific distribution (23) and assigned to each good. Then, for each importing country and each good, the low-cost supplier across countries is found, realized prices are recorded, and aggregate bilateral trade shares are computed.

Step 3.—From the realized prices, a subset of goods common to all countries is defined and the subsample of prices is recorded – i.e., we are acting as if we were collecting prices for the international organization that collects the data. We added disturbances to the predicted trade shares with the disturbances drawn from a mean zero normal distribution with the standard deviation set equal to the standard deviation of the residuals from Step 1.

These steps then provide us with an artificial data set of micro-level prices and trade shares that mimic their analogs in the data. Given this artificial data set, we can then compute moments—as functions of θ—and compare them to the moments in the data.

5.3. Estimation

We perform two estimations: an overidentified procedure with two moments and an exactly identified procedure with one moment. Below, we describe the moments we try to match and the details of our estimation procedure.

Moments. Define $\hat{\beta}_k$ as EK's method of moment estimator defined in (12) using the kth-order statistic over micro-level price differences. Then, the moments we are interested in are:

$$\beta_k = -\frac{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}{\sum_n \sum_i \left(\log \hat{\tau}_{ni}^k(L) + \log \hat{P}_i - \log \hat{P}_n\right)}, \quad k = 1, 2 \tag{24}$$

where $\hat{\tau}_{ni}^k(L)$ is computed as the kth-order statistic over L micro-level price differences between countries n and i. In the exactly identified estimation, we use β_1 as the only moment.

We denote the simulated moments by $\beta_1(\theta, u_s)$ and $\beta_2(\theta, u_s)$, which come from the analogous formula as in (24) and are estimated from artificial data generated by following **Steps 1-3** above. Note that these moments are a function of θ and depend upon a vector of random variables u_s associated with a particular simulation s. There are three components to this vector. First, there are the random productivity draws for production technologies for each good and each country. The second component is the set of goods sampled from all countries. The third component mimics the residuals ν_{ni} from equation (21), which are described in Section 5.2.

Stacking our data moments and averaged simulation moments gives us the following zero function:

$$y(\theta) = \begin{bmatrix} \beta_1 - \frac{1}{S}\sum_{s=1}^{S} \beta_1(\theta, u_s) \\ \\ \beta_2 - \frac{1}{S}\sum_{s=1}^{S} \beta_2(\theta, u_s) \end{bmatrix}. \tag{25}$$

Estimation Procedure. We base our estimation procedure on the moment condition:

$$E\left[y(\theta_o)\right] = 0,$$

where θ_o is the true value of θ. Thus, our simulated method of moments estimator is:

$$\hat{\theta} = \arg\min_{\theta}\left[y(\theta)' \, \mathbf{W} \, y(\theta)\right], \tag{26}$$

where **W** is a 2×2 weighting matrix that we discuss below.

The idea behind this moment condition is that, though β_1 and β_2 will be biased away from θ, the moments $\beta_1(\theta, u_s)$ and $\beta_2(\theta, u_s)$ will be biased by the same amount when evaluated at θ_o, in expectation. Viewed in this language, our moment condition is closely related to the estimation of bias functions discussed in MacKinnon and Smith (1998) and to indirect inference, as discussed in Smith (2008). The key issue in MacKinnon and Smith (1998) is how the bias function behaves. As we argued in Section 5.1, the bias is monotonic in the parameter of interest. Furthermore, Figure 2 shows that the bias is basically linear, so it is well behaved.

For the weighting matrix, we use the optimal weighting matrix suggested by Gouriéroux and Monfort (1996) for simulated method of moments estimators. Because the weighting matrix depends on our estimate of θ, we use a standard iterative procedure outlined in the next steps.

Step 4.—We make an initial guess of the weighting matrix $\mathbf{W}^0$ and solve for $\hat{\theta}^0$. Then, given this value we simulate the model to generate a new estimate of the weighting matrix.[10] With the new estimate of the weighting matrix we solve for a new $\hat{\theta}^1$. We perform this iterative procedure until our estimates of the weighting matrix and $\hat{\theta}$ converge. We explicitly consider simulation error because we utilize the weighting matrix suggested by Gouriéroux and Monfort (1996).

Step 5.—We compute standard errors using a bootstrap technique. We compute residuals from the data and the fitted values obtained using the estimates in (24), we resample the residuals with replacement, and we generate a new set of data using the fitted values. Using the data constructed from each resampling b, we computed new estimates β_1^b and β_2^b.

For each bootstrap b, we replace the moments β_1 and β_2 with bootstrap-generated moments β_1^b and β_2^b. To account for simulation error, a new seed is set to generate a new set of model-

[10]The computation of this matrix is described in Gouriéroux and Monfort (1996).

generated moments. Defining $y^b(\theta)$ as the difference in moments for each b, as in (25), we solve for:

$$\hat{\theta}^b = \arg\min_{\theta} \left[y^b(\theta)' \, \mathbf{W} \, y^b(\theta) \right]. \tag{27}$$

We repeat this exercise 100 times and we compute the standard error of our estimate of $\hat{\theta}$ as:

$$\text{S.E.}(\hat{\theta}) = \left[\frac{1}{100} \sum_{b=1}^{100} (\hat{\theta}^b - \hat{\theta})(\hat{\theta}^b - \hat{\theta})' \right]^{\frac{1}{2}}. \tag{28}$$

This procedure for constructing standard errors is similar in spirit to the approach of Eaton, Kortum, and Kramarz (2011), who use a simulated method of moments estimator to estimate the parameters of a trade model featuring micro-level heterogeneity from the performance of French exporters.

5.4. Performance on Simulated Data

In this section, we evaluate the performance of our estimation approach using simulated data when we know the true value of θ.

Table 3 presents the results from the following exercise. We generate two sets of artificial data on trade flows and disaggregate prices with true values of θ that are equal to 8.28 and 4.00, respectively, and then we apply our estimation routine.[11] We repeat this procedure 100 times. Table 3 reports average estimates. The sequence of artificial data is the same for both the overidentified case and the exactly identified case to facilitate comparisons across estimators.

The first row presents the average value of our simulated method of moments estimate, which is 8.29 with a standard error of 0.03. For all practical purposes, the estimation routine recovers the true value of θ that generated the data. To emphasize our estimator's performance, the next two rows of Table 3 present the approach of EK (which also corresponds to the moments used). Though not surprising given the discussion above, this approach generates estimates of θ that are significantly (in both their statistical and economic meaning) higher than the true value of θ of 8.28. The final two rows present the exactly identified case when we use only one moment to estimate θ. In this case, we use β_1. Similar to the overidentified case, the average value of our simulated method of moments estimate is 8.24 with a standard error of 0.04. Again, this is effectively the true value of θ.

The second column reports the results when the true value of θ is set equal to 4.00. The estimates using our estimator are 3.99 and 3.98 in the overidentified and the exactly identified

[11] To generate the artificial data set, we employ the same simulation procedure described in Steps 1-3 in Section 5.2 using the trade data from EK.

Table 3: Estimation Results With Artificial Data

Estimation Approach	**True θ = 8.28**	**True θ = 4.00**
Overidentified	Mean Estimate of θ (S.E.)	Mean Estimate of θ (S.E.)
SMM	8.29 (0.03)	3.99 (0.02)
Moment, β_1	12.47 (0.05)	6.03 (0.03)
Moment, β_2	15.20 (0.05)	7.34 (0.03)
Exactly Identified		
SMM	8.24 (0.04)	3.98 (0.02)
Moment, β_1	12.47 (0.05)	6.03 (0.03)

Note: In each simulation there are 19 countries, 150,000 goods and 100 simulations performed. The sequence of artificial data is the same for both the overidentified case and exactly identified case.

case, respectively. Similar to the previous results, these values are effectively the true value of θ. Furthermore, the alternative approaches that correspond to the moments that we used in our estimation are biased away from the true value of θ.

We also compare our estimation approach to an alternative statistical approach to bias reduction. Robson and Whitlock (1964) propose a way to reduce the bias when estimating the truncation point of a distribution. This problem is analogous to estimating the trade cost from price differences. This can be seen by inspecting the integral in (13) of Lemma 1. Robson and Whitlock's (1964) approach would suggest (in our notation) an estimator of the trade cost of $2\hat{\tau}^1_{ni} - \hat{\tau}^2_{ni}$, or two times the first-order statistic minus the second-order statistic. This makes intuitive sense because it increases the first-order statistic by the difference between the first- and second-order statistic. They show that this estimator is as efficient as the first-order statistic but with less bias.[12]

We apply their approach to approximate the trade friction and then use it as an input into the simple method of moments estimator. Table 4 compares the results from this estimation procedure to the results obtained using our SMM estimator. The second row reports the results when using Robson and Whitlock's (1964) approach to reduce the bias. This approach reduces the bias relative to using the first-order statistic (EK's approach) reported in the third row. It is not, however, a complete solution, as the estimates are still meaningfully higher than both the true value of θ and the estimates from our estimation approach. Moreover, we should emphasize Robson and Whitlock's (1964) approach only appeal is its computational simplicity. The

[12] Robson and Whitlock (1964) provide more-general refinements using inner-order statistics, but methods using inner-order statistics will have very low efficiency. Cooke (1979) provides an alternative bias reduction technique but only considers cases in which the sample size (L in our notation) is large.

Table 4: Comparison to Alternative Statistical Approaches to Bias Reduction

	True θ = 8.28	True θ = 4.00
Estimation Approach	Mean Estimate of θ (S.E.)	Mean Estimate of θ (S.E.)
SMM	8.29 (0.03)	3.99 (0.02)
Robson and Whitlock (1964)	10.54 (0.07)	5.11 (0.03)
Moment, β_1	12.47 (0.05)	6.03 (0.03)

Note: In each simulation there are 19 countries, 150,000 goods and 100 simulations performed. The sequence of artificial data is the same for all cases.

fact that the approach does not depend on the explicit distributional assumptions is not a benefit because without these assumptions the model does not yield a gravity equation. Without gravity, it is not clear what structural parameter is being estimated, which calls into question the entire enterprise.

Overall, we view these results as evidence supporting our estimation approach and empirical estimates of θ presented in Section 6 below.

6. Empirical Results

In this section, we apply our estimation strategy described in Section 5 to several different data sets. The key finding of this section is that our estimation approach yields an estimate around four, in contrast to previous estimation strategies, which yield estimates around eight.

6.1. Baseline Results Using New ICP 2005 Data

Our sample contains 123 countries. We use trade flows and production data for the year 2004 to construct trade shares. The price data used to compute aggregate price indices and proxies for trade costs come from basic-heading-level data from the 2005 round of the International Comparison Programme (ICP). The dataset has been employed in a number of empirical studies. For example, Bradford (2003) and Bradford and Lawrence (2004) use the ICP price data in order to measure the degree of fragmentation, or the level of trade barriers, among OECD countries. In addition, the authors provide an excellent description of the data-collection process.

The ICP collects price data on goods with identical characteristics across retail locations in the participating countries during the 2003-2005 period.[13] The basic-heading level represents a narrowly-defined group of goods for which expenditure data are available. The data set contains a total of 129 basic headings, and we reduce the sample to 62 categories based on their

[13]The ICP Methodological Handbook is available at http://go.worldbank.org/MW520NNFK0.

correspondence with the trade data employed. Appendix 1.2 provides more details.

On its own, this data set provides two contributions to the existing literature. First, because this is the latest round of the ICP, the measurement issues are less severe than in previous rounds. Furthermore, this data set provides very extensive coverage, as it includes as many as 123 developing and developed countries that account for 98 percent of world output.

The ICP provides a common list of "representative" goods whose prices are to be randomly sampled in each country over a certain period of time. A good is representative of a country if it comprises a significant share of a typical consumer's bundle there. Thus, the ICP samples the prices of a common basket of goods across countries, where the goods have been pre-selected due to their highly informative content for the purpose of international comparisons.

EK's model gives a natural common basket of goods to be priced across countries. In this model, agents in all countries consume all goods that lie within a fixed interval, $[0, 1]$. Thus, we consider this common list in the simulated model and randomly sample the prices of its goods across countries, in order to approximate trade barriers, much like it is done in the ICP data.

Table 5 presents the results.[14] The first row simply reports the moments that our estimation procedure targets. As discussed, these values correspond with EK's estimate of θ.

Table 5: Estimation Results With 2005 ICP Data

	Estimate of θ (S.E.)	β_1	β_2
Data Moments	—	7.75	9.61
Exactly Identified Case	4.12 (0.02)	7.75	—
Overidentified Case	4.06 (0.01)	7.65	9.62

The second row reports the results for exactly identified estimation, where the underlying moment used is β_1. In this case, our estimate of θ is 4.12, roughly half of EK's estimate of θ.

The third row reports the results for the overidentified estimation. The estimate of θ is 4.06—almost the same as in our exactly identified estimation and, again, roughly half of EK's estimate. The second and third columns report the resulting moments from the estimation routine, which are close to the data moments targeted, given that only one parameter is used to match two moments.

[14]The results from the Step 1 gravity regressions are presented in Table 12 and Table 15.

6.2. Estimates Using EK's Data

In this section, we apply our estimation strategy to the same data used in EK as another check of our estimation procedure. Their data set consists of bilateral trade data for 19 OECD countries in 1990 and 50 prices of manufactured goods for all countries. The prices come from a study conducted by the OECD. It is these same data that were included in a round of the ICP in the early 1990's. Similar to our data, the price data are at the basic-heading level and are for goods with identical characteristics across retail locations in the participating countries.

Table 6 presents the results.[15] The first row simply reports the moments that our estimation procedure targets. The entry in the third column corresponds with β_2, which is EK's baseline estimate of θ.

Table 6: Estimation Results With EK's Data

	Estimate of θ (S.E.)	β_1	β_2
Data Moments	—	5.93	8.28
Exactly Identified Case	3.93 (0.09)	5.93	—
Overidentified Case	4.42 (0.06)	6.64	8.10

The second row reports the results for exactly identified estimation, where the underlying moment used is β_1. In this instance, our estimate of θ is 3.93, which is, again, roughly half of EK's estimate of θ. The standard error of our estimate is fairly tight.

The third row reports the results for the overidentified estimation. Here, our estimate of θ is 4.42. Again, this is substantially below EK's estimate. Unlike our results in Table 5 with newer data, the overidentified case seems to be giving a different value than the exactly identified case gives. This contrasts with the Monte Carlo evidence, which suggests that the estimation procedure should not deliver very different estimates. Furthermore, comparing the data moments in the top row versus the implied moments in the second and third columns of the third row suggests that the estimation routine is facing challenges fitting the observed moments. We view this as pointing towards a problem with measurement error in the old data, as EK suggested.

6.3. Relation to Existing Literature

The elasticity of trade has been a focus of many studies. Below we discuss our method and results in relation to alternatives in the literature. We focus our discussion first on alternative

[15] The results from the Step 1 gravity regressions are presented in Table 13 and Table 15.

procedures that use price variation to approximate trade frictions and then on gravity-based estimators that use alternative proxies of trade frictions to estimate the trade elasticity.

EK provide a second estimate of the trade elasticity that amounts to 12.8.[16] Our critique and proposed solution apply to the estimator employed in this exercise as well. The critique applies because the alternative estimation approach is based on the same measures of trade frictions discussed above, which always underestimate the true trade friction. In Appendix C, we perform a Monte Carlo study where we find that EK's alternative methodology yields estimates that are nearly 100 percent higher than the true elasticity. Then, we employ a simulated method of moments estimator that minimizes the distance between the moments from EK's alternative approach on real and artificial data. We find that the estimate of θ is 4.39, which is essentially the same as our estimate in Table 6.

Donaldson (2009) estimates θ as well, and his approach is illuminating relative to the issues we have raised. His strategy for approximating trade costs is to study differences in the price of salt across locations in India. In principle, his approach is subject to our critique as well—i.e., how could price differences in one good be informative about trade frictions? However, he argues convincingly that in India, salt was produced in only a few locations and exported everywhere. Thus, by examining salt, Donaldson (2009) finds a "binding good". Using this approach, he finds estimates in the range of 3.8-5.2, which is consistent with our range of estimates of θ.

Anderson and van Wincoop (2004) survey the literature on trade-elasticity estimates obtained from gravity-based methods (which include EK's approach) and they find that the estimates range between five and ten. Excluding EK's results, the evidence cited in Anderson and van Wincoop (2004) comes from two alternative estimation approaches. The first uses second moments of changes in prices and expenditure shares, as in Feenstra (1994). The second uses the gravity equation with direct measures of trade barriers (i.e. tariffs), as in Head and Ries (2001). We discuss each of these approaches in turn below.

In Appendix D, we explore Feenstra's (1994) method in the context of the Ricardian model. We find that Feenstra's (1994) method, as well as papers that build on it such as Broda and Weinstein (2006), Imbs and Mejean (2010), and Feenstra, Obstfeld, and Russ (2010), does not recover the elasticity of trade in the Ricardian model with micro-level heterogeneity. In particular, we apply Feenstra's (1994) method to data generated from the Ricardian model and we show that the method recovers the utility parameter ρ that controls the elasticity of substitution across goods; not the trade elasticity θ. This utility parameter plays no role in determining aggregate trade flows and welfare gains from trade in the Ricardian model.[17] Hence, elasticity estimates

[16] Waugh (2010b) estimates the trade elasticity as well using EK's benchmark approach and hence our critique and solution applies to his approach as well.

[17] The parameter governs the elasticity of trade and welfare in models that do not feature micro-level heterogeneity such as the Armington model in Anderson and van Wincoop (2004) and Krugman's (1980) model (see

obtained using Feenstra's (1994) approach should not be used in quantitative analysis of the Ricardian model or other gravity-based international trade models with heterogeneity.[18]

The second set of estimates in Anderson and van Wincoop (2004) are obtained using direct measures of trade barriers in the gravity equation of trade. This methodology typically yields estimates in the range of five to ten and above. The estimation approach is appealing because of its simplicity and the appearance of being near assumption- and model-free, i.e. collect tariffs, run the regression in (21), and read off the coefficient on the tariff. However, this approach is neither assumption- nor model-free and in addition it suffers from significant data limitations.

First, in order to apply the estimation strategy that relies on tariff data to the Ricardian model with micro-level heterogeneity, the model must be able to generate a gravity equation. Given the model's utility specification, the assumption that productivity is drawn from a Fréchet distribution is crucial to obtain this result, as Arkolakis, Costinot, and Rodriguez-Clare (2011) argue. Similarly, should one wish to apply any of the estimation strategies to monopolistic-competition based models of heterogeneity, such as the framework of Melitz (2003), one would need to parametrize productivities as in Chaney (2008) in order to derive a gravity equation (see Arkolakis, Costinot, and Rodriguez-Clare (2011)). Hence, both our methodology and the approach that uses observed changes in tariffs share the same parametric restrictions.

Second, estimation strategies that use tariff data have to make strong assumptions on the form that both observable and unobservable components of trade frictions can take. For example, Head and Ries (2001) and Romalis (2007) assume that trade costs are symmetric, while Caliendo and Parro (2011) use triple differences across countries to avoid imposing the trade-cost symmetry assumption. Moreover, all three studies must assume that the unobservable component to trade frictions is uncorrelated with the observed component. This is likely to be violated during a typical liberalization episode, during which both tariff and (unobserved) non-tariff trade barriers fall. But, the magnitude of the change in the trade friction is critical to obtaining an unbiased estimate of the trade elasticity. For example, depending on how non-tariff barriers are controlled for in Head and Ries (2001), the estimate of the elasticity ranges from 11.4 to 7.9. Overall, since the method assigns the entire change in trade flows during trade liberalization to the change in tariffs, it is not surprising that the resulting elasticity estimates are high.

Third, data limitations arise because one needs adequate measures of tariffs to identify the elasticity from observed trade flows. To satisfy this restriction, researchers have typically concentrated on estimating the elasticity for the U.S., Canada, and other rich countries. It is not clear whether these estimates are applicable when addressing important questions such as:

Arkolakis, Costinot, and Rodriguez-Clare (2011)).

[18] Feenstra (2010) makes a similar argument in the context of the Melitz (2003) model when parametrized as in Chaney (2008).

How large are the welfare gains from trade for poor and developing countries?[19] Our method, in combination with a small sample of comparable price data, allows us to estimate the trade elasticity for 123 countries. To the best of our knowledge, this is the widest coverage across rich and poor countries for which the trade elasticity has been estimated.

Admittedly, there is a substantial difference between the low values of the elasticity that our approach yields and the high values obtained using tariff data. One piece of evidence in support of the values that we find is that our results compare favorably with alternative estimates of the productivity parameter θ that *do not* use gravity-based estimators. For example, estimates of θ using firm-level sales data, as in Bernard, Eaton, Jensen, and Kortum (2003) and Eaton, Kortum, and Kramarz (2011), are in the range of 3.6 to 4.8—exactly in the range of values that we find. Burstein and Vogel (2009) estimate θ matching moments regarding the skill intensity of trade and find a value of five. The identifying source of variation in Bernard, Eaton, Jensen, and Kortum (2003) and Eaton, Kortum, and Kramarz (2011) is that firm level data suggest that there is a lot of variation in productivity. The data in our paper is telling a similar story: price variation (once properly corrected) suggests that there is a lot of variation in productivity implying a relatively low trade elasticity.

Finally, we want to point out that, like methods that use tariffs and the gravity equation, our methodological approach is not specific to the Ricardian model. The methodology and the moments that we use to estimate the trade elasticity within the Ricardian framework can be derived from other trade models of heterogeneity that generate a gravity equation of trade. The key distinction, however, is that even with the same data, different assumptions about the particular model may give different estimates of the trade elasticity.

This point has broader implications in light of Arkolakis, Costinot, and Rodriguez-Clare's (2011) arguments. The authors argue that many trade models yield the same welfare-gain equation in (6). Hence, given a common estimator of the trade elasticity (perhaps one that uses tariffs and the gravity equation that the models generate), the welfare gains across different models are the same. In contrast, our methodological approach provides a common estimator that is gravity-based, but is not trivially the same across models. Thus, one can use the estimator to test the assumption that new models of heterogeneity have the same trade elasticity and therefore the same welfare gains from trade.

[19] Two extensions of the Eaton and Kortum (2002) framework have focused on this question. Waugh (2010b) shows that poor countries have the most to gain from international trade and asymmetry in trade barriers accounts for a third of cross-country income differences. Fieler (2010) argues that the trade elasticity for poor countries is fifty percent higher than the elasticity for rich countries, and that this difference affects the welfare gains from trade. Our elasticity estimates (relative to EK's) increase the welfare cost of asymmetries in trade frictions in Waugh's (2010b) analysis by over thirty percent. Contrary to Fieler's (2010) findings, we find evidence that trade elasticities are similar or lower for poor countries relative to rich countries.

7. Robustness

Below we discuss some computational issues regarding the number of goods in the simulation and measurement issues regarding the price data.

7.1. The Number of Goods

The estimation routine requires us to take a stand on the number of goods in the economy. We argue that the appropriate way to view this issue is to ask: how many goods are needed to numerically approximate the infinite number of goods in the model? Thus, the number of goods chosen should be judged on the accuracy of the approximation relative to the computational cost. The choice of the number of goods should not be judged on the basis of how many goods actually exist in the "real world" because this value is impossible to know or discipline.

To understand our argument, recall that our estimation routine is based on a moment condition that compares a biased estimate from the data with a biased estimate using artificial data. In Section 3.4, we argued that the bias depends largely on the expected value of the max over a finite sample of price differences—i.e., the integral of the left-hand side of equation (13). Thus, when we compute the biased estimate using artificial data, we are effectively computing this integral via simulation.[20] This suggests that the number of goods should be chosen in a way that delivers an accurate approximation of the integral. Furthermore, a way to judge if the number of goods selected delivers an appropriate approximation is to increase the number of goods until the estimate of θ does not change too much.

Table 7 reports the results of this analysis. It shows how our estimate of θ varies as the number of goods in the economy changes, using the EK data and the 2005 ICP data. For the EK data, notice that our estimates are relatively similar across all the different numbers of goods employed, ranging from 4.14 to 3.93. Moreover, the estimates are effectively the same after the number of goods is above 100,000, suggesting that this is a reasonable starting point.

The results obtained using the 2005 ICP data, which features 123 countries, vary more depending upon the number of goods used. While the change from 4.22 to 4.12 when going from 100,000 to 150,000 goods is numerically large, computational costs force our hand to settle on 150,000 as the number of goods in the economy.[21]

Table 7 also reports a side effect of using a low number of goods—zero trade flows between countries predicted by the model in places where we observe positive trade flows in the data.

[20] An alternative estimation strategy would be to use different numerical methods to compute the integrals (13) and then to adjust the EK estimator given this value.

[21] The reason is that 150,000 goods is near the maximum number of goods feasible while still being able to execute the simulation routine in parallel on a multi-core machine, which allows a speed-up of just under a factor of eight.

Table 7: Results with Different # of Goods

Number of Goods	5,000	25,000	100,000	150,000*
EK's Data, Exactly Identified Case, $\hat{\theta}$	4.14	3.99	3.93	3.93
Fraction of Wrong Zeros	0.10	0.03	0.005	0.003
Fraction of Correct Zeros	—	—	—	—
2004 ICP Data, Exactly Identified Case, $\hat{\theta}$	5.54	4.67	4.22	4.12
Fraction of Wrong Zeros	0.46	0.31	0.21	0.18
Fraction of Correct Zeros	0.85	0.72	0.55	0.50

Table 7 reports the fraction of zeros that the model produces in instances where there are positive trade flows observed in the data. With only 5,000 goods, using the 2005 ICP data set, in almost half of the instances where trade flows are observed in the data, the model generates a zero. While not as severe, ten percent of positive trade flows are assigned zeros with the EK data. Results of this nature suggest increasing the number of goods to minimize the number of wrong zeros, as well.

7.2. Measurement Issues

The price data that we employ in our benchmark analysis constitute the 2005 round of the ICP. The latest ICP round represents the highest quality cross-country price data that are publicly available. As discussed earlier, the goal of the ICP is to collect prices of comparable products in retail locations around the world. Since the data are meant to be comparable across countries, authors such as Bradford (2003) and Bradford and Lawrence (2004) have used it in order to infer the degree of fragmentation, or the level of trade barriers, among different countries.

With this in mind, measurement error in price data is a general concern in empirical work. Below, we discuss measurement error issues that relate to distribution costs, mark-ups, product quality, and aggregation, and the possible biases they may introduce in our estimation. We conclude that these sources of measurement error potentially affect our estimation in various and different directions. Thus, in order to address them, one would need to take a stand on the mechanism that potentially generates them and incorporate it in the estimation procedure. The advantage of our simulation-based estimator that relies on a model is that it can accommodate these extensions. Below, we preview how various mechanisms may affect the results and we offer further avenues of research on this topic.

Random Measurement Error. In our estimation, non-systematic measurement error in the data (mean-zero measurement error) may artificially generate *larger* maximal price differences than implied by the underlying model. This would result in estimates of θ that are biased downwards. To address this issue, one can take a stand on the form of measurement error, introduce the simulated measurement error into the artificial data set, and potentially estimate it jointly with the trade elasticity. While we feel that a formal treatment of this issue is beyond the scope of this paper, we performed Monte Carlo experiments with additive log-normal measurement error. We found that the magnitude of measurement error needed to affect our results was extreme.

Distribution Costs and Markups. The price data used in our estimation were collected at the retail level. These prices may reflect local distribution costs, sales taxes, and mark-ups. As long as the frictions are multiples over marginal costs of production and they are country- but not good-specific, they will not affect our estimates of the elasticity parameter. Mathematically, one can see this by noting that any multiplicative country-specific effect cancels out in the denominator of equation (12). This is an important reason for using β as a moment in our estimation routine rather than some other moment.

What if these effects are not multiplicative? For example, Burstein, Neves, and Rebelo (2003) present a model where distribution margins over tradable goods are additive. To understand the effects of additive distribution margins, we carry out the experiment described in Section 4: we simulate trade flows and samples of micro-level prices under a known θ and then we introduce additive distribution costs to understand the bias that these effects may introduce.

These Monte Carlo experiments show that the bias in β relative to θ is *larger* than in the cases when additive distribution costs are not present. The reason is because additive distribution costs increase low prices proportionally more then high prices, thus the maximum price difference is smaller than it would be otherwise. Because of the strong monotonicity between β and θ, this suggests that incorporating additive distribution costs to the estimation would make our estimates of θ even lower.

Mark-ups that are not only country-, but also firm/retailer-specific is an important issue that is beyond the scope of this paper. However, Simonovska and Waugh (2011) applies the estimator proposed in this paper to Bernard, Eaton, Jensen, and Kortum's (2003) Ricardian framework which has firm- and country-specific markups. An interesting feature of this analysis is that while both Bernard, Eaton, Jensen, and Kortum (2003) and EK yield the same formula for the welfare gains from trade (see Arkolakis, Costinot, and Rodriguez-Clare (2011)), viewing the price data in light of the Bernard, Eaton, Jensen, and Kortum (2003) and our estimation procedure may yield different estimates for the elasticity of trade than EK. Hence different models with different micro-market structures may have different welfare gains from trade.

Aggregation. The basic-heading price data employed in our analysis are disaggregated; but, they are not at the individual-good level. For example, a price observation titled "rice" contains the average price across different types of rice sampled, such as basmati rice, wild rice, whole-grain rice, etc. Suppose that basmati rice is the binding good for a pair of countries. In the ICP data, we compute the difference between the average price of rice between the two countries, which is smaller than the price difference of basmati rice, if the remaining types of rice are more equally priced across the two countries. In this case, trade barriers are underestimated and, consequently, the elasticity of trade is biased upwards.

In order to alleviate the aggregation problem, Table 8 presents estimates of the elasticity parameter using EIU's good-level price data set, which spans a subset of 77 countries from our original data set, but provides prices for 111 individual tradable goods in two types of retail stores. The results from the Step 1 gravity regressions are presented in Table 14 and Table 15.

Table 8: Estimation Results With EIU Data

	Estimate of θ (S.E.)	β_1	β_2
Cheap Stores			
Data Moments	—	4.17	5.11
Exactly Identified Case	2.47 (0.02)	4.17	—
Overidentified Case	2.48 (0.02)	4.19	5.09
Expensive Stores			
Data Moments	—	4.39	5.23
Exactly Identified Case	2.60 (0.02)	4.39	—
Overidentified Case	2.54 (0.02)	4.29	5.21

Table 8 suggest that aggregation causes a downward bias on trade-barrier estimates, resulting in trade-elasticity estimates that are biased upwards. Indeed, when we use the highly disaggregate EIU dataset, the elasticity of trade falls to 2.47-2.60.

Table 8 also reports the estimates for cheap and expensive stores as reported in the EIU dataset. The estimates are very similar despite the very different stores from which the price data are collected. A conclusion from this finding is that retail-specific markups within a country—which would presumably characterize the difference between the two types of stores—are not systematically biasing the results in a meaningful way.

Varying Product Quality. Although the explicit goal of price collection enterprises such as the ICP and the EIU is that comparable products are priced across countries, one may be concerned that the prices used in the estimation reflect varying quality levels.

To address this concern, we engage in two exercises. The first involves the EIU data. This dataset features prices of comparable goods across countries collected in two different retail locations: high-end stores and low-end stores. It is reasonable to assume that prices collected in a particular store type across countries are more comparable to each other than prices collected in different types of stores. Hence, in Table 8 above, we provide our estimates of the elasticity using each subsample. The estimates are very similar across the two exercises, which is reassuring of our results.

The second exercise that we perform involves the ICP data. In particular, we estimate the trade elasticity using the 2005 trade and ICP price data for the subset of nineteen developed OECD countries that EK considered in their original analysis. The presumption here is that should there exist systematic differences in product quality across countries, they will be minimized within this sample of developed and more homogeneous economies.

EK's method of moments estimator for this subset of countries yields an estimate of 7.19. Our simulated method of moments estimator yields an estimate of 5.51 with a standard error of 0.11. This value remains on the lower end of the estimates that prevail in the existing literature. The reason why the value is higher than our benchmark estimate is because, in 2005, OECD countries were more open relative to non-OECD countries as well as relative to their own 1990 counterparts. Hence, the bias associated with EK's estimation approach is lower when trade barriers are lower (recall Figure 1). An interesting point of comparison are non-OECD countries. The EK method of moments estimator for this subset is 8.31—larger than the OECD estimate. Yet, our simulated method of moments estimator yields an estimate of 4.09, which is lower than the OECD sub-sample. This is again because non-OECD countries are less open and thus the bias correction is lager.

We conclude that various sources of measurement error potentially affect our estimation in different directions. Thus, there would not be one particular bias in our estimate if we had incorporated all of these features into our estimation. However, an advantage of our estimation approach is that future research on this topic can accommodate these extensions and evaluate their potential effects.

8. Conclusion

In this paper we develop a new methodology to estimate the elasticity of trade that builds on the Ricardian model of international trade with micro-level heterogeneity. We apply our estimator to novel disaggregate price and trade-flow data for the year 2004, which spans 123 countries that account for 98 percent of world output. Across numerous exercises, we obtain estimates of the trade elasticity that range between 2.47 to 5.51. These values are both lower and fall within a narrower range relative to the existing literature. Our findings imply that the measured welfare gains from international trade are twice as high as previously documented.

More importantly, the methodology and the moments that we use to estimate the trade elasticity within the Ricardian framework can be derived from other trade models of heterogeneity that generate a gravity equation of trade. The key distinction, however, is that even with the same data, different assumptions about the particular model may give different estimates of the trade elasticity.

This point has broader implications in light of Arkolakis, Costinot, and Rodriguez-Clare's (2011) arguments. The authors argue that different trade models yield the same welfare-gain equation. Hence, given a common estimator of the trade elasticity, the welfare gains across different models are the same. In contrast, our methodological approach provides a common estimator that is gravity-based, but is not trivially the same across models. Thus, one can use the estimator to test the assumption that new models of heterogeneity have the same trade elasticity and therefore the same welfare gains from trade.

References

ALVAREZ, F., AND R. J. 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 (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93, 170–192.

——— (2004): "Trade Costs," *Journal of Economic Literature*, 42(3), 691–751.

ARKOLAKIS, C., A. COSTINOT, AND A. RODRIGUEZ-CLARE (2011): "New Trade Models, Same Old Gains?," *American Economic Review (forthcoming)*.

BERNARD, A., J. EATON, J. B. JENSEN, AND S. KORTUM (2003): "Plants and Productivity in International Trade," *American Economic Review*, 93(4), 1268–1290.

BRADFORD, S. (2003): "Paying the Price: Final Goods Protection in OECD Countries," *The Review of Economics and Statistics*, 85(1), 24–37.

BRADFORD, S., AND R. Z. LAWRENCE (2004): *Has Globalization Gone Far Enough: The Costs of Fragmented Markets*, no. 349 in Peterson Institute Press: All Books. Peterson Institute for International Economics.

BRODA, C., AND D. WEINSTEIN (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(2).

BURSTEIN, A., J. NEVES, AND S. REBELO (2003): “Distribution costs and real exchange rate dynamics during exchange-rate-based stabilizations,” *Journal of Monetary Economics*, 50(6), 1189–1214.

BURSTEIN, A., AND J. VOGEL (2009): “Globalization, Technology, and the Skill Premium,” Discussion paper.

CALIENDO, L., AND F. PARRO (2011): “Estimates of the Trade and Welfare Effects of NAFTA,” *mimeo*.

CHANEY, T. (2008): “Distorted Gravity: The Intensive and Extensive Margins of International Trade,” *American Economic Review*, 98(4), 1707–1721.

COOKE, P. (1979): “Statistical inference for bounds of random variables,” *Biometrika*, 66(2), 367.

COSTINOT, A., D. DONALDSON, AND I. KOMUNJER (2011): “What Goods Do Countries Trade? A Quantitative Exploration of Ricardo’s Ideas,” *Review of Economic Studies, forthcoming*.

DAVIDSON, R., AND J. MACKINNON (2004): *Econometric Theory and Methods*. New York: Oxford University Press.

DONALDSON, D. (2009): “Railroads of the Raj: Estimating the Economic Impact of Transportation Infrastructure,” *unpublished paper, MIT*.

EATON, J., AND S. KORTUM (2002): “Technology, Geography, and Trade,” *Econometrica*, 70(5), 1741–1779.

EATON, J., S. KORTUM, AND F. KRAMARZ (2011): “An Anatomy of International Trade: Evidence from French Firms,” *Econometrica (forthcoming)*.

FEENSTRA, R. (1994): “New product varieties and the measurement of international prices,” *The American Economic Review*, 84(1), 157–177.

——— (2010): *Product Variety and the Gains from International Trade*. MIT Press.

FEENSTRA, R. C., M. OBSTFELD, AND K. N. RUSS (2010): “In Search of the Armington Elasticity,” *University of California - Davis, unpublished mimeo*.

FIELER, A. C. (2010): “Non-Homotheticity and Bilateral Trade: Evidence and a Quantitative Explanation,” *forthcoming, Econometrica*.

GOURIÉROUX, C., AND A. MONFORT (1996): *Simulation-based econometric methods*, CORE lectures. Oxford University Press.

HAYASHI, F. (2000): *Econometrics*. Princeton.

HEAD, K., AND J. RIES (2001): "Increasing Returns versus National Product Differentiation as an Explanation for the Pattern of U.S.-Canada Trade," *American Economic Review*, 91(4), 858–876.

IMBS, J., AND I. MEJEAN (2010): "Trade Elasticities," *Paris School of Economics, unpublished mimeo*.

KRUGMAN, P. (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70(5), 950–959.

MACKINNON, J. G., AND A. A. SMITH (1998): "Approximate bias correction in econometrics," *Journal of Econometrics*, 85(2), 205–230.

MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725.

MUENDLER, M.-A. (2009): "Converter from SITC to ISIC," *University of California - San Diego, unpublished mimeo*.

ROBSON, D., AND J. WHITLOCK (1964): "Estimation of a truncation point," *Biometrika*, 51(1-2), 33.

ROMALIS, J. (2007): "NAFTA's and CUSFTA's Impact on International Trade," *Review of Economics and Statistics*, 89(3), 416–35.

SIMONOVSKA, I., AND M. WAUGH (2011): "Different Trade Models, Different Elasticities?," *New York University and Princeton University, mimeo*.

SMITH, A. (2008): "Indirect inference," *The New Palgrave Dictionary of Economics, 2nd Edition (forthcoming)*.

WAUGH, M. (2010a): "Bilateral Trade, Relative Prices, and Trade Costs," *unpublished paper, NYU*.

WAUGH, M. E. (2010b): "International Trade and Income Differences," *American Economic Review*, 100(5), 2093–2124.

YI, K.-M. (2003): "Can Vertical Specialization Explain the Growth of World Trade?," *Journal of Political Economy*, 111(1), 52–102.

A. Data Appendix

1.1. Trade Shares

To construct trade shares, we used bilateral trade flows and production data as follows:

$$\frac{X_{ni}}{X_n} = \frac{\text{Imports}_{ni}}{\text{Gross Mfg. Production}_n - \text{Exports}_n + \text{Imports}_n},$$

$$\frac{X_{nn}}{X_n} = 1 - \sum_{k \neq n}^{N} \frac{X_{ni}}{X_n}.$$

Putting the numerator and denominator together is simply computing an expenditure share by dividing the value of goods country n imported from country i by the total value of goods in country n. The home trade share $\frac{X_{nn}}{X_n}$ is simply constructed as the residual from one minus the sum of all bilateral expenditure shares.

To construct $\frac{X_{ni}}{X_n}$, the numerator is the aggregate value of manufactured goods that country n imports from country i. Bilateral trade-flow data are from UN Comtrade for the year 2004. We obtain all bilateral trade flows for our sample of 123 countries at the four-digit SITC level. We then used concordance tables between four-digit SITC and three-digit ISIC codes provided by the UN and further modified by Muendler (2009).[22] We restrict our analysis to manufacturing bilateral trade flows only—namely, those that correspond with manufacturing as defined in ISIC Rev.#2.

The denominator is gross manufacturing production minus manufactured exports (for only the sample) plus manufactured imports (for only the sample). Gross manufacturing production data are the most serious data constraint we faced. We obtain manufacturing production data for 2004 from UNIDO for a large sub-sample of countries. We then imputed gross manufacturing production for countries for which data are unavailable as follows: We first obtain 2004 data on manufacturing (MVA) and agriculture (AVA) value added, as well as population size (L) and GDP for all countries in the sample. We then impute the gross output (GO) to manufacturing value added ratio for the missing countries using coefficients resulting from the following regression:

$$\log\left(\frac{MVA}{GO}\right) = \beta_0 + \beta_{GDP}C_{GDP} + \beta_L C_L + \beta_{MVA}C_{MVA} + \beta_{AVA}C_{AVA} + \epsilon,$$

where β_x is a 1x3 vector of coefficients corresponding to C_x, an Nx3 matrix which contains

[22] The trade data often report bilateral trade flows from two sources. For example, the exports of country A to country B can appear in the UN Comtrade data as exports reported by country A or as imports reported by country B. In this case, we take the report of bilateral trade flows between countries A and B that yields a higher total volume of trade across the sum of all SITC four-digit categories.

$[\log(x), (\log(x))^2, (\log(x))^3]$ for the sub-sample of N countries for which gross output data are available.

1.2. Prices

The ICP price data we employ in our estimation procedure is reported at the basic-heading level. A basic heading represents a narrowly-defined group of goods for which expenditure data are available. For example, basic heading "1101111 Rice" is made up of prices of different types of rice, and the resulting value is an aggregate over these different types of rice. This implies that a typical price observation of "Rice" contains different types of rice, as well as different packaging options that affect the unit price of rice within and across countries.

According to the ICP Handbook, the price of the basic heading "Rice" is constructed using a transitive Jevons index of prices of different varieties of rice. To illustrate this point, suppose that the world economy consists of three countries, A, B, C and ten types of rice, 1-10. Further suppose that consumers in country A have access to all 10 types of rice; those in country B only have access to types 1-5 of rice; and those in country C have access to types 4-6 of rice. Although all types of rice are not found in all three countries, it is sufficient that each pair of countries shares at least one type of rice.

The ICP obtains unit prices for all available types of rice in all three countries and records a price of 0 if the type of rice is not available in a particular country. The relative price of rice between countries A and B, based on goods available in these two countries, $p_{AB}^{A,B}$, is a geometric average of the relative prices of rice of types 1-5

$$p_{AB}^{A,B} = \left[\prod_{j=1}^{5} \frac{p_A(j)}{p_B(j)}\right]^{\frac{1}{5}}.$$

Similarly, one can compute the relative price of rice between countries A and C (B and C) based on varieties available in both A and C (B and C). The price of the basic heading "Rice" reported by the ICP is:

$$p_{AB} = \left[p_{AB}^{A,B} p_{AB}^{A,B} \frac{p_{AC}^{A,C}}{p_{BC}^{B,C}}\right]^{\frac{1}{3}},$$

which is a geometric average that features not only relative prices of rice between countries A and B, but also cross-prices between A and B linked via country C. This procedure ensures that prices of basic headings are transitive across countries and minimizes the impact of missing prices across countries.

Thus, a basic-heading price is a geometric average of prices of varieties that is directly comparable across countries.

B. Proofs

Below, we describe the steps to proving Lemmata 1 and 2. The key part in Lemma 1 is deriving the distribution of the maximal log price difference. We then prove Propositions 1 and 2.

2.1. Proof of Lemma 1, Lemma 2, and Proposition 1

First, we derive the distribution of the maximal log price difference. The key insight is to work with direct comparisons of goods' prices (i.e., do not impose equilibrium and work from the equilibrium price distribution) and to compute the distribution of log price differences and then the distribution of the maximal log price difference.

Having obtained the distribution of the maximum log price difference, we show that the expected value of the maximum log price difference is biased in a finite sample and the estimator $\hat{\beta}$ is biased.

2.1.A. Preliminaries

In deriving the distribution of maximum log price differences, we will work with a relabeling of the production functions and exponential distributions following an argument in Alvarez and Lucas (2007). They relate the pdfs of the exponential and Frèchet distributions. The claim is that if $z_i \sim \exp(T_i)$, then $y_i \equiv z_i^{-\frac{1}{\theta}} \sim \exp(-T_i y_i^{-\theta})$. To see this, notice that since $y_i = h(z_i)$ is a decreasing function, it must be that $f(z_i)dz_i = -g(y_i)dy_i$, where f, g are the pdf's of z_i, y_i, respectively. The result will allow us to characterize moments of the log price difference by invoking properties of the exponential distribution.

2.1.B. Proof of Lemma 1

The proof of Lemma 1 follows.

Let $z_k^{-\frac{1}{\theta}} \sim \exp(-T_k(z_k^{-\frac{1}{\theta}})^{-\theta})$ be the productivity associated with good z, drawn from the Frèchet pdf in country k. By the argument above, the underlying distribution of z_k is exponential. The price for good z produced in country k and supplied to country i is $p_{ik} \equiv w_k \tau_{ik} z_k^{\frac{1}{\theta}}$. The relative price ratio of good z between countries n and i is:

$$v_{ni}(z) = \frac{\min\left\{\min_{k\neq i}[w_k \tau_{nk} z_k^{\frac{1}{\theta}}], w_i \tau_{ni} z_i^{\frac{1}{\theta}}\right\}}{\min\left\{\min_{k\neq n}[w_k \tau_{ik} z_k^{\frac{1}{\theta}}], w_n \tau_{in} z_n^{\frac{1}{\theta}}\right\}}. \tag{b.1}$$

Take this object to the power of θ:

$$(v_{ni}(z))^{\theta} = \frac{\min\left\{\min_{k\neq i}[w_k^{\theta}\tau_{nk}^{\theta}z_k], w_i^{\theta}\tau_{ni}^{\theta}z_i\right\}}{\min\left\{\min_{k\neq n}[w_k^{\theta}\tau_{ik}^{\theta}z_k], w_n^{\theta}\tau_{in}^{\theta}z_n\right\}}. \tag{b.2}$$

We want to characterize the distribution of (b.2), so we will first derive the pdf's of its components. Define $\tilde{z}_{ik} = w_k^{\theta}\tau_{ik}^{\theta}z_k$. Since $z_k \sim \exp(T_k)$, it must be that $\tilde{z}_{ik} \sim \exp(T_k w_k^{-\theta}\tau_{ik}^{-\theta})$. Let $\tilde{\lambda}_{ik} \equiv T_k w_k^{-\theta}\tau_{ik}^{-\theta}$.

Next, we derive the distribution of $\tilde{z}_i \equiv \min_{k\neq n}[w_k^{\theta}\tau_{ik}^{\theta}z_k] = \min_{k\neq n}[\tilde{z}_{ik}]$. Since each $\tilde{z}_{ik} \sim \exp(\tilde{\lambda}_{ik})$ and independent across countries k, $\tilde{z}_i \sim \exp(\sum_{k\neq n}\tilde{\lambda}_{ik})$. Define $\tilde{\lambda}_i \equiv \sum_{k\neq n}\tilde{\lambda}_{ik}$. Repeat the procedure for importer n in the numerator.

Given these definitions, (b.2) can be rewritten as:

$$(v_{ni}(z))^{\theta} = \frac{\min\left\{\tilde{z}_n, w_i^{\theta}\tau_{ni}^{\theta}z_i\right\}}{\min\left\{\tilde{z}_i, w_n^{\theta}\tau_{in}^{\theta}z_n\right\}}. \tag{b.3}$$

Define $\epsilon_{ni}(z) = \log(v_{ni}(z))$. Taking logs of expression (b.3) gives:

$$\begin{aligned} \theta\epsilon_{ni}(z) = \quad & \min\left\{\log(\tilde{z}_n), [\theta\log(w_i) + \theta\log(\tau_{ni}) + \log(z_i)]\right\} \\ - \quad & \min\left\{\log(\tilde{z}_i), [\theta\log(w_n) + \theta\log(\tau_{in}) + \log(z_n)]\right\}. \end{aligned} \tag{b.4}$$

Next, we argue that $\theta\epsilon_{ni}(z) \in [-\theta\log(\tau_{in}), \theta\log(\tau_{ni})]$. For any good z, $\theta\epsilon_{ni}(z)$ can satisfy one and only one of the following three cases:

1. Countries n and i buy good z from two different sources. Then,

$$\theta\epsilon_{ni}(z) = \log(\tilde{z}_n) - \log(\tilde{z}_i) \tag{b.5}$$

2. Country n buys good z from country i. Assuming that trade barriers don't violate the triangle inequality, it must be that i buys the good from itself. Then,

$$\theta\epsilon_{ni}(z) = \theta\log(w_i) + \theta\log(\tau_{ni}) + \log(z_i) - \theta\log(w_i) - \log(z_i) = \theta\log(\tau_{ni}). \tag{b.6}$$

3. Country i buys good z from n. Then it must be that n buys the good from itself, so:

$$\theta\epsilon_{ni}(z) = \theta\log(w_n) + \log(z_n) - \theta\log(w_n) - \theta\log(\tau_{in}) - \log(z_n) = -\theta\log(\tau_{in}). \tag{b.7}$$

We claim that the following ordering occurs: $-\theta\log(\tau_{in}) \leq \log(\tilde{z}_n) - \log(\tilde{z}_i) \leq \theta\log(\tau_{ni})$. To show this, we need to consider the following two scenarios:

1. Countries n and i buy good z from the same source k. Then,

$$\begin{aligned}\log(\tilde{z}_n) - \log(\tilde{z}_i) &= \log(w_k^\theta \tau_{nk}^\theta z_k) - \log(w_k^\theta \tau_{ik}^\theta z_k) \\ &= \theta(\log(\tau_{nk}) - \log(\tau_{ik})). \end{aligned} \tag{b.8}$$

   Clearly,

$$\theta(\log(\tau_{nk}) - \log(\tau_{ik})) \geq -\theta \log(\tau_{in}) \iff \tau_{in}\tau_{nk} \geq \tau_{ik},$$

   where the latter inequality is true under the triangle inequality assumption.

   Similarly,

$$\theta(\log(\tau_{nk}) - \log(\tau_{ik})) \leq \theta \log(\tau_{ni}) \iff \tau_{nk} \leq \tau_{ni}\tau_{ik},$$

   again true by triangle inequality.

2. Country n buys good z from source a and country i from source b, $a \neq b$. We want to show that $-\theta \log(\tau_{in}) \leq \log(w_a^\theta \tau_{na}^\theta z_a) - \log(w_b^\theta \tau_{ib}^\theta z_b) \leq \theta \log(\tau_{ni})$.

   Since n imported from a over b, it must be that:

$$w_a^\theta \tau_{na}^\theta z_a \leq w_b^\theta \tau_{nb}^\theta z_b \tag{b.9}$$

   Similarly, since i imported from b over a, it must be that:

$$w_b^\theta \tau_{ib}^\theta z_b \leq w_a^\theta \tau_{ia}^\theta z_a \tag{b.10}$$

   To find the upper bound, take logs of (b.9) and subtract $\log(w_b^\theta \tau_{ib}^\theta z_b)$ from both sides:

$$\log(w_a^\theta \tau_{na}^\theta z_a) - \log(w_b^\theta \tau_{ib}^\theta z_b) \leq \log(w_b^\theta \tau_{nb}^\theta z_b) - \log(w_b^\theta \tau_{ib}^\theta z_b) \tag{b.11}$$

   It suffices to show that the right-hand side is itself below the upper bound since, by transitivity, so is the left-hand side (which is the object of interest).

$$\begin{aligned} &\log(w_b^\theta \tau_{nb}^\theta z_b) - \log(w_b^\theta \tau_{ib}^\theta z_b) \leq \theta \log(\tau_{ni}) \\ \iff &\theta \log(\tau_{nb}) - \theta \log(\tau_{ib}) \leq \theta \log(\tau_{ni}) \\ \iff &\tau_{nb} \leq \tau_{ni}\tau_{ib}, \end{aligned} \tag{b.12}$$

   which is true by triangle inequality.

   The argument for the lower bound is similar. Take logs of (b.10), multiply by -1 (and reverse inequality) and add $\log(w_a^\theta \tau_{na}^\theta z_a)$ to both sides:

$$\log(w_a^\theta \tau_{na}^\theta z_a) - \log(w_b^\theta \tau_{ib}^\theta z_b) \geq \log(w_a^\theta \tau_{na}^\theta z_a) - \log(w_a^\theta \tau_{ia}^\theta z_a) \tag{b.13}$$

It suffices to show that the right-hand side is itself above the lower bound since, by transitivity, so is the left-hand side (which is the object of interest).

$$\begin{aligned} & \log(w_a^\theta \tau_{na}^\theta z_a) - \log(w_a^\theta \tau_{ia}^\theta z_a) \geq -\theta \log(\tau_{in}) \\ \iff & \theta \log(\tau_{na}) - \theta \log(\tau_{ia}) \geq -\theta \log(\tau_{in}) \\ \iff & \tau_{in}\tau_{na} \geq \tau_{ia}, \end{aligned} \tag{b.14}$$

which is true by triangle inequality.

Hence, $\theta\epsilon_{ni}(z) \in [-\theta\log(\tau_{in}), \theta\log(\tau_{ni})]$.

Next, we proceed to derive the distribution of $\theta\epsilon_{ni}(z) = \log(\tilde{z}_n) - \log(\tilde{z}_i)$. First, we derive the pdfs of its two components.

Let $y_i \equiv \log(\tilde{z}_i)$. Then $\tilde{z}_i = \exp(y_i)$. The pdf of y_i must satisfy:

$$\begin{aligned} f(y_i)dy_i = g(\tilde{z}_i)d\tilde{z}_i \;\Rightarrow\quad & f(y_i) = \tilde{\lambda}_i \exp(-\tilde{\lambda}_i \tilde{z}_i)\frac{d\tilde{z}_i}{dy_i} \\ \Rightarrow\quad & f(y_i) = \tilde{\lambda}_i \exp(-\tilde{\lambda}_i \exp(y_i))\exp(y_i) \\ \Rightarrow\quad & F(y_i) = 1 - \exp(-\tilde{\lambda}_i \exp(y_i)) \end{aligned} \tag{b.15}$$

The same holds for n.

Now that we have the pdf's of the two components, we can define the pdf of $\epsilon \equiv \theta\epsilon_{ni}(z) \in [-\theta\log(\tau_{in}), \theta\log(\tau_{ni})]$ as follows:

$$f(\epsilon) \equiv f_{y_n - y_i}(x) = \int_{-\infty}^{\infty} f_{y_n}(y) f_{y_i}(y - x)dy, \tag{b.16}$$

where we have used the fact that y_n and y_i are independently distributed hence, the pdf of their difference is the convolution of the pdfs of the two random variables.

Substituting the pdfs of y_n and y_i into (b.16) yields:

$$\begin{aligned} f(\epsilon) = & \int_{-\infty}^{\infty} \tilde{\lambda}_n \exp(-\tilde{\lambda}_n \exp(y))\exp(y)\tilde{\lambda}_i \exp(-\tilde{\lambda}_i \exp(y-\epsilon))\exp(y-\epsilon)dy \\ = & \frac{-\tilde{\lambda}_n\tilde{\lambda}_i}{(\tilde{\lambda}_n \exp(\epsilon) + \tilde{\lambda}_i)^2} \left[\frac{\tilde{\lambda}_n \exp(y+\epsilon) + \tilde{\lambda}_i \exp(y) + \exp(\epsilon)}{\exp\left\{\exp(y)(\tilde{\lambda}_n + \tilde{\lambda}_i \exp(-\epsilon))\right\}} \right]_{y=-\infty}^{y=+\infty} \end{aligned} \tag{b.17}$$

Let $v(y)$ be the expression in the bracket.

$$\lim_{y\to-\infty} v(y) = \frac{0+0+\exp(\epsilon)}{\exp\{0\}} = \exp(\epsilon) \tag{b.18}$$

For the upper bound, we use l'Hopital rule:

$$\begin{aligned}
\lim_{y\to\infty} v(y) = \lim_{y\to\infty} & \frac{\tilde{\lambda}_n \exp(y+\epsilon) + \tilde{\lambda}_i \exp(y)}{\exp\left\{\exp(y)(\tilde{\lambda}_n + \tilde{\lambda}_i \exp(-\epsilon))\right\} \exp(y)(\tilde{\lambda}_n + \tilde{\lambda}_i \exp(-\epsilon))} \\
= \lim_{y\to\infty} & \frac{\tilde{\lambda}_n \exp(\epsilon) + \tilde{\lambda}_i}{\exp\left\{\exp(y)(\tilde{\lambda}_n + \tilde{\lambda}_i \exp(-\epsilon))\right\} (\tilde{\lambda}_n + \tilde{\lambda}_i \exp(-\epsilon))} \\
= 0 &
\end{aligned} \tag{b.19}$$

Thus, (b.17) becomes:

$$f(\epsilon) = \frac{\tilde{\lambda}_n \tilde{\lambda}_i \exp(\epsilon)}{(\tilde{\lambda}_n \exp(\epsilon) + \tilde{\lambda}_i)^2} \tag{b.20}$$

The corresponding cdf is:

$$F(\epsilon) = 1 - \frac{\tilde{\lambda}_i}{\tilde{\lambda}_n \exp(\epsilon) + \tilde{\lambda}_i} \tag{b.21}$$

Given that ϵ is bounded, we can compute the truncated pdf as:

$$\begin{aligned}
f_T(\epsilon) =& \frac{f(\epsilon)}{F(\theta \log(\tau_{ni})) - F(-\theta \log(\tau_{in}))} \\
=& \gamma^{-1} \frac{\tilde{\lambda}_n \tilde{\lambda}_i \exp(\epsilon)}{(\tilde{\lambda}_n \exp(\epsilon) + \tilde{\lambda}_i)^2},
\end{aligned} \tag{b.22}$$

where:

$$\gamma = -\frac{\tilde{\lambda}_i}{\tilde{\lambda}_n \exp(\theta \log(\tau_{ni})) + \tilde{\lambda}_i} + \frac{\tilde{\lambda}_i}{\tilde{\lambda}_n \exp(-\theta \log(\tau_{in})) + \tilde{\lambda}_i} \tag{b.23}$$

Similarly, the truncated cdf is:

$$F_T(\epsilon) = \gamma^{-1} \int_{-\theta \log(\tau_{in})}^{\epsilon} f(t)dt \tag{b.24}$$

Now that we have these distributions, we compute order statistics from them, which allow us to characterize the trade barriers estimated from price data. We use the following result: Given L observations drawn from pdf $h(x)$, the pdf of the r-th order statistic (where $r = L$ is the max and $r = 1$ is the min) is:

$$h_r(x) = \frac{L!}{(r-1)!(L-r)!} H(x)^{r-1}(1 - H(x))^{L-r} h(x) \tag{b.25}$$

The pdf of the max reduces to:

$$h_{\max}(x, L) = L H(x)^{L-1} h(x)$$

With this pdf defined, we can compute the expectation of the maximum statistic:

$$E[\max_{z \in L}(x_z)] = \int_{-\infty}^{\infty} x h_{\max}(x, L) dx \tag{b.26}$$

Recall that we are interested in computing the expectation of the maximum logged price difference between countries n and i. But, so far, we have derived the truncated pdf and cdf of $\epsilon = \theta \log(v_{ni}(z))$. Our object of interest is actually $\log(v_{ni}(z)) = \frac{1}{\theta}\epsilon$. The expectation of this object, which represents the maximum log price difference, for L draws, is given by:

$$E[\max_{z \in L}(\log(p_n(z)) - \log(p_i(z)))] = \frac{1}{\theta} \int_{-\theta \log(\tau_{in})}^{\theta \log(\tau_{ni})} \epsilon f_{\max}(\epsilon, L) d\epsilon, \tag{b.27}$$

where:

$$\begin{aligned} f_{\max}(\epsilon, L) =& L F_T(\epsilon)^{L-1} f_T(\epsilon) \\ =& L \left[\gamma^{-1} \int_{-\theta \log(\tau_{in})}^{\epsilon} f(t) dt \right]^{L-1} \gamma^{-1} \frac{\tilde{\lambda}_n \tilde{\lambda}_i \exp(\epsilon)}{(\tilde{\lambda}_n \exp(\epsilon) + \tilde{\lambda}_i)^2} \end{aligned} \tag{b.28}$$

Hence, the expectation of the maximum of the log price difference is proportional to $1/\theta$, where the proportionality object comes from gravity,

$$E[\max_{z \in L}(\log(p_n(z)) - \log(p_i(z)))] = \Psi_{ni}(L; \mathbf{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{n}}), \tag{b.29}$$

where:

$$\Psi_{ni}(L; \mathbf{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{n}}) \equiv \frac{1}{\theta} \int_{-\theta \log(\tau_{in})}^{\theta \log(\tau_{ni})} \epsilon f_{\max}(\epsilon, L) d\epsilon, \tag{b.30}$$

and the values $\mathbf{S}$ and $\tilde{\boldsymbol{\tau}}_{\boldsymbol{n}}$ correspond with the definitions outlined in Definition 1. It is worth emphasizing the nature of this integral: Other than the scalar in the front, it depends completely on objects that can be recovered from the standard gravity equation in (21).

Finally, one can rewrite equation (b.30) via integration by parts as:

$$E[\max_{z\in L}(\log(p_n(z)) - \log(p_i(z)))] = \log \tau_{ni} - \frac{1}{\theta}\int_{-\theta\log(\tau_{in})}^{\theta\log(\tau_{ni})} F_{\max}(\epsilon, L)d\epsilon \tag{b.31}$$

$$\log \tau_{ni} = E[\max_{z\in L}(\log(p_n(z)) - \log(p_i(z)))] + \frac{1}{\theta}\int_{-\theta\log(\tau_{in})}^{\theta\log(\tau_{ni})} F_{\max}(\epsilon, L)d\epsilon, \tag{b.32}$$

which implies the following strict inequality:

$$\log \tau_{ni} > E[\max_{z\in L}(\log(p_n(z)) - \log(p_i(z)))] = \Psi_{ni}(L; \mathbf{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{n}}), \tag{b.33}$$

where the strict inequality simply follows from the inspection of the CDF $F_{\max}(\epsilon, L)$ which has positive mass below the point $\theta\log(\tau_{ni})$. This then proves claim 1. in Lemma 1.

To prove claim 2. in Lemma 1, we compute the difference in the expected values of log prices between two countries. We show that they are equal to the (scaled) difference in the price parameters Φ.

Rather than working with the distribution described above, it is more convenient to directly compute the expectation of log prices using the equilibrium price distribution. Note that EK show that the cdf and pdf of prices in country i are $G(p) = 1 - \exp(-\Phi_i p^\theta)$ and $g(p) = p^{\theta-1}\theta\Phi_i\exp(-\Phi_i p^\theta)$, respectively.

For any country i, define the expectation of logged prices as

$$E[\log(p_i(z))] = \int_0^\infty \log(p)g(p)dp \tag{b.34}$$

Substituting the pdf of prices and then utilizing some algebra to find an appropriate change in variables, expression (b.34) yields

$$\begin{aligned} E[\log(p_i(z))] &= \int_0^\infty \log(p)p^{\theta-1}\theta\Phi_i\exp(-\Phi_i p^\theta)dp \\ &= \int_0^\infty \log(p)\theta\Phi_i\exp(\theta\log(p))\exp(-\Phi_i\exp(\theta\log(p)))\frac{dp}{p} \end{aligned}$$

Our change of variables will set $x = \log(p)$, which yields $dx/dp = 1/p$. Then, integration by change of variables allows us to rewrite the above as

$$\begin{aligned} E[\log(p_i(z))] &= \int_0^\infty \log(p)\theta\Phi_i \exp(\theta\log(p))\exp(-\Phi_i\exp(\theta\log(p)))\frac{dp}{p} \\ &= \int_0^\infty x\theta\Phi_i \exp(\theta x)\exp(-\Phi_i\exp(\theta x))\frac{\theta}{\theta}dx \end{aligned}$$

Let $y = \theta x$, so that $dy/dx = \theta$; then, another change of variables gives

$$E[\log(p_i(z))] = \frac{1}{\theta}\int_0^\infty y\Phi_i \exp(y)\exp(-\Phi_i\exp(y))dy$$

Let $t = \Phi_i \exp(y)$, so that $dt/dy = \Phi_i \exp(y)$ and $y = \log(t/\Phi_i)$. Then,

$$\begin{aligned} E[\log(p_i(z))] &= \frac{1}{\theta}\int_0^\infty \log\left(\frac{t}{\Phi_i}\right)\exp(-t)dt \\ &= \frac{1}{\theta}\left\{\int_0^\infty \log(t)\exp(-t)dt - \int_0^\infty \log(\Phi_i)\exp(-t)dt\right\} \\ &= -\frac{1}{\theta}\left\{\gamma + \log(\Phi_i)\right\}, \end{aligned} \tag{b.35}$$

where γ is the Euler-Mascheroni constant. Finally, using (b.35) and taking the expected difference in log prices between country n and country i, the scaled Euler-Mascheroni constant cancels between the two countries and leaves the following expression

$$\begin{aligned} E[\log(p_n(z))] - E[\log(p_i(z))] \qquad &= -\frac{1}{\theta}\left\{\log(\Phi_n) - \log(\Phi_i)\right\} \\ &\equiv \Omega_{ni}(\mathbf{S}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{n}}, \tilde{\boldsymbol{\tau}}_{\boldsymbol{i}}), \end{aligned} \tag{b.36}$$

which then proves claim 2. in Lemma 1.

2.1.C. Proof of Lemma 2 and Proposition 1

To prove Lemma 2 and Proposition 1, we invert EK's estimator for the elasticity of trade:

$$\frac{1}{\hat{\beta}} = -\frac{\sum_n \sum_i \left(\log\hat{\tau}_{ni} + \log\hat{P}_i - \log\hat{P}_n\right)}{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)} \tag{b.37}$$

Given the assumption that the trade data are fixed, equation (b.37) is linear in the random variables $\log\hat{\tau}_{ni}$ and $(\log\hat{P}_n - \log\hat{P}_i)$. With this observation, taking expectations of these random variables yields

$$E\left(\frac{1}{\hat{\beta}}\right) = \frac{1}{\theta}\left\{-\frac{\sum_n \sum_i (\theta\Psi_{ni}(L) - (\log\Phi_i - \log\Phi_n))}{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}\right\} < \frac{1}{\theta}, \tag{b.38}$$

by substituting in for the expectation of the maximum log price difference using (b.30), and the difference in expectations of log prices using (b.36). Inspection of the bracketed term above implies that the following strict inequality must hold,

$$1 \; > \; \left\{-\frac{\sum_n \sum_i (\theta\Psi_{ni}(L) - (\log\Phi_i - \log\Phi_n))}{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}\right\} \; > \; 0, \tag{b.39}$$

with the reason being that $\Psi_{ni}(L) < \log\tau_{ni}$ from Lemma 1; otherwise, the bracketed term would correspond exactly with equation (5) in logs and, thus, equal one. Now, inverting the expression above and applying Jensen's inequality results in the following:

$$E(\hat{\beta}) \; \geq \; \left[E\left(\frac{1}{\hat{\beta}}\right)\right]^{-1} = \theta\left\{-\frac{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}{\sum_n \sum_i (\theta\Psi_{ni}(L) - (\log\Phi_i - \log\Phi_n))}\right\} > \theta, \tag{b.40}$$

with the strict inequality following from (b.38) and (b.39). This proves Proposition 1.

2.2. Proof of Proposition 2

In this subsection, we prove Proposition 2. To prove the claims in Proposition 2, we start with claim 1.

To prove claim 1., we argue that the sample maximum of scaled log price differences is a consistent estimator of the scaled trade cost. In particular, we argue that as the sample size becomes infinite, the probability that the sample scaled trade cost is arbitrarily close to the true scaled trade cost is one.

To see this, consider an estimate of the scaled trade barrier, given a sample of L goods' prices,

$$\theta\log\hat{\tau}^L_{ni} = \theta\left\{\max_{\ell=1,\ldots,L}(\log p_n(\ell) - \log p_i(\ell))\right\}. \tag{b.41}$$

The cdf of this random variable is the integral of its pdf, which is given in expression (b.28), over the compact interval in which the scaled logged price difference lies, $[-\theta\log\tau_{in}, \theta\log\tau_{ni}]$. Denote this cdf by F^L_{max}. From (b.28), $F^L_{max} \equiv (F_T)^L$, where F_T is the truncated distribution of the scaled log price difference over the domain $[-\theta\log\tau_{in}, \theta\log\tau_{ni}]$. By definition, F_T and F^L_{max} take on values between zero and one, as they are cdfs. In particular, for any realization $x < \theta\log\tau_{ni}$,

$F_T(x) < 1$. For any $L > 1$, $F_{max}^L(x) = (F_T(x))^L \leq F_T(x) < 1$.

Take $L \to \infty$. Then, for any $x \in [-\theta \log \tau_{in}, \theta \log \tau_{ni})$, $F_{max}^L = (F_T)^L$ becomes arbitrarily close to zero since $F_T < 1$. Hence, all the mass of the cdf F_{max}^L becomes concentrated at $\theta \log \tau_{ni}$. Thus, as the sample size becomes infinite, the estimated scaled trade barrier converges to the true scaled trade barrier, in probability. Rescaling everything by $\frac{1}{\theta}$ then implies

$$\operatorname*{plim}_{L\to\infty} \log \hat{\tau}_{ni}^L = \operatorname*{plim}_{L\to\infty} \max_{\ell=1,\ldots,L} (\log p_n(\ell) - \log p_i(\ell)) = \log \tau_{ni}. \tag{b.42}$$

This proves claim 1. of Proposition 2.

To show consistency of the estimator $\hat{\beta}$, we argue that

$$\operatorname*{plim}_{L\to\infty} \hat{\beta}\left(L; \mathbf{S}, \tilde{\boldsymbol{\tau}}, \mathbb{X}\right) = \theta, \tag{b.43}$$

or, equivalently, that $\forall \delta > 0$,

$$\lim_{L\to\infty} \Pr\left[\| \hat{\beta}\left(L; \mathbf{S}, \tilde{\boldsymbol{\tau}}, \mathbb{X}\right) - \theta \| < \delta\right] = 1. \tag{b.44}$$

Basically, we will argue that, by sampling the prices of an ever-increasing set of goods and applying the estimator β over these prices, with probability one, we will obtain estimates that are arbitrarily close to θ.

Inverting the expression for the estimator $\hat{\beta}$ in expression (12), rearranging, and multiplying and dividing by the scalar θ yields

$$\frac{1}{\hat{\beta}} = \frac{1}{\theta} \frac{\sum_n \sum_i \left(\theta \log \hat{\tau}_{ni}^L - \theta[\log \hat{P}_n - \log \hat{P}_i]\right)}{-\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}. \tag{b.45}$$

By assumption, the denominator is trade data and is not a random variable.

In the numerator, $\log \hat{P}_n - \log \hat{P}_i$ is the difference in the average of logged prices for countries n and i, given a sample of L goods. In particular,

$$\log \hat{P}_n - \log \hat{P}_i \equiv \frac{1}{L} \sum_{\ell=1}^{L} \log p_n(\ell) - \frac{1}{L} \sum_{\ell=1}^{L} \log p_i(\ell) \tag{b.46}$$

We refer the reader to Davidson and MacKinnon (2004) for a proof of the well known result that the sample average is both an unbiased and consistent estimator of the mean. Since the difference operator is continuous, the difference in the sample average of logged price is an unbiased and consistent estimator of the difference in mean logged prices. Finally, multiplying these sample averages by a scalar θ, a continuous operation, ensures convergence to true difference in the price terms Φ.

We have argued that the two components in the numerator converge in probability to their true parameter counterparts, as the sample size becomes infinite. Taking the difference of these two components, summing over all country pairs (n, i), and dividing by the scalar $\left[-\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)\right]$, all of which are continuous operations, allows us to conclude that

$$\operatorname*{plim}_{L\to\infty} \frac{\sum_n \sum_i \left(\theta \log \hat{\tau}_{ni}^L - \theta[\log \hat{P}_n - \log \hat{P}_i]\right)}{-\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}$$

$$= \operatorname*{plim}_{L\to\infty} \frac{\sum_n \sum_i \left(\theta \max_{\ell=1,\dots,L}(\log p_n(\ell) - \log p_i(\ell)) - \theta\left[\frac{1}{L}\sum_{\ell=1}^L \log p_n(\ell) - \frac{1}{L}\sum_{\ell=1}^L \log p_i(\ell)\right]\right)}{-\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}$$

$$= \frac{-\sum_n \sum_i (\theta \log \tau_{ni} - [\log \Phi_i - \log \Phi_n])}{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}. \tag{b.47}$$

To complete the argument, consider the log of expression (5), which involves Φ. Summing this expression over all (n, i) country pairs gives:

$$\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right) = -\sum_n \sum_i (\theta \log \tau_{ni} - [\log \Phi_i - \log \Phi_n]). \tag{b.48}$$

Substituting expression (b.48) in the denominator of (b.47) above makes the fraction in that expression equal to unity. Hence, $1/\hat{\beta}$ converges to $1/\theta$ in probability. Since, for $\beta \in (0, \infty)$, $1/\hat{\beta}$ is a continuous function of $\hat{\beta}$, $\hat{\beta}$ converges to θ in probability. This proves claim 2. of Proposition 2.

Claim 3. of Proposition 2 follows from the fact that $\hat{\beta}$ is a consistent estimator of θ (see Hayashi (2000) for a discussion).

2.3. Deriving the Inverse Marginal Cost Distribution

To simulate the model, we argue that by using the coefficients **S** estimated from the gravity regression (21), we have enough information to simulate prices and trade flows. The key insight is that the S's are sufficient to characterize the inverse marginal cost distribution. Thus, we can sample from this distribution and then compute equilibrium prices and trade flows.

To see this argument, let $z_i \sim F_i(z_i) = \exp(-T_i z_i^{-\theta})$ and define $u_i \equiv z_i/w_i$. The pdf of z_i is

$f_i(z_i) = \exp(-T_i z_i^{-\theta})\theta T_i z_i^{-\theta-1}$. To find the pdf of the transformation u_i, $m_i(u_i)$, use the fact that $f_i(z_i)dz_i = m_i(u_i)du_i$, or $m_i(u_i) = f_i(z_i)(du_i/dz_i)^{-1}$. Let $\tilde{S}_i = T_i w_i^{-\theta}$. Using $f_i(z_i)$, $\tilde{S}_i$, and the fact that $du_i/dz_i = 1/w_i$, we obtain:

$$
\begin{aligned}
m_i(u_i) = f_i(z_i)\left(\frac{du_i}{dz_i}\right)^{-1} &= \exp(-T_i z_i^{-\theta})\theta T_i z_i^{-\theta-1}\left(\frac{1}{w_i}\right)^{-1} \\
&= \exp\left(-T_i z_i^{-\theta}\frac{w_i^{-\theta}}{w_i^{-\theta}}\right)\theta T_i z_i^{-\theta-1}\left(\frac{1}{w_i}\right)^{-1}\frac{w_i^{-\theta}}{w_i^{-\theta}} \\
&= \exp\left(-\tilde{S}_i\frac{z_i^{-\theta}}{w_i^{-\theta}}\right)\theta\tilde{S}_i\frac{z_i^{-\theta-1}}{w_i^{-\theta-1}} \\
&= \exp\left(-\tilde{S}_i u_i^{-\theta}\right)\theta\tilde{S}_i u_i^{-\theta-1}
\end{aligned}
$$

Clearly $m_i(u_i)$ is the pdf that corresponds to the cdf $M_i(u_i) = \exp(-\tilde{S}_i u_i^{-\theta})$, which concludes the argument.

C. EK's Alternative Estimators of θ

EK use two other alternative methods to estimate θ than the one described in the main body of the paper. Through these alternative methods they are able to establish a range from 3.6 to 12.86. In this section, we explore the properties of one of these alternative estimators. We show that the estimator associated with the estimate of 12.86 is biased by economically meaningful magnitudes for the same reasons as the estimator discussed in the paper. Similar to our earlier arguments, we then use the moments associated with the biased estimator as the basis for our estimation. Doing so allows us to establish a range from 3.6 to 4.3 with EK's data rather than the range between 3.6 and 12.86.

Before proceeding, we should note that we have little to say about EK's estimation approach that leads to an estimate of 3.6. To arrive at this estimate, they use wage data and proxies for the productivity parameters, T, and find a value of 3.6.[23] While one may have objections to the particular statistics that they employ, the resulting estimate is in line with the estimates that we obtain, which we view as reassuring.

[23] Similarly, Costinot, Donaldson, and Komunjer (2011) estimate θ using trade data and proxies for productivity at the industry level for 21 developed countries. They provide a wide range of estimates depending on the specification, with a preferred estimate of 6.53.

3.1. EK's 2SLS Approach

EK propose an alternative estimator for θ that uses the same variation in price data discussed in the text. First, they use the object D_{ni} defined as

$$D_{ni} = \log\left(\frac{\hat{P}_i \hat{\tau}_{ni}}{\hat{P}_n}\right) \tag{b.49}$$

$$\text{where} \quad \log \hat{\tau}_{ni}(L) = \max_{\ell \in L} \left\{ \log p_n(\ell) - \log p_i(\ell) \right\},$$

$$\text{and} \quad \log \hat{P}_i = \frac{1}{L} \sum_{\ell=1}^{L} \log(p_i(\ell)),$$

to proxy trade costs in the gravity equation (21). By using this measure in the gravity equation (rather than using distance and fixed effects), they can then interpret the coefficient on D_{ni} as an estimate of the trade elasticity.

When using (21) and (b.49) to approximate trade costs, EK are concerned about measurement error, so they employ instrumental variables to alleviate this concern. Specifically, they use the geography variables (distance, border, language) in (22) as instruments for D_{ni}. The resulting two stage least squares (2SLS) estimate of θ is 12.86.

3.2. A Monte Carlo Study of EK's 2SLS Approach

Here we apply the same experiment described in Section 4: we simulate trade flows and samples of micro-level prices under a known θ. Then, we apply EK's 2SLS estimator to the artificial data. We employ the same simulation procedure described in Steps 1-3 in Section 5.2 and we estimate all parameters (except for θ) using the trade data from EK. We set the true value of θ equal to 8.28. The sample size of prices is set to $L = 50$, which is the number of prices EK had access to in their data set.

Table 9 presents the results. The first row shows that the estimates using EK's 2SLS approach are almost 100 percent larger than the true θ of 8.28. Comparing theses results with the study of the method of moment estimator in Table 1, the bias takes the same form, but in the 2SLS approach the bias is significantly larger (12.5 vs. 15.9). This suggests (and our estimation below confirms) that any difference between EK's original results using method of moments vs. 2SLS arises because of how the particular estimator interacts with the bias in the approximation of the trade friction.

The next three rows show how these results change as the number of sampled prices increases. Here increasing the sample size systematically reduces the bias similar to the method of mo-

Table 9: Monte Carlo Results, EK's 2SLS Approach, True θ = 8.28

Approach = 2SLS, Gravity	Mean Estimate of θ (S.E.)	Median Estimate of θ
50 sampled prices	15.9 (0.24)	15.6
500 sampled prices	10.5 (0.05)	10.4
5,000 sampled prices	8.72 (0.02)	8.73
50,000 sampled prices	8.33 (0.01)	8.33

Note: S.E. is the standard error of the mean. In each simulation there are 19 countries and 500,000 goods. 100 simulations performed.

ment results in Table 2. This shows that the key problem with EK's approach is not the estimator per se, but, instead, the poor approximation of the trade costs. Once the sample size of prices becomes large enough, trade costs are better approximated and the bias in the estimate of θ is reduced.

Recall that the purpose of EK's 2SLS estimator was to alleviate an error-in-variables problem. However, 2SLS only works if the error-in-variables problem is classical in the sense that the measurement error is mean zero. The issue identified in this paper is a situation where the measurement error is not classical. The approximated trade friction always underestimates the true trade friction and the approximation error is never mean zero, thus it is not obvious that 2SLS corrects the problem. In fact, the results in Table 9 suggest that 2SLS makes the bias in θ worse when compared to alternative estimators.

3.3. Using EK's 2SLS Estimates as a Basis For Estimation

The estimates from EK's 2SLS approach can be used as the basis for our estimation rather than the estimates from EK's method of moments approach. Specifically, in the exactly identified case, we compare the empirical moment from EK's 2SLS estimation to the averaged simulation moment, which yields the following zero function:

$$y(\theta) = \left[\beta_{2SLS} - \frac{1}{S}\sum_{s=1}^{S}\beta_{2SLS}(\theta, u_s)\right]. \tag{b.50}$$

Our estimation procedure is based on the same moment condition described in the main text:

$$E\left[y(\theta_o)\right] = 0,$$

where θ_o is the true value of θ. Thus, our simulated method of moments estimator is

$$\hat{\theta} = \arg\min_{\theta} \left[y(\theta)' y(\theta) \right], \tag{b.51}$$

where we abstract from the weighting matrix since we focus on the exactly identified case in this section.

Table 10: Estimation Results: 2SLS Moments, EK Data

	Estimate of θ (S.E.)	β_{2SLS}
Data Moments	—	8.03
2SLS Moments, Exactly Identified	4.39 (0.86)	8.03

Table 10 presents the result using EK's data. The first row presents the data moments. Here the estimate of β_{2SLS} is 8.03. This differs from EK's number of 12.86 only because we are using the maximum price difference rather than the second order statistic used in EK. The second row presents the estimate of θ which is 4.39, the standard error, and the model-implied moment.

Note that, while a very different moment is the basis of our estimation, the estimate is nearly identical to the exactly identified results in Table 6, i.e. 4.39 vs. 4.42. On its own, this is a reassuring result because it shows that alternative moments are giving similar answers. Moreover, it suggests that any difference between the results using method of moments vs. 2SLS in EK arises primarily because of how the particular estimator interacts with the bias in the approximation of the trade friction. Yet once this bias is corrected for, we find similar results independent of the particular moment used.

D. Feenstra's 1994 Methodology in the Ricardian Model

In this section we analyze Feenstra's (1994) method to estimate the elasticity of substitution from cross-country data in the context of the Ricardian model. We show that Feenstra's (1994) method recovers the *elasticity of substitution across goods*, i.e. the ρ parameter in CES preferences. *It does not recover the θ parameter controlling the trade elasticity*, i.e. how trade flows change in response to changes in trade costs and the welfare gains from trade. Thus, using the estimates from Feenstra (1994) or Broda and Weinstein (2006) to calibrate the θ parameter in the Ricardian model is inappropriate.

We show this result by asking the following question: given prices and shares generated from the Ricardian model, what would Feenstra's (1994) method recover — the θ or the ρ? To answer this question we will briefly describe Feenstra's (1994) method and its application to simulated data from the Ricardian model. In the description, we will mainly follow Feenstra (2010).

First, we will define an individual variety in Feenstra's (1994) language as a specific good j on the zero-one interval. In the Ricardian model, the expenditure share for good j in country k at date t is given by the following formula:

$$s(j)_{kt} = \frac{p(j)_{kt}^{1-\rho}}{\left\{ \int_0^1 p(\ell)_{kt}^{\frac{1}{1-\rho}} d\ell \right\}^{1-\rho}} \tag{b.52}$$

which is the standard formula for expenditure shares from CES demand structures. Recall that the prices $p(j)$ and $p(\ell)$ are optimal, i.e. they correspond to the lowest cost producer. Aggregate expenditure shares in (4) come from integrating (b.52) over country pair combinations.

Taking logs, differencing, and expressing the denominator (b.52) as a time fixed effect yields

$$\Delta \log s(j)_{kt} = \phi_t - (\rho - 1)\Delta \log p(j)_{kt} + \epsilon(j)_{kt}. \tag{b.53}$$

The final term $\epsilon(j)_{kt}$ represents both trade cost shocks and productivity shocks that will generate variation in shares and prices across time/simulations. Equation (b.53) is the same equation that Feenstra's (1994) methodology exploits.

Feenstra (1994) introduces an upward sloping log-linear supply curve into the estimation of (b.53). Define the "reduced form" supply elasticity as η.[24] By differencing the supply and demand equations with respect to a reference county i and then multiplying these equations together (see Feenstra (2010)) he arrives at the following equations:

$$Y_{kt} = \theta_1 X_{kt} + \theta_2 X_{2kt} + u_{kt}, \tag{b.54}$$

[24] This raises an often overlooked conceptual issue as well. Modulo general equilibrium effects, the Ricardian model of Eaton and Kortum (2002) or more generally the models of Krugman (1980) or Melitz (2003) do not have upward sloping supply curves. In the Ricardian model this is clear. With constant returns to scale technologies and perfect competition, firms price at marginal cost whether the quantity supplied is negligible or infinite. This is problematic because researchers typically take the estimated demand elasticity (that has an associated supply elasticity) and use these estimates in models that do not have an upward-sloping supply curve. Thus, it is not clear whether the change of quantities in response to a change in trade frictions implied by the estimation is the same as that implied by the model.

where

$$Y_{kt} = (\Delta \log p(j)_{kt} - \Delta \log p(j)_{it})^2, \tag{b.55}$$

$$X_{1kt} = (\Delta \log s(j)_{kt} - \Delta \log s(j)_{it})^2, \tag{b.56}$$

$$X_{1kt} = (\Delta \log p(j)_{kt} - \Delta \log p(j)_{it})(\Delta \log s(j)_{kt} - \Delta \log s(j)_{it}), \tag{b.57}$$

$$\theta_1 = \frac{\nu}{(\rho-1)^2(1-\nu)}, \quad \theta_2 = \frac{2\nu - 1}{(\rho-1)^2(1-\nu)} \tag{b.58}$$

u_{kt} is an error term composed of the shocks to the demand curve and the supply curve. Then averaging these equations across time yields:

$$\bar{Y}_k \;\; = \;\; \theta_1 \bar{X}_{1k} + \theta_2 \bar{X}_{2k} + \bar{u}_k. \tag{b.59}$$

Equation (b.59) relates second moments of price and share changes that linearly depend on demand and supply elasticities. Given the appropriate assumptions on the variances of the error terms across countries and across demand and supply shocks, least squares estimates of (b.59) are consistent. Finally, given the estimates of θ_1 and θ_2 one can recover the demand and supply elasticity by using (b.58).

There is an important point to note here. First — and this should be clear from equations (b.58) and (b.59) — Feenstra's (1994) method can only speak to and recover the parameter ρ, which our Monte-Carlo experiment confirms below. This is an important observation because the parameter ρ does not affect aggregate trade flows or measures of the welfare gains from trade in the Ricardian model.[25]

4.1. Monte-Carlo Study of Feenstra's Method in the Ricardian Model

To further illustrate what Feenstra's (1994) method recovers, we performed the following exercise: We simulated prices and expenditure shares for individual varieties from the Ricardian model when calibrated as in Section 4. To generate time series variation we introduced trade cost shocks, cost shocks, and measurement error in the prices. These shocks are independent across time and countries. All the shocks are multiplicative and log normally distributed with the mean of the associated normal distribution set equal to zero and a standard deviation parameter picked by us.

[25] We suspect that a similar result can be derived for the Melitz (2003) model as articulated in Chaney (2008) because the aggregate trade elasticity there relates to the underlying shape parameter of the Pareto distribution of firm productivity.

Table 11: Estimates of Demand Elasticity, Feenstra's Method

	Mean Estimate	Median Estimate
Model, $\theta = 4, \rho = 1.5$	1.51 (0.001)	1.51
Model, $\theta = 4, \rho = 2.5$	2.52 (0.003)	2.52
Model, $\theta = 8, \rho = 1.5$	1.51 (0.001)	1.51
Model, $\theta = 8, \rho = 2.5$	2.51 (0.004)	2.51

Note: In the simulation there are 19 countries with trade frictions and productivity parameters calibrated to fit Eaton and Kortum's (2002) data. 29 periods of data were generated and used, which is consistent with the time series in Broda and Weinstein (2006). Means and medians are over 100 simulations.

Given a sequence of prices and shares as described above we apply Feenstra's (1994) method. Mechanically we implement Feenstra's (1994) method by estimating (b.59) by least-squares while constraining $\theta_1 > 0$. This constraint ensures that the recovered demand elasticity is a real number.

Table 11 presents the results for different θ's and ρ's. In all cases, the mean and median elasticity correspond essentially with the ρ parameter in the calibrated model. In no case does Feenstra's (1994) method correspond with the θ parameter in the model. Thus, Feenstra's (1994) method can only speak to and recover the parameter ρ.

As mentioned earlier, this observation is important because the parameter ρ does not affect aggregate trade flows or measures of the welfare gains of trade in the Ricardian model. Moreover, this shows that using the estimates from Broda and Weinstein (2006) obtained by following Feenstra's (1994) method to calibrate the θ parameter in the Ricardian model is inappropriate.

Table 12: 2005 ICP Data, Step 1 Country-Specific Estimates

Country	$\hat{S}_i$	S.E.	ex_i	S.E.	Country	$\hat{S}_i$	S.E.	ex_i	S.E.	Country	$\hat{S}_i$	S.E.	ex_i	S.E.
Angola	−1.04	0.21	−2.67	0.35	Fiji	−0.58	0.20	−2.06	0.31	Nepal	0.48	0.24	−3.00	0.32
Argentina	1.13	0.18	2.34	0.25	Finland	1.09	0.17	2.15	0.23	New Zealand	−0.25	0.30	3.17	0.24
Armenia	0.83	0.20	−3.91	0.29	France	0.39	0.16	5.09	0.22	Nigeria	−0.85	0.25	−1.00	0.29
Australia	0.24	0.16	3.59	0.23	Gabon	−1.07	0.18	−1.52	0.27	Norway	0.33	0.37	1.88	0.23
Austria	0.39	0.16	2.71	0.22	Gambia, The	−2.40	0.22	−2.32	0.34	Oman	−0.19	0.36	−0.74	0.26
Azerbaijan	−0.03	0.20	−2.76	0.28	Georgia	−2.78	0.19	0.70	0.27	Pakistan	0.55	0.29	2.03	0.23
Bangladesh	0.76	0.18	0.46	0.24	Germany	0.40	0.16	5.57	0.22	Paraguay	0.04	0.36	−0.74	0.28
Belarus	1.27	0.18	−0.98	0.25	Ghana	−1.32	0.21	0.44	0.29	Peru	0.47	0.24	1.10	0.25
Belgium	−2.75	0.16	8.26	0.22	Greece	0.78	0.16	0.58	0.23	Philippines	−0.34	0.39	2.64	0.24
Benin	−0.62	0.22	−3.66	0.36	Guinea	−1.76	0.22	−2.16	0.33	Poland	0.84	0.34	1.76	0.23
Bhutan	0.37	0.30	−5.45	0.43	Guinea-Bissau	−0.40	0.28	−5.77	0.48	Portugal	−0.20	0.24	2.71	0.23
Bolivia	0.28	0.19	−1.65	0.29	Hungary	0.86	0.17	0.98	0.23	Romania	0.60	0.25	0.75	0.23
Bosnia and Herzegovina	1.14	0.23	−3.68	0.32	Iceland	−0.26	0.18	−0.55	0.26	Russian Federation	1.32	0.34	2.12	0.23
Botswana	0.97	0.25	−3.73	0.37	India	0.94	0.16	3.53	0.25	Rwanda	0.09	0.27	−5.05	0.36
Brazil	1.30	0.16	3.67	0.23	Indonesia	1.34	0.16	3.07	0.23	Sierra Leone	−0.97	0.25	−3.61	0.41
Brunei Darussalam	1.68	0.25	−5.15	0.37	Iran, Islamic Rep.	1.02	0.21	−0.85	0.28	Saudi Arabia	0.70	0.30	0.70	0.28
Bulgaria	0.30	0.17	0.39	0.24	Ireland	−3.21	0.16	6.39	0.22	Senegal	−0.86	0.27	−0.63	0.25
Burkina Faso	0.32	0.20	−4.07	0.31	Israel	0.59	0.17	1.70	0.24	Slovak Republic	−0.31	0.26	1.34	0.23
Burundi	−1.52	0.20	−3.12	0.34	Italy	0.58	0.16	4.56	0.22	Slovenia	1.02	0.38	−0.20	0.24
Cameroon	1.54	0.21	−3.34	0.30	Japan	1.51	0.16	4.89	0.23	South Africa	0.41	0.25	3.61	0.23
Canada	−0.27	0.16	4.59	0.22	Jordan	−0.25	0.18	−0.65	0.25	Spain	0.29	0.31	4.09	0.22
Cape Verde	−0.37	0.21	−4.86	0.38	Kazakhstan	0.28	0.18	−0.03	0.26	Sri Lanka	−0.14	0.42	0.65	0.25
Central African Republic	0.55	0.25	−4.67	0.36	Kenya	−0.53	0.16	−0.07	0.23	Sudan	−0.12	0.33	−3.47	0.32
Chad	0.54	0.24	−6.49	0.40	Korea, Rep.	1.04	0.16	4.38	0.22	Swaziland	2.10	0.38	−3.30	0.33
Chile	0.27	0.18	1.96	0.25	Kyrgyz Republic	0.03	0.20	−2.86	0.30	Sweden	0.75	0.31	3.34	0.22
China	1.13	0.16	5.74	0.23	Lao PDR	1.43	0.27	−3.92	0.35	Switzerland	0.10	0.25	3.69	0.27
Colombia	0.38	0.17	0.50	0.24	Latvia	−0.46	0.19	−0.10	0.26	Syrian Arab Republic	−0.34	0.31	−0.86	0.26
Comoros	−0.84	0.27	−4.54	0.42	Lebanon	0.60	0.20	−2.29	0.28	Tajikistan	1.10	0.37	−3.19	0.34
Congo, Dem. Rep.	−0.65	0.24	−2.31	0.34	Lesotho	1.09	0.30	−5.44	0.44	Tanzania	−1.01	0.26	−1.41	0.31
Congo, Rep.	−0.95	0.21	−1.08	0.30	Lithuania	0.67	0.21	−0.88	0.29	Thailand	0.86	0.29	3.57	0.28
Cte d'Ivoire	0.78	0.21	−1.22	0.30	Macedonia, FYR	0.41	0.18	−2.71	0.27	Togo	−1.40	0.25	−1.34	0.27
Croatia	1.08	0.16	−1.29	0.24	Malawi	−0.63	0.19	−2.59	0.28	Tunisia	0.34	0.36	−0.30	0.24
Cyprus	−0.86	0.17	0.45	0.24	Malaysia	−1.43	0.16	6.58	0.22	Turkey	0.93	0.28	2.38	0.23
Czech Republic	0.43	0.16	2.02	0.23	Mali	−1.03	0.23	−2.66	0.32	Uganda	−0.71	0.29	−2.30	0.26
Denmark	−0.24	0.16	3.63	0.23	Mauritania	−1.97	0.23	−1.79	0.33	Ukraine	1.41	0.24	0.88	0.28
Djibouti	−2.04	0.24	−2.37	0.38	Mauritius	−1.63	0.17	1.44	0.24	United Kingdom	−0.29	0.32	5.59	0.22
Ecuador	−0.24	0.18	0.12	0.26	Mexico	0.21	0.16	2.61	0.24	United States	0.06	0.34	6.87	0.22
Egypt, Arab Rep.	0.44	0.17	0.62	0.23	Moldova	−0.47	0.19	−2.12	0.29	Uruguay	−0.51	0.29	1.40	0.27
Equatorial Guinea	0.47	0.24	−4.24	0.39	Morocco	−0.39	0.17	1.32	0.23	Venezuela, RB	0.72	0.29	−0.60	0.26
Estonia	−1.74	0.17	1.61	0.24	Mozambique	−0.16	0.22	−2.06	0.33	Vietnam	−0.44	0.24	2.69	0.28
Ethiopia	−0.66	0.21	−2.15	0.31	Namibia	1.09	0.23	−3.64	0.33	Zambia	−3.99	0.30	2.59	0.27

Table 13: EK Data, Step 1 Country-Specific Estimates

Country	$\hat{S}_i$	S.E.	ex_i	S.E.	Country	$\hat{S}_i$	S.E.	ex_i	S.E.
Australia	−0.20	0.15	0.54	0.24	Japan	2.54	0.13	1.74	0.21
Austria	0.50	0.12	−1.65	0.18	Netherlands	−3.09	0.12	0.80	0.18
Belgium	−4.38	0.12	0.98	0.18	New Zealand	−1.42	0.15	0.37	0.24
Canada	−0.46	0.13	1.06	0.22	Norway	−0.34	0.12	−1.01	0.18
Denmark	−1.16	0.12	−0.67	0.18	Portugal	−0.28	0.12	−1.38	0.19
Finland	0.82	0.12	−1.33	0.18	Spain	1.56	0.12	−1.35	0.18
France	1.15	0.12	0.05	0.18	Sweden	0.05	0.12	−0.06	0.18
Germany	1.44	0.12	0.82	0.18	United Kingdom	0.52	0.12	0.89	0.18
Greece	−0.38	0.12	−2.51	0.18	United States	1.34	0.13	2.83	0.22
Italy	1.81	0.12	−0.12	0.18					

Table 14: EIU Data, Step 1 Country-Specific Estimates

Country	$\hat{S}_i$	S.E.	ex_i	S.E.	Country	$\hat{S}_i$	S.E.	ex_i	S.E.	Country	$\hat{S}_i$	S.E.	ex_i	S.E.
Argentina	0.71	0.17	0.74	0.24	Iceland	−0.06	0.17	−3.09	0.24	Poland	0.56	0.16	0.26	0.23
Australia	−0.23	0.16	2.19	0.24	India	0.54	0.16	1.76	0.24	Portugal	−0.03	0.16	0.31	0.23
Austria	0.11	0.16	1.31	0.23	Indonesia	1.01	0.16	1.39	0.23	Romania	0.13	0.16	−0.42	0.23
Azerbaijan	0.06	0.17	−5.28	0.25	Iran, Islamic Rep.	0.75	0.18	−2.71	0.26	Russian Federation	1.17	0.16	0.55	0.24
Belgium	−2.79	0.16	6.16	0.23	Ireland	−3.13	0.16	4.65	0.23	Saudi Arabia	0.22	0.18	−0.59	0.26
Brazil	0.59	0.16	2.33	0.23	Israel	0.08	0.17	0.37	0.24	Senegal	−0.70	0.17	−3.76	0.25
Brunei Darussalam	1.59	0.21	−7.49	0.32	Italy	0.35	0.16	2.93	0.23	Slovak Republic	−0.45	0.16	−0.21	0.23
Bulgaria	0.03	0.17	−1.23	0.24	Japan	1.09	0.16	3.49	0.23	South Africa	0.00	0.16	1.76	0.23
Canada	−0.44	0.16	2.89	0.23	Jordan	−0.81	0.17	−1.88	0.24	Spain	0.03	0.16	2.48	0.23
Central African Republic	0.69	0.21	−7.07	0.31	Kazakhstan	0.55	0.17	−2.45	0.24	Sri Lanka	−0.25	0.17	−1.06	0.24
Chile	−0.04	0.17	0.56	0.24	Kenya	−0.65	0.16	−2.69	0.24	Sweden	0.42	0.16	1.86	0.23
China	0.71	0.16	4.04	0.23	Korea, Rep.	0.58	0.16	3.06	0.23	Switzerland	−0.04	0.18	2.05	0.25
Colombia	0.11	0.16	−1.18	0.24	Malaysia	−1.93	0.16	5.33	0.23	Syrian Arab Republic	−0.41	0.17	−3.04	0.24
Cote d'Ivoire	0.80	0.18	−3.42	0.27	Mexico	−0.23	0.16	1.43	0.24	Thailand	0.51	0.18	1.87	0.25
Czech Republic	0.05	0.16	0.68	0.23	Morocco	−0.47	0.16	−0.65	0.23	Tunisia	0.10	0.16	−2.02	0.24
Denmark	−0.37	0.16	1.74	0.23	Nepal	0.43	0.21	−5.07	0.28	Turkey	0.71	0.16	0.51	0.23
Ecuador	−0.32	0.17	−1.84	0.24	New Zealand	−0.62	0.17	1.75	0.24	Ukraine	1.52	0.18	−1.21	0.26
Egypt, Arab Rep.	0.29	0.16	−1.30	0.23	Nigeria	−1.02	0.18	−2.98	0.26	United Kingdom	−0.38	0.16	3.71	0.23
Ethiopia	−0.68	0.18	−4.15	0.27	Norway	0.35	0.16	0.00	0.23	United States	−0.25	0.16	5.19	0.23
Finland	0.60	0.16	0.93	0.23	Oman	−0.36	0.17	−2.94	0.25	Uruguay	−0.53	0.18	−0.59	0.25
France	0.38	0.16	3.04	0.23	Pakistan	0.30	0.16	0.09	0.23	Venezuela, RB	0.85	0.17	−2.51	0.24
Germany	0.11	0.16	3.95	0.23	Paraguay	−0.03	0.18	−3.03	0.26	Vietnam	−0.37	0.18	0.63	0.26
Greece	0.33	0.16	−0.92	0.23	Peru	0.22	0.17	−0.73	0.24	Zambia	−1.91	0.17	−1.84	0.26
Hungary	0.59	0.16	−0.17	0.23	Philippines	−0.72	0.16	1.50	0.23					

Table 15: Step 1 Trade Cost Estimates and Summary Statistics

Geographic Barriers	**ICP 2005 Data**		**EK Data**		**EIU Data**	
Barrier	Parameter Estimate	S.E.	Parameter Estimate	S.E.	Parameter Estimate	S.E.
$[0, 375)$	-5.30	0.21	-2.89	0.14	-5.02	0.19
$[375, 750)$	-6.29	0.14	-3.56	0.10	-5.28	0.11
$[750, 1500)$	-7.27	0.09	-3.87	0.07	-5.71	0.07
$[1500, 3000)$	-8.50	0.06	-4.10	0.15	-6.63	0.05
$[3000, 6000)$	-9.65	0.04	-6.15	0.09	-7.70	0.04
$[6000, \text{maximum}]$	-10.35	0.05	-6.60	0.10	-8.41	0.04
Shared border	1.25	0.12	0.44	0.14	1.04	0.16

Summary Statistics			
	ICP 2005 Data	**EK Data**	**EIU Data**
No. Obs	10,513	342	4,607
TSS	152,660	2,936	47,110
SSR	30,054	76.56	8,208
σ_ν^2	2.93	0.25	1.84