

The Elasticity of Trade: Estimates and Evidence

Ina Simonovska
University of California, Davis, United States
NBER, United States

Michael E. Waugh
New York University, United States

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 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. 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 2003-2005 ICP round. We thank George Alessandria, Alexander Aue, Dave Donaldson, 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. Ina Simonovska thanks Princeton University for their hospitality and financial support through the Peter B. Kenen fellowship.

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 two examples: First, for any pair of countries, the estimate of the tariff equivalent of a border effect is inversely proportional to the assumed elasticity of trade with respect to trade frictions. Thus, observed reductions in tariffs across countries can explain almost all or none of the growth in world trade, depending on this elasticity. Second, the trade elasticity is one of only two statistics needed to measure the welfare cost of autarky in a large and important class of structural gravity models of international trade. 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 EK's 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 span 123 countries that account for 98 percent of world GDP. Our benchmark estimate for the elasticity of trade is 4.14, rather than approximately eight, as EK's estimation strategy suggests. This difference doubles the measured welfare gains from trade.

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. We show that the reason behind the difference is that their estimator is biased in finite samples

¹The class of models includes Armington as articulated in [Anderson \(1979\)](#), [Krugman \(1980\)](#), [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.

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 2.79 and 4.46. 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 threefold. First, we provide a precise point estimate of the trade elasticity in the context of EK's Ricardian model that doubles the welfare gains from trade predicted by EK's estimation. Since EK's model is a canonical model of international trade and it is widely used in quantitative trade studies, providing a precise point estimate of the trade elasticity in the context of this model is important. Moreover, our findings suggest a range for the trade elasticity of 2.79 to 4.46, which is both lower and narrower relative to EK's estimates of 3.6 to 12.8. In particular, our critique also applies to EK's estimate of 12.8, which was obtained using an alternative approach. After correcting for biases in 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 2.79 to 4.46, relative to EK's wide range of estimates.

Second, we develop a methodology that is applicable to a wide class of trade models. The method and the moments that we use to estimate the trade elasticity within EK's Ricardian framework can be derived for other structural gravity models of trade. In [Simonovska and Waugh \(2013\)](#), we show how the new estimation strategy applies to models with product differentiation such as Armington as articulated in [Anderson \(1979\)](#) and [Krugman \(1980\)](#), 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.

Third, the estimates that we obtain using the newly-developed methodology contribute to a large and important literature that aims to measure the trade elasticity. [Anderson and van Wincoop \(2004\)](#) 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 \(2004\)](#) report is obtained using [Feenstra's \(1994\)](#) method. However, in heterogeneous frameworks with constant-elasticity-of-substitution (CES) preferences, such as EK's Ricardian model, [Feenstra's \(1994\)](#) 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 EK's Ricardian model with micro-level heterogeneity.

Another set of estimates that [Anderson and van Wincoop \(2004\)](#) document relies on time-series and cross-industry variation in tariffs and trade flows during trade liberalization episodes as in [Head and Ries \(2001\)](#) and [Romalis \(2007\)](#), or time-series and cross-country variation in tariffs and trade flows for developed economies during the post-war period as in [Baier and Bergstrand \(2001\)](#). Recently, [Caliendo and Parro \(2012\)](#) build on these approaches and estimate sectoral trade elasticities from cross-sectional variations in trade flows and tariffs. The methods that rely on variations in tariffs and trade flows in order to identify the trade elasticity are applicable to a variety of structural gravity models, including EK's Ricardian model. Hence, the estimates obtained using these methods are comparable to our estimates of the trade elasticity.

Admittedly, there are two outstanding issues in our analysis. First, there is a difference between the low values of the elasticity that our approach yields and the high values typically obtained using tariff data. In particular, [Head and Ries \(2001\)](#), [Romalis \(2007\)](#), and [Baier and Bergstrand \(2001\)](#) find values in the range of five to ten, while our benchmark estimates center around four. The corollary is that the low values of the elasticity we find imply large deviations between observed trade frictions (tariffs, transportation costs, etc.) and those inferred from trade flows.

However, there are two pieces of evidence in support of the values that we find. First, [Parro](#)

(2013) uses the tariff based approach of [Caliendo and Parro \(2012\)](#) to estimate an aggregate trade elasticity for capital goods and non-capital, traded goods. He finds estimates of 4.6 and 5.2 which are only modestly larger than ours. Second, our results compare favorably with alternative estimates of the shape parameter of the productivity distribution, which governs the trade elasticity in models with micro-level heterogeneity, that are not obtained from gravity-based estimators. For example, estimates of the shape parameter from 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. Identification of the parameter in these papers comes from firm-level data, which suggest that there is a lot of variation in firm productivity. The data in our paper are 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.

Second, there are concerns about the quality of the price data that we use in our analysis and we address them to the best of our ability within the scope of the paper. As in EK, we use cross-country micro-level price data from the International Comparison Program (ICP). Obvious concerns with these data are the degree of product comparability (especially across rich and poor countries), aggregation, general measurement error, the role of distribution mark-ups, etc. To make headway, we incorporate these issues into our estimation to determine the direction and quantitative effects that they could have on our estimates. Finally, we provide an additional set of estimates using cross-country price data from the Economist Intelligence Unit (EIU), which suggest a trade elasticity that is even lower than our baseline results.

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:

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

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 :

$$p_n(j) = \min_{k=1, \dots, N} \left\{ p_{nk}(j) \right\}. \quad (2)$$

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, $\gamma = [\Gamma(\frac{\theta+1-\rho}{\theta})]^{-\frac{1}{1-\rho}}$ 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 Armington as articulated in [Anderson \(1979\)](#) and [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 \(2012\)](#) 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:

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

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:

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

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

²This is the most common case, though Donaldson (2012) 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.

country i must satisfy the following inequality:

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

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:

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

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

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

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

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:

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

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

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

³They also propose two other estimators. One uses the approximation in (11) and the gravity equation in (22). 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.

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

Assumption regarding the key source of randomness. Before describing the properties of the estimator $\hat{\beta}$, we state the assumptions that we maintain throughout the theoretical analysis regarding the sources of error in equation (12). Following EK, we assume that trade barriers and price indices are approximated from price data using the last two equations in (12). These two objects are potentially measured with error because of the approximation. Hence, approximation error is the key source of error that we examine in the theoretical analysis. 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 expression (12).

In practice, there may be other sources of error. First, trade shares also appear in equation (12). Throughout the theoretical analysis, we assume that bilateral trade shares are observable statistics that are not measured with error. Therefore, we treat them as constants. We recognize that in practice this may not be the case, so we relax this assumption in the quantitative analysis. Second, prices may be measured with error in the data. Consequently, in the quantitative analysis in Section 7.4, we consider a number of sources of price variation outside of the model. We find that different sources of price variation affect the estimates of the trade elasticity in different directions. Crucially, however, approximation error in trade barriers remains to be the key source of bias in the estimates. Therefore, we turn to the theoretical characterization of the approximation error next.

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

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 $\mathbf{S} = \{\log(T_1 w_1^{-\theta}), \dots, \log(T_N w_N^{-\theta})\}$.

⁴To 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.

3. Let the vector $\tilde{\tau}_i = \{\theta \log(\tau_{i1}), \dots, \theta \log(\tau_{iN})\}$ and let $\tilde{\tau}$ be a matrix with typical row, $\tilde{\tau}_i$.
4. Let $g(p_i; \mathbf{S}, \tilde{\tau}_i)$ be the pdf of prices of individual goods in country i , $p_i \in (0, \infty)$.
5. Let $f_{\max}(\epsilon_{ni}; L, \mathbf{S}, \tilde{\tau}_i, \tilde{\tau}_n)$ 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,*

$$\Psi_{ni}(L; \mathbf{S}, \tilde{\tau}_i, \tilde{\tau}_n) \equiv \frac{1}{\theta} \int_{-\theta \log(\tau_{in})}^{\theta \log(\tau_{ni})} \epsilon_{ni} f_{\max}(\epsilon_{ni}; L, \mathbf{S}, \tilde{\tau}_i, \tilde{\tau}_n) d\epsilon_{ni} < \log(\tau_{ni}). \quad (13)$$

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,

$$\Omega_{ni}(\mathbf{S}, \tilde{\tau}_n, \tilde{\tau}_i) \equiv \int_0^\infty \log(p_n) g(p_n; \mathbf{S}, \tilde{\tau}_n) dp_n - \int_0^\infty \log(p_i) g(p_i; \mathbf{S}, \tilde{\tau}_i) dp_i = \frac{1}{\theta} (\log \Phi_i - \log \Phi_n), \quad (14)$$

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

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

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,*

$$\text{plim}_{L \rightarrow \infty} \max_{\ell=1, \dots, L} (\log p_n(\ell) - \log p_i(\ell)) = \log \tau_{ni}. \quad (18)$$

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

$$\text{plim}_{L \rightarrow \infty} \hat{\beta}(L; \mathbf{S}, \tilde{\tau}, \mathbb{X}) = \theta, \quad (19)$$

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

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

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 is 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 1000 simulations. The rows present two different estimation approaches: method 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

Table 1: Monte Carlo Results, True $\theta = 8.28$

Approach	Mean Estimate of θ (S.E.)	Median Estimate of θ
EK's Estimator	12.5 (0.06)	12.5
Least Squares	12.1 (0.06)	12.1
True Mean $\tau = 1.83$		Estimated Mean $\tau = 1.50$

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 . 10000 simulations performed.

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.⁵ 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 12.1 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 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 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

⁵We 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.

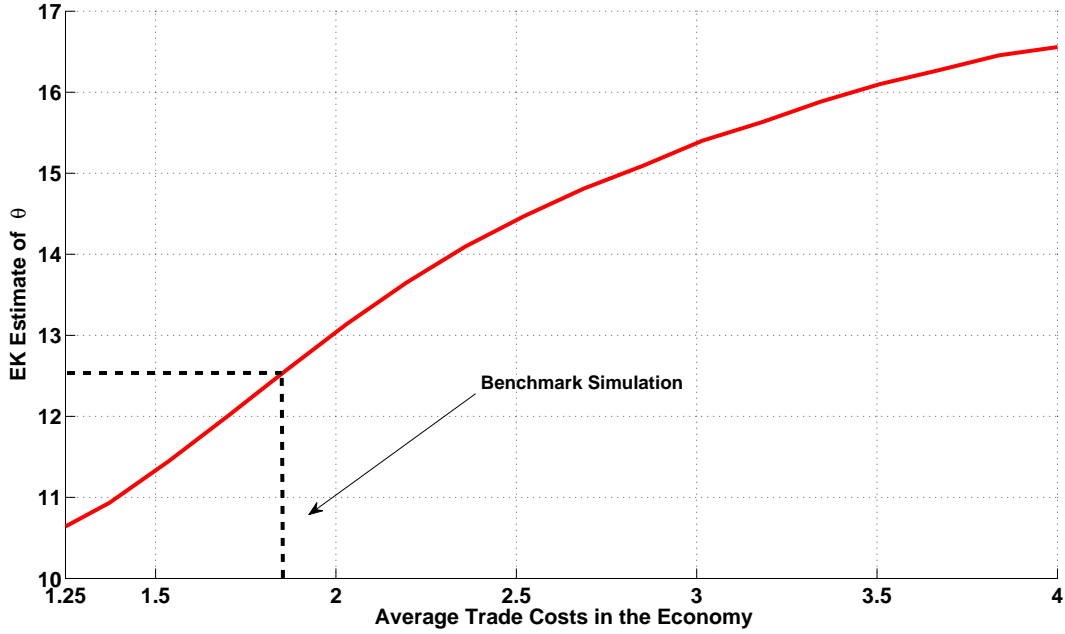


Figure 1: EK’s Estimator and the Level of Trade Costs, True $\theta = 8.28$

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 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. 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.

Table 2: Increasing the Sample of Prices Reduces the Bias, True $\theta = 8.28$

Sample Size of Prices	Mean θ (S.E.)	Median θ	Mean τ
50	12.51 (0.06)	12.51	1.50
500	9.42 (0.02)	9.42	1.69
5,000	8.47 (0.01)	8.47	1.80
50,000	8.30 (0.01)	8.30	1.83

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. 10000 simulations performed. True Mean $\tau = 1.83$.

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 $\tilde{\tau}$. 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) (in Appendix B) 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 $\tilde{\tau}$ 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 .⁶

⁶Lemma 1 established that the expected value of $1/\hat{\beta}$ is proportional to $1/\theta$. Hence, modulo effects from Jensen's




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° 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 .⁷ 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.

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

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

inequality, this suggests that $\hat{\beta}$ is roughly proportional to θ . Figure 2 confirms this.

⁷Another reason for using the moment β is that $\hat{\beta}$ is a consistent estimator of θ , as argued in Proposition 2.

where S_i is defined as $\log [T_i w_i^{-\theta}]$ and is the same value in the parameter vector \mathbf{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 .

The goal is to estimate the objects S_i for all $i = 1, \dots, N$ and $\theta \log \tau_{ni}$ for all country pairs n and i such that $n \neq i$. To do so we first derive an empirical gravity equation that corresponds to the theoretical expression in (21). It is given by

$$\log \left(\frac{X_{ni}/X_n}{X_{nn}/X_n} \right) = \hat{S}_i - \hat{S}_n - \hat{\theta} \log \hat{\tau}_{ni} + \nu_{ni}. \quad (22)$$

\hat{S}_i 's are recovered as the coefficients on country-specific dummy variables given the restrictions on how trade costs can covary across countries. Trade costs take the following functional form

$$\log \hat{\tau}_{ni} = d_k + b_{ni} + ex_i. \quad (23)$$

Here, trade costs are a logarithmic function of distance, where d_k with $k = 1, 2, \dots, 6$ is the effect of distance between country i and n lying in the k -th distance interval.⁸ 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 which allows the trade-cost level to vary depending upon the exporter. This object can be separately identified from \hat{S}_i . The key observation is that the \hat{S}_i 's are a common component to a country that appears in the estimating equation both when the country is exporting *and* importing. The exporter fixed effect appears only when that country is on the exporting side. This distinction then identifies the difference between a high exporting cost versus the \hat{S}_i .

To see the argument formally, combine equation (22) and (23) to obtain

$$\begin{aligned} \log \left(\frac{X_{ni}/X_n}{X_{nn}/X_n} \right) &= \hat{S}_i - \hat{S}_n - \hat{\theta} (d_k + b_{ni} + ex_i) + \nu_{ni} \\ &= \bar{S}_i - \hat{S}_n - \hat{\theta} (d_k + b_{ni}) + \nu_{ni}, \end{aligned}$$

where $\bar{S}_i = \hat{S}_i - \hat{\theta} ex_i$. Abstracting from distance and border effects, we estimate two distinct effects per country (up to a normalization): An importer fixed effect \hat{S}_n and an exporter fixed effect \bar{S}_i . Then to recover the exporter specific component of trade costs, we use the restriction that $\hat{S}_i - \bar{S}_i = \hat{S}_i - (\hat{S}_i - \hat{\theta} ex_i) = \hat{\theta} ex_i$. Thus, the difference in the country specific importer and exporter fixed effects tells us about the exporter specific component of the trade cost.

⁸Intervals are in miles: [0, 375); [375, 750); [750, 1500); [1500, 3000); [3000, 6000); and [6000, 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 \(2010a\)](#). Section 7.2 shows that our results are robust to using this alternative approach.

The exporter-specific trade-cost function was introduced by [Vaugh \(2010b\)](#) who shows that including this term helps EK-type models match cross-country variation in trade flows and aggregate prices. Given that price data are central to our analysis, we wanted to ensure that the model replicates these features of the data. In [Section 7.2](#), we show that using importer fixed effects (as in EK) or aggregate price data to exactly identify trade costs does not change our estimates by economically meaningful amounts. Finally, our results are robust to incorporating bilateral colonial, language, and legal origin ties as well as countries' geographical attributes.

We assume that ν_{ni} reflects other factors and it is orthogonal to the regressors and normally distributed with mean zero and standard deviation σ_ν . This error term simply captures the fact that the observed trade flows are not entirely explained by the gravity equation of trade in practice.⁹ We use least squares to estimate equations (22) and (23). Finally, we explored estimating equations (22) and (23) with the Poisson pseudo-maximum-likelihood estimator advocated by [Silva and Tenreyro \(2006\)](#) and we found that our results for our estimate of θ are robust to this approach.

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, $\hat{\theta}$.

The relationship is obvious in the estimation of trade barriers since $\log \hat{\tau}_{ni}$ is scaled by $\hat{\theta}$ in (22). To see that we can simulate micro-level prices as a function of $\hat{\theta}$ only, notice that for any good j , the model implies that $p_{ni}(j) = \tau_{ni}w_i/z_i(j)$. 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:

$$M_i(u_i) = \exp\left(-\tilde{S}_i u_i^{-\theta}\right), \quad \text{with } \tilde{S}_i = \exp(S_i) = T_i w_i^{-\theta}. \quad (24)$$

Thus, having obtained the coefficients \hat{S}_i from the first-stage gravity regression, we can simulate the inverse of marginal costs and prices.¹⁰

To simulate the model, we assume that there are a large number (150,000) of potentially tradable

⁹We explored a specification in which we included the error term in expression (23) instead of (22), i.e. a structural shock to trade barriers rather than measurement error. The results were nearly identical to our benchmark estimates.

¹⁰Since the \hat{S}' s (and the $\hat{\theta} \log \hat{\tau}'$ s) are sufficient to simulate the model, it is easy to see why the presence of intermediate goods, as in the original EK model, does not affect our estimates of θ . Let P_i denote the price index of tradable (intermediate) goods and let $1 - \alpha$ be the intermediate goods share. Then, the gravity equation is identical to expression (22), except that now $\tilde{S}_i = T_i (w_i^\alpha P_i^{1-\alpha})^{-\theta}$. The estimate of the \tilde{S}' s is the same whether the model has intermediates or not, thus the unit costs drawn from the distributions that reflect the \tilde{S}' s are exactly the same. All that changes is the interpretation of what the unit costs are composed of, but not the unit costs themselves. To see that the estimate of θ is unchanged, derive the trade shares for the model with intermediates, $X_{ni}/X_n = T_i (\tau_{ni} w_i^\alpha P_i^{1-\alpha})^{-\theta} / \sum_{k=1}^N T_k (\tau_{nk} w_k)^{-\theta}$. Dividing this expression by X_{ii}/X_i obtains the same estimating equation for θ given by expression (5). What occurs is that the unit cost component, $w_i^\alpha P_i^{1-\alpha}$, “differences out”, so that the presence of intermediates does not affect the moment that we use in our estimation.

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 (24) 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. Let $\hat{\beta}_k$ be EK’s method of moment estimator defined in (12) using the k th-order statistic over micro-level price differences. Then, the moments we are interested in are:

$$\beta_k = -\frac{\sum_n \sum_i \log\left(\frac{X_{ni}/X_n}{X_{ii}/X_i}\right)}{\sum_n \sum_i \left(\log \hat{\tau}_{ni}^k(L) + \log \hat{P}_i - \log \hat{P}_n\right)}, \quad k = 1, 2 \quad (25)$$

where $\hat{\tau}_{ni}^k(L)$ is computed as the k th-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 (25) 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 (22), which are described in Section 5.2.

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

function:

$$y(\theta) = \begin{bmatrix} \beta_1 - \frac{1}{S} \sum_{s=1}^S \beta_1(\theta, u_s) \\ \beta_2 - \frac{1}{S} \sum_{s=1}^S \beta_2(\theta, u_s) \end{bmatrix}. \quad (26)$$

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

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

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

$$\hat{\theta} = \arg \min_{\theta} [y(\theta)' \mathbf{W} y(\theta)], \quad (27)$$

where \mathbf{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 start with the identity matrix as an initial guess for 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 following the approach described in [Gouriéroux and Monfort \(1996\)](#). With the new estimate of the weighting matrix we solve for a new $\hat{\theta}^1$. We perform this iterative procedure twice. We found that iterating until some convergence criteria gave effectively the same results but with substantial increases in computing time.

Step 5.—We compute standard errors using a parametric bootstrap technique. Given our estimate of θ , estimates of trade costs and S_s , and the error variance σ_ν , we have a completely specified data generating process. We then proceed to simulate micro-level data, add error terms to the trade data, collect a sample of prices, and compute new estimates β_1^b and β_2^b . Next, we estimate the model using moments β_1^b and β_2^b and obtain θ^b . We repeat this procedure 100 times and construct standard errors accordingly.

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: Estimation Results With Artificial Data

Estimation Approach	True $\theta = 8.28$	True $\theta = 4.00$
Overidentified	Mean Estimate of θ (S.E.)	Mean Estimate of θ (S.E.)
SMM	8.27 (0.03)	4.00 (0.02)
Moment, β_1	12.52 (0.06)	6.04 (0.03)
Moment, β_2	15.20 (0.06)	7.34 (0.03)
Exactly Identified		
SMM	8.28 (0.04)	4.00 (0.02)
Moment, β_1	12.52 (0.06)	6.04 (0.03)

Note: S.E. is the standard error of the mean. 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.

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 8.27 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 8.28 with a standard error of 0.04. Again, this is the same as the true value of θ .

¹¹To 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.

The second column reports the results when the true value of θ is set equal to 4.00. The estimates using our estimator are 4.00 in both the overidentified and the exactly identified case, respectively. Similar to the previous results, these values are equivalent to 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 $2\hat{\tau}_{ni}^1 - \hat{\tau}_{ni}^2$, 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

Table 4: Comparison to Alternative Statistical Approaches to Bias Reduction

	True $\theta = 8.28$	True $\theta = 4.00$
Estimation Approach	Mean Estimate of θ (S.E.)	Mean Estimate of θ (S.E.)
SMM	8.28 (0.04)	4.00 (0.02)
Robson and Whitlock (1964)	10.63 (0.06)	5.14 (0.03)
Moment, β_1	12.52 (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.

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 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

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.

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 data sets. The key finding of this section is that our approach yields an estimate around four, in contrast to Eaton and Kortum's (2002) estimation strategy, which results in an estimate around eight.

6.1. New ICP Data

Our sample contains 123 countries. We use trade flows and production data for the year 2004 to construct trade shares. The trade and production data and the construction of trade shares are standard in the literature, so we relegate the details to Appendix 1.1. Instead, in this section, we focus on the price data. To compute aggregate price indices and proxies for trade costs we use basic-heading-level price data from the 2003-2005 round of the International Comparison Program (ICP). Bradford (2003) and Bradford and Lawrence (2004) use earlier rounds of the ICP price data to measure the degree of fragmentation, or the level of trade barriers, among OECD countries. The authors, as well as Deaton and Heston (2010), 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. Appendix 1.2 provides details on the procedure that the ICP uses to construct the price of a basic heading. We reduce the sample to 62 tradable categories based on their correspondence with the trade data that we employ. Table 20 (in the Online Appendix) lists the 62 basic headings used in the analysis.

We choose to work with this dataset for three reasons. First, the database provides extensive coverage, as it includes as many as 123 developing and developed countries that account for 98 percent of world GDP. Second, the sampled goods span all categories of GDP and therefore reflect a number of industries. Third, and most important, because this is the latest round of the ICP, the measurement issues are less severe than in previous rounds.

We recognize that sources of price variation that are outside of the EK model are present in our micro-level price dataset. We consider a number of these sources in later sections as part of our

¹³The ICP Methodological Handbook is available at <http://go.worldbank.org/6VPHKOKHG0>

robustness analysis. In the remainder of this section, however, we describe the nature of the ICP price dataset and we provide useful references for the reader, since this database has not been used extensively by the international trade literature in the past.

Considerable care is taken to ensure that the price data are properly collected and recorded. [Roberts \(2012\)](#) details the extensive preparation and post-collection validation checks that are performed on the ICP price data at both the national and the international level with the explicit purpose of eliminating error in the pricing of products. The ICP addresses two types of errors: product error, which refers to the failure to survey comparable products, and price error, which refers to the failure to record the price of a product correctly.

To minimize product error, the products that appear on the survey are defined very precisely. [Deaton and Heston \(2010\)](#) provide the following three examples of products (within basic headings) whose prices are sampled across countries: (a) Nescafé classic: product presentation, tin or glass jar, 100 grams: type, 100 percent Robusta: variety, instant coffee, caffeine, not decaffeinated: brand, Nestlé-Nescafé classic; (b) Boubou (item within womens clothing): product specification, no package, 1 unit: fibre type, cotton 100 percent: production, small scale: type, boubou: sleeve length, sleeveless: fabric design, brocade: details/features, embroidery; (c) light bulb: product presentation, carton, 1 piece: type, regular: power 40 watts: brand name, indicate brand.

If a product is unavailable, price collectors are instructed to collect the price of a substitute product and record its characteristics. It is sometimes possible to adjust the price for quality differences between the product priced and the product specified. Alternatively, if other countries report prices for the same substitute product, price comparisons can be made for the substitute product as well as for the product originally specified. If neither of these options is available, the price is discarded.

To minimize price error, the ICP validates the data via statistical methods aimed to identify potential outlying prices. Prices are then reconfirmed through additional data collection and/or adjustments if the initial price cannot be verified. Outliers that are large in the statistical sense are discarded. Thus, while measurement error is always a concern, the methods and approaches of the ICP are intended to minimize this error subject to feasibility.

Finally, the ICP does not randomly sample prices from the entire set of produced goods in the world economy. Instead, it 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 that appear across countries, where the goods have been pre-selected due to their highly informative content for the purpose of international comparisons.

It is important to account for this feature of the ICP data in the estimation procedure that relies on the EK model. We argue that EK’s model gives a natural common basket of goods to be priced across countries. In the 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.

6.2. Baseline Results—ICP 2004 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 2004 ICP Data

	Estimate of θ (S.E.)	β_1	β_2
Data Moments	—	7.75	9.64
Exactly Identified Case	4.14 (0.09)	7.75	—
Overidentified Case	4.10 (0.08)	7.67	9.65

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 4.14, roughly half of EK’s estimate of θ .

The third row reports the results for the overidentified estimation. The estimate of θ is 4.10—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.

It should be noted that our estimates of the trade elasticity imply that trade costs are very large. Using the estimate that EK would arrive at from the ICP database—7.75—the median trade cost across all countries is 3.3 with an inter-quartile range between 2.4 and 4.6. Using our estimate of 4.14, the implied trade costs are substantially larger. The median trade cost is now 9.2 with an inter-quartile range between 5.1 and 17.2. The overall size of these costs hides the relatively low frictions between rich countries (recall that there are 123 countries in the ICP dataset, many of which are very poor). For example, the median trade cost between only OECD countries is only 3.2. This estimate is consistent with our estimates of trade frictions using the EK data, which focuses only on rich/developed countries.

¹⁴The results from the Step 1 gravity regressions are presented in Table 16 and Table 19 in the Online Appendix.

6.3. 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.¹⁵ 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: 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 (0.17)	5.93	—
Overidentified Case	4.27 (0.16)	6.45	7.84

Table 6 presents the results.¹⁶ 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 .

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 4.27. Again, this is substantially below EK's estimate. Unlike our results in Table 5 with newer data, the overidentified case is giving a slightly 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.

Once again, there is a substantial difference between the implied trade barriers from EK's and our estimation. Using EK's original estimate of 8.28, the median trade cost is 1.9 with an interquartile range between 1.5 and 2.1. Using our estimate of 3.93, the median trade cost is 3.7 with

¹⁵The data are available here: <http://home.uchicago.edu/kortum/papers/tgt/tgtprogs.htm>.

¹⁶The results from the Step 1 gravity regressions are presented in Table 17 and Table 19 in the Online Appendix.

an inter-quartile range between 2.4 and 4.7. This estimate is about 100 percentage points larger than [Anderson and van Wincoop's \(2004\)](#) estimate that the total trade barrier is equivalent to a 170 percent ad-valorem tax equivalent, or a trade cost of 2.7.

6.4. 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 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 12.8.¹⁷ 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 \(2012\)](#) 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 \(2012\)](#) 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 \(2001\)](#), [Baier and Bergstrand \(2001\)](#), and [Romalis \(2007\)](#). We discuss each of these approaches in turn.

In Appendix D, we explore [Feenstra's \(1994\)](#) method in the context of EK's Ricardian model with micro-level heterogeneity. We find that [Feenstra's \(1994\)](#) method, as well as papers that

¹⁷[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.

build on it such as [Broda and Weinstein \(2006\)](#), [Imbs and Mejean \(2010\)](#), and [Feenstra, Obstfeld, and Russ \(2012\)](#), does not recover the elasticity of trade in EK's Ricardian model with micro-level heterogeneity. In particular, we apply [Feenstra's \(1994\)](#) method to data generated from EK's 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 EK's Ricardian model with micro-level heterogeneity.¹⁸ Hence, elasticity estimates obtained using [Feenstra's \(1994\)](#) approach should not be used in quantitative analysis of Ricardian and monopolistic-competition models with micro-level heterogeneity.¹⁹

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 estimates that [Anderson and van Wincoop \(2004\)](#) report are obtained using time-series and cross-industry variation in tariffs and trade flows during bilateral trade liberalization episodes as in [Head and Ries \(2001\)](#) and [Romalis \(2007\)](#), or time-series variation in tariffs in the cross-section of countries as in [Baier and Bergstrand \(2001\)](#). Recently, [Caliendo and Parro \(2012\)](#) build on these approaches and estimate sectoral trade elasticities from cross-sectional variations in trade flows and tariffs.

The tariff-based estimation approach is appealing because of its simplicity and the appearance of being model-free. However, in order to apply it to EK's 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 \(2012\)](#) argue. Hence, both our methodology and the approach that uses tariffs share the same parametric restrictions in order to be applied to EK's Ricardian model.

Admittedly, there is a substantial difference between the values of the elasticity that are typically obtained using the tariff-based versus our approach. In particular, [Head and Ries \(2001\)](#), [Romalis \(2007\)](#), and [Baier and Bergstrand \(2001\)](#) find values in the range of five to ten, while our benchmark estimates center around four. The corollary is that low values of the elasticity imply large deviations between observed trade frictions (tariffs, transportation costs, etc.) and those inferred from trade flows.

There are two pieces of evidence in support of the values that we find. First, [Parro \(2013\)](#) uses the approach of [Caliendo and Parro \(2012\)](#) to estimate an aggregate trade elasticity for capital

¹⁸The parameter governs the elasticity of trade and welfare in models that do not feature micro-level heterogeneity such as the Armington model as articulated in [Anderson \(1979\)](#) and the [Krugman \(1980\)](#) model (see [Arkolakis, Costinot, and Rodriguez-Clare \(2012\)](#)). In [Simonovska and Waugh \(2013\)](#), we show that our method can also recover this key parameter in the two frameworks that do not feature micro-level heterogeneity.

¹⁹[Feenstra \(2010\)](#) makes a similar argument in the context of the [Melitz \(2003\)](#) model when parameterized as in [Chaney \(2008\)](#).

goods and non-capital, traded goods. He finds estimates of 4.6 and 5.2 with a standard error of 0.27 and 0.29. These point estimates are only modestly larger than ours. Thus, there are tariff based estimates of the elasticity that are consistent with our findings.

Second, our results compare favorably with alternative estimates of the shape parameter of the productivity distribution, θ , that *are not* obtained from gravity-based estimators. For example, estimates of the shape parameter from 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 \(2010\)](#) 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 \(2011\)](#) comes from firm-level data, which suggest that there is a lot of variation in firm productivity. The data in our paper are 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 EK's model. The methodology and the moments that we use to estimate the trade elasticity within EK's Ricardian framework can be derived from other structural gravity models of trade that feature product differentiation. In [Simonovska and Waugh \(2013\)](#), we demonstrate the applicability of our estimator to the Armington model as articulated in [Anderson \(1979\)](#) and the [Krugman \(1980\)](#) model, the Ricardian framework with variable mark-ups of [Bernard, Eaton, Jensen, and Kortum \(2003\)](#), and the monopolistic-competition model of [Melitz \(2003\)](#) as articulated in [Chaney \(2008\)](#).

Our main finding is that different models—estimated to fit the same moments in the data—imply different trade elasticities. The key insight behind the result is that the different margins of adjustment in new trade models, e.g., an extensive margin or variable markups, alter the mapping from the price data to the estimate of the trade elasticity. For example, in [Bernard, Eaton, Jensen, and Kortum \(2003\)](#), mark-ups are negatively correlated with marginal costs implying that the same degree in cost variation results in smaller variation of prices relative to EK. Hence, the maximal price difference lies further below the trade friction relative to the EK model and an even lower trade elasticity is needed to rationalize the same amount of trade observed in the data. Thus, the presence of variable markups yields a lower estimate of the trade elasticity and, hence, larger welfare gains from trade.

7. Robustness

In this section, we conduct a variety of exercises to verify the robustness of our benchmark estimates of the trade elasticity. First, we discuss some computational issues regarding the number of goods in the simulation. Second, we explore the sensitivity of our results to different

assumptions regarding the functional form that trade costs take. Third, we apply estimation approaches that are based on different moment conditions than the ones employed above. Fourth, we allow for additional sources of variation in the price data that are not captured by the EK model. In particular, we discuss issues that relate to distribution costs, mark-ups, good-specific trade costs, product quality, and aggregation, and the possible biases that they may introduce in our estimation.

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 2004 ICP data. For the EK data, notice that our estimates vary little above 5,000 goods. 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 2004 ICP data, which feature 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

²⁰An 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.

²¹The 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.

Table 7: Results with Different # of Goods

Number of Goods	1,000	5,000	25,000	100,000	150,000*
EK’s Data, Exactly Identified Case, $\hat{\theta}$	4.57	4.13	3.99	3.93	3.93
Fraction of Wrong Zeros	0.29	0.10	0.03	0.005	0.003
Fraction of Correct Zeros	—	—	—	—	—
2004 ICP Data, Exactly Identified Case, $\hat{\theta}$	6.54	5.43	4.67	4.22	4.14
Fraction of Wrong Zeros	0.63	0.46	0.31	0.21	0.18
Fraction of Correct Zeros	0.93	0.85	0.72	0.55	0.50

countries predicted by the model in places where we observe positive trade flows in the data. The reason for zero trade flows is that the probability of no trade occurring between any two country pairs is increasing with a more “sparse” approximation of the continuum of goods. This result is consistent with [Armenter and Koren \(2010\)](#) and their emphasis on the sparsity of trade data in accounting for zeros in trade flows.

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 2004 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, though at the cost of getting some observed zeros incorrect.

7.2. Trade Cost Functional Form

This section explores how the results depend upon our assumptions about the particular functional form that trade costs take. Overall, we find that our main results are robust to alternative trade-cost representations.

In **Step 2.** of our estimation we specified a particular functional form (equation (23)) for the trade costs. To explore the sensitivity of our results to this assumption we consider two alternative approaches.

First, we rely on the functional form for trade costs that was employed by EK

$$\log \hat{\tau}_{ni} = d_k + b_{ni} + m_n. \tag{28}$$

The functional form mimics the one used throughout the paper except for the presence of an importer fixed effect rather than the exporter fixed effect (m_n vs. ex_i). This term allows the trade-cost level to vary depending upon the importer rather than the exporter as in the main results of the paper. As discussed in [Vaugh \(2010b\)](#), the model-implied trade flows using this functional form will not differ from the baseline specification; the differences lie in the way in which the trade costs co-vary with country pairs and the \hat{S} 's estimated from the gravity regression.

Second, we re-estimate the trade elasticity without imposing any functional form on the trade costs. In particular, we back out the scaled trade costs as a residual from Equation (7)

$$\log \hat{\tau}_{ni} = -\frac{1}{\hat{\theta}} \log \left(\frac{X_{ni}/X_n}{X_{ii}/X_i} \right) - \log \hat{P}_i + \log \hat{P}_n, \quad (29)$$

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

So given value of $\hat{\theta}$, trade flow data, and proxies for aggregate price indices, we recover trade costs as a residual from (29). What remains is to recover the implied \hat{S} parameters (necessary to simulate marginal costs) that are consistent with these trade costs. To do so, note that there is a log-linear relationship between the \hat{S} 's, a country's price index, and its home trade share,

$$\hat{S}_i = -\hat{\theta} \log \hat{P}_i + \log (X_{ii}/X_i), \quad (30)$$

which allows us to recover the \hat{S} parameters. We should note that these \hat{S} parameters are recovered in a way so that the model will *exactly* (up to any simulation error) replicate the observed trade flows and aggregate price indices.

Table 8 presents the results. The top panel reports the results using the EK data under the baseline specification, the EK functional form, and the residual-based specification. Across all cases, the estimate of θ is effectively the same as in the baseline case, i.e. around 4. The ICP results tell a similar story. The bottom panel reports the results using the ICP data set under the three different trade-cost specifications. Similar to the results with the EK data set, the estimate using the EK functional form is virtually indistinguishable from the one obtained using the [Vaugh \(2010b\)](#) functional form; the residual based specification is only slightly larger (4.46 vs. 4.14) than the baseline.²²

²²The similarity in the results obtained using exporter versus importer fixed effects is not coincidental. Note that the estimator $\hat{\beta}$ relates average trade flows to average trade costs adjusted by aggregate price differences. Using importer or exporter fixed effects will always return the same average trade costs adjusted by aggregate price differences because both importer and exporter fixed effects will return the same average predicted trade flows. The two specifications differ in the way in which the separate components (aggregate prices and trade costs) covary with trade flows, which is the argument made in [Vaugh \(2010b\)](#). Because our estimation is based on

Table 8: Alternative Trade Cost Functional Forms: Estimation Results

Data Set	Functional Form	Estimate of θ (S.E.)
EK Data	Baseline	3.93 (0.17)
	EK Functional Form	3.93 (0.17)
	No Functional Form/Residual	3.96 (0.18)
ICP Data	Baseline	4.14 (0.09)
	EK Functional Form	4.14 (0.09)
	No Functional Form/Residual	4.46 (0.12)

Note: Exactly identified case used across all specifications.

7.3. Alternative Moment Conditions

In this section, we explore an alternative moment condition to estimate θ . The alternative moment condition that we explore uses more information about price dispersion between countries to estimate θ . The idea is to compute the average variance in prices between two countries rather than to study the maximal price difference only. In other words, we are using more information about the entire price distribution than just the max. Note that it is not a priori obvious how our baseline estimate of θ should relate to the estimate that is based upon the alternative moment condition. Hence, this robustness exercise is a strong cross-check on our results.

To operationalize this idea, we define the moment

$$\psi = \frac{1}{N^2 - N} \sum_n \sum_i \text{var}_{ni} (\log p_n(\ell) - \log p_i(\ell)).$$

In particular, for a country pair, we compute the variance of log price differences and then we average across all country pairs (not including own pairs) to arrive at ψ . Given the moment ψ constructed from the data, we then form the following zero function

$$y^A(\theta) = \left[\psi - \frac{1}{S} \sum_{s=1}^S \psi(\theta, u_s) \right], \quad (31)$$

matching averages, for which both specifications have the same predictions, the results should be similar.

with the alternative moment condition

$$E [y^A(\theta_o)] = 0.$$

Table 9: Alternative Moment Conditions: Estimation Results

Data Set	Estimation Approach	Estimate of θ (S.E.)	Data Moments
	Baseline	3.93 (0.17)	5.93
EK Data	Alternative	4.40 (0.16)	0.10
	Baseline	4.14 (0.09)	7.75
ICP Data	Alternative	3.59 (0.06)	0.19

Note: Exactly identified case used across all specifications.

Table 9 presents the results. The first panel reports the results using the EK data set for the baseline and alternative moment conditions. The second panel reports the results using the ICP data set. The results using the average variance in bilateral price differences are in the same ballpark as the baseline estimates. As discussed above, we should not expect them to be the same given that the alternative approach incorporates significantly more information about price dispersion than does the baseline. However, it is reassuring that the different features of the data are telling a similar story.

Finally, we explored an alternative estimation based on the moment $1/\beta_1$, instead of β_1 . The idea behind this moment condition is to bring the estimation closer to our theoretical results in Section 3 which centered around $1/\beta_1$. In the exactly identified estimation, we found that the baseline estimates and the estimates based on the inverse of the moment β_1 are effectively the same regardless of the data set used. The details and results from this exercise are available upon request.

7.4. Additional Sources of Price Variation

The price data that we employ in our benchmark analysis constitute the 2003-2005 round of the ICP. Although the goal of the ICP is to minimize measurement error and to collect prices of comparable products across countries, the reported prices likely reflect additional sources of variation that are not captured by the EK model. In this section, we discuss issues that relate to measurement error, distribution costs, mark-ups, good-specific trade costs, product quality, and aggregation, and the possible biases that they may introduce in our estimation. We conclude that these sources of price variation 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 price 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, we introduce additive log-normal measurement error to the artificial prices in our simulation routine and we re-estimate θ as described in Section 5. This exercise allows us to quantify the sensitivity of our results to assumptions about the extent of measurement error in the data.

Table 10: Results with Measurement Error

Measurement Error, σ	0	0.01	0.05	0.10	0.20
EK's Data, Exactly Identified Case, $\hat{\theta}$	3.93	3.94	4.27	4.38	6.44
2004 ICP Data, Exactly Identified Case, $\hat{\theta}$	4.14	4.15	4.20	4.39	5.59

Note: Estimates of θ are under the assumption that observed prices $p_i(j)$ in logs equal $\log(p_i(j)) = \log(\hat{p}_i(j)) + \epsilon$ with $\epsilon \sim N(0, \sigma)$ where $\hat{p}_i(j)$ is the true price.

Table 10 presents the results from this exercise. Each column presents the estimates of θ for different magnitudes of measurement error, σ , where σ is the standard deviation of the error. The leftmost column with $\sigma = 0$ reproduces our benchmark exactly-identified results. Table 10 shows that our estimates increase with the extent of measurement error, which is consistent with the intuition described above. However, only when measurement error is large do our estimates change in economically meaningful amounts. For example, with $\sigma = 0.20$ our estimate of θ is 6.44 in EK's data and 5.59 in the 2004 ICP data. This amount of measurement error is large as it implies that 30 percent of all prices are mismeasured by 20 percent or more.

While measurement error can meaningfully affect our results, we argue that substantial measurement error (i.e. $\sigma = 0.20$) is implausible for two reasons. First, as discussed in Section 6, the ICP performs extensive validation checks of the price data (at least in the 2003-2005 round) at both the national and international level with the explicit purpose of eliminating error in the pricing of products. Second, if maximal price differences reflected measurement error, then they should be relatively uncorrelated with distance, other gravity variables, and trade shares.

However, we do not observe this in the data.

Table 11: Max Price Differences and Distance: Data and Model

Elasticity w.r.t. Distance	Data	Model ($\sigma = 0$)	Model ($\sigma = 0.20$)
EK Data	0.12 (0.015)	0.08 (0.024)	0.03 (0.015)
2004 ICP Data	0.10 (0.004)	0.07 (0.014)	0.05 (0.010)

Note: The dependent variable is the logarithm of $\hat{\tau} = \max_{\ell \in L} \left\{ \frac{p_n(\ell)}{p_i(\ell)} \right\}$ and this is regressed on the logarithm of distance, a shared border indicator, and country fixed effects. Standard errors are in parenthesis.

The first row of Table 11 reports the results from the regression of log maximal price differences on log distance and other controls. The first column of Table 11 shows that maximal price differences are positively correlated with distance with an elasticity of 0.12 and 0.10 in the EK and in the 2004 ICP data. Both coefficients are statically different from zero and the magnitudes of the elasticity are similar to the findings in Donaldson (2012). The second and third columns of Table 11 report the results of the same regressions using artificial data generated from the estimated model (both to EK data and 2004 ICP data) with and without measurement error. Table 11 shows that the model with large measurement error predicts elasticities that are 30 to 60 percent lower relative to the model with no measurement error. Relative to the data, the model with measurement error yields elasticities that are 50 to 75 percent lower. These results suggest that substantial measurement error is inconsistent with observed correlations between maximal price differences and distance.

Good-Specific Trade Costs. Our estimation abstracts from good-specific trade costs by following the approach of EK and focusing on the estimation of an aggregate θ . The potential cost of this abstraction is that good-specific trade costs may bias the estimate of θ in various ways.

Mode choice (i.e. air, land, sea) to transport goods is a natural mechanism which will lead trade costs to differ across goods. As Hummels (2007) documents, different goods are shipped by different means across borders. Some goods are transported expensively by air, while others are transported cheaply by sea. Oftentimes, the characteristics of the good motivate the differences in shipping modes.

Lux (2012) studies the endogenous choice of shipping mode in an extended EK framework with good- and mode-specific trade costs. The key feature of his model is that the good-specific component of trade costs depends upon the shipping mode. In the model, consumers who source a good choose both the low-cost supplier (as in EK) and the low-cost mode of shipment. An

implication of Lux's (2012) model is that the no-arbitrage condition in (10) should be modified with relative prices bounded by the *lowest* trade cost across different modes. This implies that our estimate of θ (which abstracts from good-specific trade costs) is an *upper bound*. Lux (2012), however, argues that the variation in his estimates of the mode-specific component is small and thus our estimate of θ is unlikely to be biased by economically meaningful amounts.

It is important to note, however, that the presence of good-specific trade costs does not alter our critique of EK's estimator and their estimates of θ . Moreover, our method is applicable to an industry- or a narrowly-defined product-level where presumably good-specific trade costs are less of a concern. One could use a multi-sector EK model (along the lines of Levchenko and Zhang (2011) or Caliendo and Parro (2012)), estimate θ 's at the industry level, and then compute an aggregate θ parameter. In fact, recently, Giri, Yi, and Yilmazkuday (2012) rely on OECD cross-country micro data on prices for 1400 goods, each of which is mapped into one of 21 manufacturing sectors, as well as bilateral trade flows for each of these sectors to estimate sectoral trade elasticities using the method that we develop in this paper.

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 and study the bias that may arise.

The Monte Carlo exercise shows that the bias in β relative to θ is *larger* than in the cases when additive distribution costs are not present. The reason is that additive distribution costs increase low prices proportionally more than high prices, so 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, in Simonovska and Waugh (2013), we apply the estimator developed in this paper to Bernard, Eaton, Jensen, and Kortum's (2003) Ricardian framework which features variable markups. The exercise yields lower estimates of the trade

elasticity relative to the EK model. The crucial observation is that because mark-ups are negatively correlated with marginal costs in [Bernard, Eaton, Jensen, and Kortum \(2003\)](#), the same degree in cost variation results in smaller variation of prices relative to EK. Hence, the maximal price difference lies further below the trade friction relative to the EK model and an even lower trade elasticity is needed to rationalize the same amount of trade observed in the data.

Variable mark-ups may also arise in monopolistic-competition frameworks with micro-level heterogeneity and non-homothetic consumer preferences as in [Simonovska \(2010\)](#). In this model, the relative price of identical goods across countries reflects trade barriers and minimum productivity thresholds necessary to serve each destination, where the latter summarize the effects that destination-specific characteristics have on the level of competition in each market. Two important points need to be made with respect to this model. First, although maximal price differences are not necessarily bounded above by trade barriers in this model, order statistics from the relative price distribution remain to be informative about trade barriers. Second, since preferences are non-homothetic, the trade elasticity and the domestic expenditure share are no longer sufficient statistics to quantify the welfare gains from trade (see [Arkolakis, Costinot, Donaldson, and Rodriguez-Clare \(2012\)](#)). Consequently, we leave it for future research to extend the methodology developed in the present paper to trade models that feature non-homothetic preferences.

Varying Product Quality. Our estimation relies on cross-country variation in prices of goods in order to identify trade costs and ultimately trade elasticities. One may be concerned that relative prices of goods across countries reflect not only trade costs, but also varying product quality. For example, [Baldwin and Harrigan \(2011\)](#), [Johnson \(2012\)](#), and [Manova and Zhang \(2012\)](#) use free-on-board unit value data to argue that richer countries import goods of higher quality from a given source. It is important to note that these studies employ unit values from highly disaggregated trade data as proxies for prices of individual goods.

In contrast, we use the basic-heading-level price data from the 2003-2005 round of the ICP in our study. As described in Section 6, the ICP collects prices of precisely-defined products with identical characteristics across retail locations in the participating countries. With the methodologies and practices employed by the ICP in mind, it is reasonable to argue that varying product quality is not a first-order concern in our data. Moreover, varying product quality should be even less of a concern among countries of similar levels of development. To illustrate this point, we repeat the overidentified estimation using price and trade-flow data for the thirty largest countries in terms of absorption in our dataset. The resulting estimate of the trade elasticity is 4.21, which compares favorably to the benchmark estimates obtained using data on all 123 countries.

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.

To quantify the bias in our estimates that is associated with aggregation, we conduct a robustness exercise using the Economist Intelligence Unit (EIU) data. The EIU surveys the prices of individual goods across various cities in two types of retail stores: mid-priced, or branded stores, and supermarkets, or chain stores. The dataset contains the nominal prices of goods and services, reported in local currency, as well as nominal exchange rates relative to the US dollar, which are recorded at the time of the survey. [Crucini, Telmer, and Zachariadis \(2003\)](#) and [Crucini and Yilmazkuday \(2013\)](#) use the same data to study the determinants of the deviations from the law of one price across cities and countries.

The database spans a subset of 71 countries from our original data set, but provides prices for 110 individual tradable goods. Table 21 in the Online Appendix lists the 110 goods used in the analysis. While in the majority of the countries, price surveys are conducted in a single major city, in 17 of the 71 countries multiple cities are surveyed.²³ For these countries, we use the price data from the city which provided the maximum coverage of goods. In most instances, the location that satisfied this requirement was the largest city in the country. We use prices collected in mid-priced stores in the year 2004 and we combine them with the observations on trade and output from the benchmark analysis.²⁴

Table 12: Estimation Results With EIU Data

	Estimate of θ (S.E.)	β_1	β_2
Data Moments	—	4.39	5.23
Exactly Identified Case	2.82 (0.06)	4.39	—
Overidentified Case	2.79 (0.05)	4.39	5.24

Table 12 presents estimates of the elasticity parameter using EIU’s good-level price data set. The results from the Step 1 gravity regressions are presented in Table 18 and Table 19 in the

²³These countries are Australia, Canada, China, France, Germany, India, Italy, Japan, New Zealand, Russian Federation, Saudi Arabia, South Africa, Spain, Switzerland, United Kingdom, USA, and Vietnam.

²⁴The results are robust to using supermarket price data for the same year.

Online Appendix. The estimates of the elasticity of trade that we obtain from the good-level price data for the subset of 71 countries range between 2.79 and 2.82. As a point of comparison, using the same 71 countries but the ICP price data yields estimates of 3.98 and 4.05, respectively. This finding suggests that the benchmark estimates are potentially biased upwards due to aggregation.

To quantify the bias, we aggregate the individual goods in the EIU dataset to the basic-heading level at which the ICP data are reported. To do so, we first assigned each good in the EIU dataset to one of the 62 basic headings employed in the ICP data. The procedure yields 41 non-zero good categories, with the mode category having 12 products. We then aggregated prices by taking the geometric average across the price of each good within each category. The elasticity estimates increase modestly to 3.04-3.08. This result suggests that aggregation bias is not a first-order concern in the benchmark analysis which relies on the ICP data.

While aggregation is a (small) shortcoming of the ICP data, key benefits of the database are that it offers a wide coverage of both goods and countries, it is readily comparable to the price data employed by EK, and it is least prone to measurement error as discussed in [Roberts \(2012\)](#). Hence, it is reasonable to favor our benchmark estimates of the trade elasticity, which center around four.

8. Conclusion

In this paper we develop a new methodology to estimate the elasticity of trade that builds on EK's 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 span 123 countries that account for 98 percent of world GDP. Across numerous exercises, we obtain estimates of the trade elasticity that range between 2.79 and 4.46. 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.

References

- ANDERSON, J., AND E. VAN WINCOOP (2004): "Trade Costs," *Journal of Economic Literature*, 42(3), 691–751.
- ANDERSON, J. E. (1979): "A Theoretical Foundation for the Gravity Equation," *American Economic Review*, 69(1), 106–16.
- ARKOLAKIS, C., A. COSTINOT, D. DONALDSON, AND A. RODRIGUEZ-CLARE (2012): "The Elusive Pro-Competitive Effects of Trade," MIT, *mimeo*.

- ARKOLAKIS, C., A. COSTINOT, AND A. RODRIGUEZ-CLARE (2012): "New Trade Models, Same Old Gains?," *American Economic Review*, 102(1), 94–130.
- ARMENTER, R., AND M. KOREN (2010): "A Balls-and-Bins Model of Trade," CEPR Discussion Papers 7783, C.E.P.R. Discussion Papers.
- BAIER, S., AND J. BERGSTRAND (2001): "The Growth of World Trade: Tariffs, Transport Costs, and Income Similarity," *Journal of International Economics*, 53(1), 1–27.
- BALDWIN, R., AND J. HARRIGAN (2011): "Zeros, Quality, and Space: Trade Theory and Trade Evidence," *American Economic Journal: Microeconomics*, 3(2), 60–88.
- 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 (2010): "Globalization, Technology, and the Skill Premium: A Quantitative Analysis," NBER Working Papers 16459, National Bureau of Economic Research, Inc.
- CALIENDO, L., AND F. PARRO (2012): "Estimates of the Trade and Welfare Effects of NAFTA," Working Paper 18508, National Bureau of Economic Research.
- 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.
- CRUCINI, M. J., C. TELMER, AND M. ZACHARIADIS (2003): "Price dispersion: The role of distance, borders and location," GSIA Working Papers 2004-E25, Carnegie Mellon University, Tepper School of Business.

- CRUCINI, M. J., AND H. YILMAZKUDAY (2013): "Understanding Long-run Price Dispersion," Working Paper 18811, National Bureau of Economic Research.
- DEATON, A., AND A. HESTON (2010): "Understanding PPPs and PPP-Based National Accounts," *American Economic Journal: Macroeconomics*, 2(4), 1–35.
- DONALDSON, D. (2012): "Railroads of the Raj: Estimating the Economic Impact of Transportation Infrastructure," *Massachusetts Institute of Technology*, *mimeo*.
- 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*, 79(5), 1453–1498.
- 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 (2012): "In Search of the Armington Elasticity," *University of California - Davis*, *unpublished mimeo*.
- GIRI, R., K.-M. YI, AND H. YILMAZKUDAY (2012): "Gains from Trade: Does Sectoral Heterogeneity Matter?," *mimeo*.
- 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.
- HUMMELS, D. (2007): "Transportation Costs and International Trade in the Second Era of Globalization," *Journal of Economic Perspectives*, 21(3).
- IMBS, J., AND I. MEJEAN (2010): "Trade elasticities," *Proceedings*, (Oct).
- JOHNSON, R. C. (2012): "Trade and prices with heterogeneous firms," *Journal of International Economics*, 86(1), 43 – 56.
- KRUGMAN, P. (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70(5), 950–959.

- LEVCHENKO, A. A., AND J. ZHANG (2011): "The Evolution of Comparative Advantage: Measurement and Welfare Implications," Working Paper 16806, National Bureau of Economic Research, Inc.
- LUX, M. (2012): "Defying Gravity: The Substitutability of Transportation in International Trade," *unpublished manuscript*.
- MACKINNON, J. G., AND A. A. SMITH (1998): "Approximate bias correction in econometrics," *Journal of Econometrics*, 85(2), 205–230.
- MANOVA, K., AND Z. ZHANG (2012): "Export Prices Across Firms and Destinations," *The Quarterly Journal of Economics*, 127(1), 379–436.
- MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725.
- PARRO, F. (2013): "Capital-Skill Complementarity and the Skill Premium in a Quantitative Model of Trade," *American Economic Journal: Macroeconomics*, 5(2), 72–117.
- ROBERTS, D. (2012): "Validation of ICP Regional Prices and Basic Heading PPPs," chap. 9, *Measuring the Real Size of the World Economy The Framework, Methodology, and Results of the International Comparison Program (ICP)*.
- 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.
- SILVA, J. M. C. S., AND S. TENREYRO (2006): "The Log of Gravity," *The Review of Economics and Statistics*, 88(4), 641–658.
- SIMONOVSKA, I. (2010): "Income Differences and Prices of Tradables," Working Paper 16233, National Bureau of Economic Research.
- SIMONOVSKA, I., AND M. WAUGH (2013): "Trade Models and Trade Elasticities," *New York University and University of California - Davis, mimeo*.
- SMITH, A. (2008): "Indirect inference," *The New Palgrave Dictionary of Economics, 2nd Edition*.
- 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.