

The Elasticity of Trade: Estimates and Evidence

Ina Simonovska

University of California, Davis and NBER

Michael E. Waugh

New York University

December 2009

PRELIMINARY AND INCOMPLETE,
Comments Welcome

ABSTRACT

Quantitative results from a large class of structural gravity models of international trade depend critically on a single parameter governing the elasticity of trade with respect to trade frictions. We provide a new method to estimate this elasticity and illustrate the merits of our approach relative to the estimation strategy of Eaton and Kortum (2002). We employ this method on data for 123 developed and developing countries for the year 2004 using new disaggregate price and trade flow data. Our benchmark estimate for all countries is approximately 4.5, nearly 50 percent lower than the alternative estimation strategy would suggest, but within the range of existing estimates obtained using a variety of data. This difference implies a doubling of the measured welfare costs of autarky across a large class of widely used trade models.

Email: inasimonovska@ucdavis.edu, mwaugh@stern.nyu.edu. We are grateful to the World Bank for generously providing us with the price data from the 2005 ICP round. We thank George Alessandria, Robert Feenstra, Timothy Kehoe, B. Ravikumar, and the International Reading Group at NYU for their feedback.

1 Introduction

Quantitative results from a large class of models of international trade depend critically on a single parameter that governs the elasticity of trade with respect to trade frictions.¹ Eaton and Kortum (2002) provided an innovative approach to estimating this elasticity and their estimate is widely used in quantitative applications. We argue that the approach of Eaton and Kortum (2002) yields systematically biased estimates of this elasticity. We develop a new method to estimate this elasticity and we apply our method using new disaggregate price and trade flow data with 123 countries representing 98 percent of world GDP for the year 2004.

To illustrate how important this parameter is consider three examples: Anderson and van Wincoop (2003) find that the estimate of the tariff equivalent of the U.S.-Canada border varies between 48 and 19 percent depending upon the assumed elasticity of trade with respect to trade frictions. Yi (2003) points out that observed reductions in tariffs can explain almost all or none of the growth in world trade depending upon this elasticity. Arkolakis, Costinot, and Rodríguez-Clare (2009) argue that this parameter is one of only two parameters needed to measure the welfare cost of autarky in a large and important class of trade models. Therefore this elasticity is key to understanding the size of the frictions to trade, the response of trade to changes in tariffs, and the welfare gains or losses from trade.

Estimating this parameter is difficult because quantitative trade models can rationalize small trade flows with either large trade frictions and small elasticities or small trade frictions and large elasticities. Thus, one needs satisfactory measures of trade frictions *independent* of trade flows to estimate this elasticity. Eaton and Kortum (2002) solved 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 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 argue that the approach of Eaton and Kortum (2002) leads to systematically biased estimates of the elasticity of trade. We show this by performing a simple monte carlo experiment by discretizing the Eaton and Kortum (2002) model, simulating trade flows and product-level prices under an assumed elasticity of trade, and then applying their approach. We find that one cannot recover the true elasticity of trade and that the estimates are biased upward by economically significant magnitudes.

The approach of Eaton and Kortum (2002) fails to recover the true parameter because the sample size of prices (typically 50-70 depending on the data set) is small relative to the number of

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

goods in the economy. This is a problem because the probability that the max operator over a small sample of prices actually recovers the true trade cost is close to zero and the estimated trade cost will always be less than the true trade cost. Because the trade costs are almost always underestimated, this leads to systematic upward estimates of the elasticity of trade.

We develop a new method to estimate this elasticity when the sample size of prices is small. Our approach exploits the ability to simulate trade flows and prices as a function of the parameter of interest. Given our ability to simulate these objects, we employ a simulated method of moments estimator that minimizes the distance between the regression coefficients from the approach of Eaton and Kortum (2002) on real and artificial data. We explore the properties of this estimator using simulated data and show that it can recover the true elasticity of trade in contrast to the alternative.

We apply our method to a new and unique data set. The new data set we employ has 123 countries representing 98 percent of world GDP using new disaggregate price and trade flow data. The innovative feature of this data set is its coverage of developing countries. Previous estimates of this elasticity often come from small samples of developed countries.² Thus the applicability of these estimates in the analysis of trade with both developed and developing countries is an important issue we can address.

Our benchmark estimate using this data is approximately 4.5. In contrast, the approach of Eaton and Kortum (2002) would yield estimates between 7.5 and 9.5 depending on if the max or the second order statistic is used to approximate the trade friction.³ We also apply our method using the same data set of Eaton and Kortum (2002) with only developed countries and estimate this elasticity to be approximately 4.5. This is in contrast to their preferred estimate of 8.28. Thus our results provide strong evidence that the estimated elasticity of trade is in the range of 4.5 not 7-9 as the approach of Eaton and Kortum (2002) would suggest. Our results also provide suggestive evidence that this elasticity does not vary depending upon countries' level of development.

Why does this matter? As noted earlier, this matters because the welfare gains in these models depend critically on this elasticity. Our new estimate of this elasticity implies a doubling of the percentage change in real income necessary to compensate a representative consumer for going to autarky, i.e. the welfare cost of autarky. Thus while new heterogeneous firm and production models may yield no larger welfare gains over simpler models as Arkolakis, Costinot, and Rodríguez-Clare (2009) argues, only with the structure of a heterogeneous production model

²See, for example, Head and Ries (2001) for the United States and Canada, Baier and Bergstrand (2001) and Eaton and Kortum (2002) for OECD countries, or the survey of these and several other studies in Anderson and van Wincoop (2004).

³Our approach is robust to using either the max or the second order statistic, while the approach of Eaton and Kortum (2002) always generates larger estimates using the second order statistic.

such as Eaton and Kortum (2002), could we have used both measurement and theory to arrive at a more robust and better estimate of the elasticity of trade and hence the welfare gains from trade.

2 Model

In the following subsections, we describe several popular models of international trade and show how they all relate trade shares, prices and trade costs in the same exact manner. Furthermore, across all these models one parameter shows up in these relationships that controls how trade shares respond to changes in trade frictions or what we term the elasticity of trade.

2.1 Benchmark: Ricardian Model With Heterogeneity

We analyze a version of the multi-country Ricardian model of trade introduced by Eaton and Kortum (2002). 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 Z_i (drawn independently for each j) from its country-specific Fréchet probability distribution $F_i(z) = \exp(-T_i z^{-\theta})$. The country-specific parameter $T_i > 0$ governs the location of the distribution, thus higher values of it imply that a high productivity draw for any good j is more likely. The parameter $\theta > 1$ is assumed to be common across countries and if higher, it generates less variability within the distribution.⁴

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 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

⁴In our quantitative analysis, we estimate values for this parameter for different sets of countries and conclude that they are fairly similar, a finding that supports this assumption.

the triangle inequality: For any three countries i, k, n , $\tau_{ni} \leq \tau_{nk}\tau_{ki}$. With these in mind, the marginal cost of production and delivery of good j from country i to destination n is given by:

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

International markets are perfectly competitive, so consumers in destination n would pay $p_{ni}(j)$, should they decide to buy good j from country i . Thus, the actual price consumers in n pay for good j is the minimum price across all sources k :

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

Substituting the pricing rule into the productivity distribution allows us to obtain the following price index for each destination n :

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

In the above equation,

$$\gamma = \left[\Gamma \left(\frac{\theta + 1 - \rho}{\theta} \right) \right]^{\frac{1}{1-\rho}},$$

where Γ is the Gamma function and parameters are restricted such that $\theta > \rho - 1$.

Furthermore, 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 country n spends on goods from country i or, as we will call it, the trade share is

$$\frac{X_{ni}}{X_n} = \frac{T_i (\tau_{ni} w_i)^{-\theta}}{\sum_{k=1}^N T_k (\tau_{nk} w_k)^{-\theta}}. \quad (2)$$

Note that the sum across k for a fixed n must add up to one.

Expressions (1) and (2) allow us to relate observed expenditure shares to bilateral trade frictions and the price indices of each trading partner via the following equation:

$$\frac{X_{ni}/X_n}{X_{ii}/X_i} = \left(\frac{\tau_{ni} P_i}{P_n} \right)^{-\theta}. \quad (3)$$

2.2 Armington Model Without Heterogeneity

In principal there is nothing unique about equation (3) to the model of Eaton and Kortum (2002). The model of Anderson and van Wincoop (2003) generates equation (3) as well. To do so, assume that each country has constant returns technologies with competitive firms producing a good which is defined by its country of origin, i.e., the Armington assumption. These assumptions imply the unit cost (and price) to deliver a country i good to destination n is $p_{ni} = \tau_{ni} T_i^{\frac{1}{\theta}} w_i$. Similarly to above, w_i is the unit labor cost in country i and $T_i^{\frac{1}{\theta}}$ is total factor productivity there.

Preferences are equally simple. Each country has symmetric constant elasticity preferences over all the (country-specific) goods with common elasticity of substitution $\theta + 1 > 1$. The model yields expenditure shares

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

Given preferences, destination n faces the following price index of tradable goods:

$$P_n = \Upsilon \left[\sum_{k=1}^N T_k (\tau_{nk} w_k)^{-\theta} \right]^{\frac{1}{\theta}}, \quad (5)$$

where Υ contains constant terms.

Expressions (4) and (5) allow us to relate observed expenditure shares to bilateral trade frictions and the price indices of each trading partner via the following equation:

$$\frac{X_{ni}/X_n}{X_{ii}/X_i} = \left(\frac{\tau_{ni} P_i}{P_n} \right)^{-\theta}. \quad (6)$$

This is the same expression as in (3) relating the bilateral trade shares to trade costs and the relative aggregate price of tradables.

2.3 Monopolistic Competition Model With Heterogeneity

Monopolistic competition models of trade in the spirit of Melitz (2003), under the parametrization proposed by Chaney (2008), turn out to generate an identical relationship between prices, trade frictions and trade flows. As in previous sections, consumers are assumed to derive utility from the consumption of varieties originating from different source countries, combined in an aggregate symmetric CES bundle with constant elasticity of substitution $\rho > 1$. Each variety,

however, is produced by a single firm, where firms are differentiated by their productivity, z , and country of origin, i . In every country i , there exists a pool of potential entrants who pay a fixed cost, $f_i > 0$, and subsequently draw a productivity from a Pareto distribution, $T_i z^{-\theta}$, with support $[T_i^{1/\theta}, \infty)$.⁵ Only a measure J_i of them enter in equilibrium and firm entry and exit drives average profits in each country to zero.

This model gives rise to the following expenditure share for each destination n on goods from source i :

$$\frac{X_{ni}}{X_n} = \frac{J_i T_i (\tau_{ni} w_i)^{-\theta}}{\sum_{k=1}^N J_k T_k (\tau_{nk} w_k)^{-\theta}}, \quad (7)$$

where the equilibrium number of entrants is proportional to the fixed cost of entry in each country, $J_i = (\rho - 1)/\rho\theta L_i/f_i$. Given preferences, destination n faces the following price index of tradable goods:

$$P_n = \Upsilon \left[\sum_{k=1}^N J_k T_k (\tau_{nk} w_k)^{-\theta} \right]^{-\frac{1}{\theta}} \left(\frac{f_n}{L_n} \right)^{\frac{-\theta-1+\rho}{-\theta(\rho-1)}}, \quad (8)$$

where Υ contains constant terms. Assuming that fixed costs of entry are proportional to market size, $(\forall k) f_k = AL_k$, equations (7) and (8) yield expression (3) as in the model of Eaton and Kortum (2002) and (6) using the Armington model.

2.4 Monopolistic Competition Model Without Heterogeneity

Variants of the monopolistic competition model of Krugman (1980) also generate an identical relationship between prices, trade frictions and trade flows as above. These models can be thought of as assuming degenerate firm productivity distributions in the frameworks of Melitz (2003) and Chaney (2008) outlined above. Moreover, they give rise to trade shares and prices that much resemble the ones suggested by the Armington Ricardian model of Anderson and van Wincoop (2003). Hence, expression (3) or (6) follows.

⁵That the two models yield identical relationships between prices, trade flows and trade costs is not a surprise given the identical preference assumptions as well as the link between the Fréchet and Pareto distributions. This link is further explained in the appendix.

2.5 The Elasticity of Trade

As seen in previous subsections, a key equation arising from a large class of models is

$$\frac{X_{ni}/X_n}{X_{ii}/X_i} = \left(\frac{\tau_{ni}P_i}{P_n} \right)^{-\theta}. \quad (9)$$

The parameter of interest is θ . To see how this parameter is interesting, take logs of equation (9) yielding

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

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 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.

This elasticity takes on an even larger role than merely controlling trade's response to trade frictions. Arkolakis, Costinot, and Rodríguez-Clare (2009) argue further that this elasticity is one of only two objects that control the welfare gains from trade in the same class of models we discussed above. Thus this elasticity is absolutely critical in any quantitative study of international trade in a large class of models.

3 Estimating θ : Eaton and Kortum's (2002) Approach

Equation (9) suggests that one could easily identify θ if one had data on trade shares, aggregate prices, and trade costs. However, the identification problem that one faces is that trade costs are not observed. That is one 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. Eaton and Kortum (2002) employ an innovative approach to approximate trade costs τ_{ni} . They exploit *disaggregate* price information across countries by arguing that the maximum price difference between two countries bounds the trade costs between the two countries via simple no-arbitrage arguments.

To illustrate Eaton and Kortum's (2002) argument, consider the following example: Suppose there are two countries (home and foreign) and two goods (TVs and DVD players) and prices for each of these goods are observed as in Table 3.

Table 3: Two countries and Two Prices

	TV's	DVD
Price Home	150	125
Price Foreign	100	100

Table 3 provides the following information about trade costs between the two countries. First, notice that if the trade cost $\tau_{h,f} < 1.50$, than someone in the home country could simply import TV's from the foreign country and sell them at a profit and bid away the price difference. Thus the trade friction is no less than 1.50. Notice—and this is a key point to understand for our argument—that this is only a lower bound. Only if the home country actually imports TV's does one know that the trade friction is 1.50. If the home country is not importing TV's then the trade friction may be greater than or equal to 1.5.

In general, it must be the case that for a given good ℓ , $\frac{p_n(\ell)}{p_i(\ell)} \leq \tau_{ni}$; otherwise, there would be an arbitrage opportunity as described above. This suggests that an estimate of τ_{ni} is the maximum of relative prices over goods ℓ . To summarize, Eaton and Kortum's (2002) proxy for τ_{ni} , in logs, is

$$\log \hat{\tau}_{ni} = \max_{\ell} \{ \log (p_n(\ell)) - \log (p_i(\ell)) \}, \quad (10)$$

where the max operator is over all ℓ goods.

Using (10), trade data, and the average over disaggregate price data to approximate p_i , Eaton and Kortum (2002) exploit the structural relationship in (9) to estimate θ . Details specific to their estimate are that they use a method of moments estimator and the second order statistic rather than the max. This approach yields their preferred estimate of 8.28.

4 Monte Carlo Evidence

In this section, we study Eaton and Kortum's (2002) approach to estimating θ as described in section 3. We study their approach by simulating a data set under an assumed value for θ and see if Eaton and Kortum's (2002) approach can recover the true value of θ that generated the data. Our main finding is that their approach cannot and that their estimates of θ are biased upward by quantitatively significant amounts. We argue that this failure arises because of a limited sample of prices to estimate trade costs.

4.1 Simulation Approach

We want to simulate a data set from a stochastic Ricardian model along the lines of Eaton and Kortum (2002) that resembles data.⁶ We use the approach described in the steps below. This simulation approach also provides the foundations for the simulated method of moments estimator we propose in the next section.

Step 1.—We estimate parameters for the country specific Fréchet distributions and trade costs from bilateral trade flow data. We perform this step by following Eaton and Kortum (2002) and Waugh (2009) by deriving the following gravity equation from equations 2 by dividing the bilateral trade share by the importing countries home trade share,

$$\log \left(\frac{X_{ni}/X_n}{X_{nn}/X_n} \right) = S_i - S_n - \theta \log \tau_{ni}, \quad (11)$$

in which S_i is defined as $\log [w_i^{-\theta} T_i]$. Note that this is a different equation than that used to estimate θ in (9) which is derived by dividing the bilateral trade share by the exporting countries home trade share. S_i s are recovered as the coefficients on country-specific dummy variables given the imposed restrictions on how trade costs can covary across countries. Following the arguments of Waugh (2009), trade costs take the following functional form:

$$\log(\tau_{ni}) = d_k + b_{ni} + ex_i + \epsilon_{ni}. \quad (12)$$

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 intervals.⁷ b_{ni} is the effect of a shared border in which $b_{ni} = 1$, if country i and n share a border and zero otherwise. The term ex_i is an exporter fixed effect and allows for the trade cost to vary in level depending upon the exporter. We assume ϵ_{ni} reflects barriers to trade arising from all other factors and is orthogonal to the regressors. We use least squares to estimate equations (11) and (12) to the bilateral trade shares.

Before proceeding, note that what we are doing here is exploiting the fact that we can estimate all necessary parameters to simulate trade flows and prices up to a constant, θ . This allows us to be able to simulate data as a function of the parameter θ only.

Step 2.—With an assumed θ , the estimated \hat{S}_i parameterize the Fréchet distributions for each country and the level of trade costs. In the simulations that follow, we set θ equal to 8.28—the

⁶In all the monte-carlo experiments, we use the trade data in Eaton and Kortum (2002) in Step 1 of the simulation procedure. Section 6.2.1 describes their data in more detail.

⁷Intervals are in miles: [0, 375); [375, 750); [750, 1500); [1500, 3000); [3000, 6000); and [6000, maximum]. We are currently exploring the robustness of our results to alternative trade cost specifications.

preferred estimate of Eaton and Kortum (2002). With the parameterized distributions and trade costs, we can then simulate the model.

To simulate the model, we assumed there is a large number (100,000) of potentially tradable goods. For each country, good-level efficiencies are drawn from the country-specific distribution and assigned to the production technology for each good. Then, for each importing country and each good, the low-cost supplier across countries is found, realized prices are recorded, and the 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 are 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, ϵ_{ni} , from Step 1.

Step 4.—Given the prices and trade shares, we then employ the estimation strategy suggested by Eaton and Kortum (2002).

We should note that the most important variable in the simulation is the sample size of the prices. It is important because small samples of prices will lead to significantly biased estimates of θ . In our baseline simulation, we use a sample size of 50. This is the same sample size of prices used in Eaton and Kortum (2002).

4.2 Monte Carlo Results

Table 2 presents the results from the steps outlined above. The columns of Table 2 present the mean and median estimates of θ over the 100 simulations. The rows present different estimation approaches, i.e. simple least squares and method of moments (the preferred approach of Eaton and Kortum (2002)) all with intercepts suppressed. The top panel uses the first order statistic. The bottom panel uses the second order statistic as used in the preferred approach of Eaton and Kortum (2002).

The key result from Table 2 is that the estimates of θ are significantly larger than the true θ that generated the data. As discussed, the underlying θ was set equal to 8.28 and the estimated θ 's in the simulation are between 12 and 15. This suggests the approach of Eaton and Kortum (2002) cannot recover the assumed value of θ and that this approach generates estimates that are biased upward by quantitatively significant amounts.

Table 2: Monte Carlo Results, True $\theta = 8.28$

Approach	Mean Estimate of θ (S.E.M.)	Median Estimate of θ
First Order Statistic		
Least Squares	12.1 (0.06)	12.1
Method of Moments	12.5 (0.06)	12.5
Second Order Statistic		
Least Squares	14.7 (0.06)	14.7
Method of Moments	15.2 (0.06)	15.2

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

4.3 Why the Failure?

The problem is that the sample size of prices used to construct estimates of trade costs is small. A small sample is problematic because estimates of the trade costs are approximated by the *maximum* price difference across realized prices. In a small sample of prices, the maximal price difference is likely to be far from the true maximal price difference. Put another way, in a small sample of prices, its likely that the inequality $\frac{p_n(\ell)}{p_i(\ell)} \leq \tau_{ni}$ is not binding. The implication of this is that the estimates of trade costs are likely to be less than the true trade cost. Because the size of the estimated trade costs are critical to estimating the parameter θ , the estimated θ s are larger than those really generating the data.

To concretely illustrate this, reconsider the same example from Section 3 but with three goods (TVs and DVD players and Xbox's) and prices for each of these goods are observed as in Table 3. The new information from Table 3 suggests a new estimate of the trade cost to be 1.65. The

Table 3: Two countries and Three Prices

	TV's	DVD	XBox
Price Home	150	125	165
Price Foreign	100	100	100

previous estimate of $\tau_{h,f} = 1.50$ with only two prices is biased downward by 0.15 when three prices are considered.

To see how a downward biased estimate of τ leads to an upward biased θ , consider Eaton and

Table 4: Results Increasing the Sample of Prices, True $\theta = 8.28$

Sample Size of Prices	Mean Estimate of θ (S.E.M.)	Median Estimate of θ
50	12.14 (0.06)	12.15
500	9.41 (0.02)	9.40
5,000	8.47 (0.01)	8.46
50,000	8.29 (0.01)	8.29

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

Kortum’s (2002) method of moments estimator for simplicity:

$$\hat{\theta} = -\frac{\frac{1}{M} \sum \log \left(\frac{X_{ni}/X_n}{X_{ii}/X_i} \right)}{\frac{1}{M} \sum \log \left(\frac{P_i \tau_{ni}}{P_n} \right)}.$$

The numerator is the average over the log of relative trade shares. The denominator is the average over the log of relative prices and trade costs. Notice that if trade costs are systematically downward biased, then this lowers the denominator and increases the estimate of θ .⁸

Evidence supporting this argument is seen in the estimated trade costs which are smaller relative to the trade costs generating the data. For example, the average over the simulations of the median estimated trade cost across all country pairs equals a 70 percent tariff rate equivalent. However, the true median trade costs across all country pairs equals a 200 percent tariff rate equivalent. With only 100 prices, the estimated trade costs are biased downwards resulting in estimates of θ that are biased upwards.

To further advance this argument, we performed the same exercise with 500, 5,000, and 50,000 sampled prices. Table 4 presents the results. Notice how the estimate of θ becomes less biased and begins to approach the true value of θ as the sample of prices becomes larger. However, the rate of convergence is extremely slow; even with a sample size of 5,000 the estimate of θ is larger than the value generating the data. Only when 50,000 prices are sampled—one half of all goods in the economy—does the estimate converge to the same value.

Table 4 suggests that data requirements needed to yield an unbiased estimate of θ are extreme. This observation motivates an estimation approach that will solve the problem we have identified from estimating θ using a limited sample of prices to construct estimates of trade costs.

⁸Though we focused on their method of moment estimator in this example, similar logic applies if least squares is used instead.

4.4 Characterizing the Properties of $\hat{\tau}$

The prices of individual goods $p_i(\ell)$ inherit distributional properties from the distributions over technologies. This implies that in principal one can characterize the distribution of relative prices across destinations $p_n(\ell)/p_i(\ell)$. This suggests that one can characterize the distribution and expected value of the maximum over relative prices in a finite sample. In this section, we pursue this approach and hope to use these results to analytically solve the problems discussed above.

TO BE COMPLETED ...

5 Solution: A New Approach To Estimating θ

In this section we suggest a new approach to estimating θ and discuss its performance on simulated data. The basic idea is to exploit the ability to simulate from the model and propose a simple simulated method of moments estimator that uses the regression coefficients from the approach of Eaton and Kortum (2002).

5.1 Simulation

Steps 1-3 in Section 4.1 outline our approach to simulate data, such as trade shares and good-level prices, as a function of our parameter of interest θ .

5.2 Moments

Here we will define the moments of interest. Define α and β as the intercept and slope from the regression:

$$\log \left(\frac{X_{ni}/X_n}{X_{ii}/X_i} \right) = \alpha + \beta \log \left(\hat{\tau}_{ni} \times \frac{\hat{P}_i}{\hat{P}_n} \right) + v_{ni} \quad (13)$$

where “hat” terms denote that they are estimated from good-level price data. The data moments α and β are the moments we are interested in.

We will denote the simulated moments as $\alpha(\theta, u_s)$ and $\beta(\theta, u_s)$ which come from the analogous regression as in (13), except that the trade shares, estimated trade costs, and estimated price indices are from simulated data as 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 that are sampled from all countries. The third component mimics the residuals ϵ_{ni} from equation (11) and described in Section 4.1.

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

$$y(\theta) = \begin{bmatrix} \alpha - \frac{1}{S} \sum_{s=1}^S \alpha(\theta, u_s) \\ \beta - \frac{1}{S} \sum_{s=1}^S \beta(\theta, u_s) \end{bmatrix}. \quad (14)$$

5.3 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

$$\hat{\theta} = \arg \min_{\theta} [y(\theta)' \mathbf{W} y(\theta)], \quad (15)$$

where \mathbf{W} is a 2×2 weighting matrix which we discuss below. The idea behind this moment condition is that though α and β will be biased away from 0 and θ , the moments $\alpha(\theta, u_s)$ and $\beta(\theta, u_s)$ 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 is closely related to indirect inference as discussed in Smith (2008).⁹

For the weighting matrix, we use the inverse of the estimated variance-covariance matrix Ω of the moments α and β estimated from the data.¹⁰ To compute Ω , we used a simple bootstrap procedure outlined in the following steps.

Step 1.—Using the residuals v_{ni} from the regression in (13) and the fitted values, we resampled the residuals v_{ni} with replacement and generated a new set of data using the fitted values.

⁹A key issue in MacKinnon and Smith (1998) is how the bias function behaves. One can numerically show that the bias function is approximately log-linear suggesting the bias function is well behaved. Using the results of Section 4.4, we hope to be able to prove this. Furthermore, the performance of our estimator on simulated data suggests our estimator can correctly recover the true value of θ .

¹⁰This weighting matrix makes sense for the following arguments: First, the optimal weighting matrix should be the inverse of the variance-covariance matrix of $y(\theta_o)$. Second, note that $\text{Var}(y(\theta_o)) = \text{Var}([\alpha, \beta]) + \frac{1}{S} \text{Var}([\alpha(\theta_o, u_s), \beta(\theta_o, u_s)]) = (1 + \frac{1}{S}) \text{Var}([\alpha, \beta])$. Thus the appropriate weighting matrix is $\{(1 + \frac{1}{S}) \text{Var}([\alpha, \beta])\}^{-1}$. See Davidson and MacKinnon (2004) for more details.

Using the data constructed from each resampling b , we computed an intercept term α^b and β^b .

Step 2.—Define the difference between the bootstrap generated moments and data moments as:

$$m^b = \begin{bmatrix} \alpha - \alpha^b \\ \beta - \beta^b \end{bmatrix} \quad (16)$$

we then computed the variance-covariance matrix as

$$\Omega = \frac{1}{B} \sum_{b=1}^B (m^b) \times (m^b)' \quad (17)$$

then the weighting matrix W is set equal to Ω^{-1} .

We compute standard errors using a bootstrap technique. Here it is important to take into account both sampling error and simulation error. To account for sampling error, each bootstrap b replaces the moments α and β with bootstrap generated moments α^b and β^b . Then to account for simulation error, a new seed is generating a new set of model generated moments: $\frac{1}{S} \sum_{s=1}^S \alpha(\theta, u_s)^b$ and $\frac{1}{S} \sum_{s=1}^S \beta(\theta, u_s)^b$. Then defining $y^b(\theta)$ as the difference in moments for each b as in (16), we solve for

$$\hat{\theta}^b = \arg \min_{\theta} [y^b(\theta)' \mathbf{W} y^b(\theta)]. \quad (18)$$

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

$$\text{S.E.}(\hat{\theta}) = \left[\frac{1}{100} \sum_{b=1}^{100} (\hat{\theta}^b - \hat{\theta})(\hat{\theta}^b - \hat{\theta})' \right]^{\frac{1}{2}} \quad (19)$$

This procedure to constructing standard errors is similar in spirit to the approach employed in Eaton, Kortum, and Kramarz (2008) who use a simulated method of moments estimator to estimate the parameters of a similar trade model from the performance of French exporters.

5.4 Performance on Simulated Data

In this section, we evaluate the performance of our estimation approach using simulated data when we know the true value of θ . In all the results that followed, we set the true value of θ equal to 8.28.

Table 5 presents the results from this exercise. The first row presents our simulated method

of moments estimate which is 8.47 with a standard error of 0.21. This is not far from the true value of θ generating the data. Furthermore, the deviation of our estimate from the true value is normal given the standard error.

Table 5: Estimation Results With Simulated Data

Estimation Approach	Estimate of θ	Standard Error
First Order Statistic	True $\theta = 8.28$	
SMM	8.47	0.21
Least Squares	12.47	0.25
Method of Moments	13.06	0.25
Second Order Statistic	True $\theta = 8.28$	
SMM	8.35	0.21
Least Squares	14.78	0.31
Method of Moments	15.39	0.31

To emphasize the performance of our estimator, the next two rows of Table 5 present the approach of Eaton and Kortum (2002). Though not surprising given the discussion above, both approaches generate estimates of θ around 13 which is significantly (in its economic meaning) higher than the true value of θ of 8.28.

An interesting feature of our estimator is that it is robust to using either the first or second order statistic over prices. The bottom panel of Table 5 illustrates this point. Unseing the second order statistic, the SMM estimator yields an estimate of 8.35 with a standard error of 0.21—consistent with the true value of θ . While alternative approaches using the second order statistic result in estimates that increase from around 12 to around 15.

Figure 1 further summarizes these results by plotting the loss function, $y(\theta)' \mathbf{W} y(\theta)$, for different values of θ . Note that the minimum of the loss function lies in the ballpark of the true value of θ . In contrast, the least squares estimate lies to the right of the minimum of the loss function and the true value of θ . Also plotted is the loss function $y(\theta)' y(\theta)$ which simply sets \mathbf{W} equal to the identity matrix.

We view these results as evidence supporting our estimation approach and empirical estimate of θ presented in Section 6.




Figure 1: Loss Functions: $y(\theta)' \mathbf{W} y(\theta)$ and $y(\theta)' y(\theta)$, and θ , True $\theta = 8.28$

6 Empirical Results

In this section, we apply our estimation strategy described in section 5 to several different data sets. The key finding of this section is that our estimation approach yields an estimate around 4.5 in contrast to previous estimation strategies which yield estimates around 8.

6.1 Baseline Results Using New ICP 2005 Data

6.1.1 New ICP 2005 Data

Our sample contains 123 countries. We use trade flows and production data for the year 2004 to construct trade shares. The price data used to compute aggregate price indices and proxies for trade costs comes from basic heading level data from the 2005 round of the International Comparison Programme (ICP). The ICP collects price data on goods with identical characteristics across retail locations in the participating countries during the 2003-2005 period.¹¹ The basic heading level represents a narrowly-defined group of goods for which expenditure data are available. In the data set there are a total of 129 basic headings and we reduce it to 62 based on its correspondence with the trade data employed. Appendix 10 provides more details.

¹¹The ICP Methodological Handbook is available at <http://go.worldbank.org/MW520NNFK0>.

On its own this data set provides two contributions to the existing analysis. First, because this is the latest round of the ICP the measurement issues are probably less severe than previous rounds. Furthermore, this data set includes both developed and developing countries and allows us to study questions regarding how the elasticity of trade may vary depending upon the level of development.

6.1.2 Results—New ICP 2005 Data

Table 6 presents the results.

Table 6: Estimation Results With 2005 ICP Data

Estimation Approach	Estimate of θ	Standard Error
First Order Statistic		
SMM	4.57	0.08
Least Squares	7.35	0.03
Method of Moments	7.75	0.03
Second Order Statistic		
SMM	4.63	???
Least Squares	9.24	0.04
Method of Moments	9.61	0.04

The top panel reports results using the first order statistic and the bottom panel reports the results using the second order statistic. In both instances, our estimation procedure delivers estimates of around 4.6 with a fairly small standard error.¹² This is in contrast to estimates using the Eaton and Kortum (2002) approach which deliver estimates of 7.5 to 9.5 depending upon if the first order statistic or second order statistic is used.

6.2 Estimates Using Eaton and Kortum’s (2002) Data

In this section, we apply our estimation strategy to the same data used in Eaton and Kortum (2002) as another check of our estimation procedure. Furthermore, because it includes only OECD countries it allows us to preliminarily consider if estimates from developed countries differ than estimates using data with developed and developing countries.

¹²We are still computing standard errors because the routine to create standard errors takes substantial time.

6.2.1 Eaton and Kortum's (2002) Data

Their data set consists of bilateral trade data for 19 OECD countries in 1990 and 50 prices of manufactured goods for all countries. The prices come from an earlier round of the ICP which considered only OECD countries. Similar to our data, the price data is at the basic heading level and is for goods with identical characteristics across retail locations in the participating countries.

6.2.2 Results—Eaton and Kortum's (2002) Data

Table 7 presents the results. The top panel reports results using the first order statistic and the bottom panel reports the results using the second order statistic. In both cases, our estimation strategy generates results substantially below previous estimates; 3.6 relative to 5ish numbers when using the first order statistic. 4.5 relative to 8ish numbers when using the second order statistic. In all cases, the standard errors are fairly tight.

Table 7: Estimation Results With EK (2002) Data

Estimation Approach	Estimate of θ	Standard Error
First Order Statistic		
SMM	3.57	0.13
Least Squares	5.22	0.15
Method of Moments	5.93	0.15
Second Order Statistic		
SMM	4.46	0.14
Least Squares	8.03	0.18
Method of Moments	8.28	0.18

It is interesting to note that the estimate using the first versus the second order statistic differ substantially. This is in contrast to the monte-carlo evidence that suggests the estimation procedure should not deliver different estimates depending upon if the first or second order statistic are different. Furthermore, the results using new ICP data (Section 6.1.2) also bore this out, i.e. similar estimates using the first or second order statistic. This suggests perhaps there really is a problem with measurement error in the data as Eaton and Kortum (2002) suggested.

6.3 Estimates Using Additional Data Sources

We hope to extend our analysis to two additional data sources. The first is a data set provided by the EIU Worldwide Cost of Living Survey, which features a large subset of the original 123 countries we consider. More importantly, the data comprises of 228 tradable price observations per country, among which we observe 105 products whose prices are recorded once in a supermarket and once in a mid-price store in each country. We can use this additional dimension of the data to check whether our estimates are potentially biased by the presence of retail mark-ups. In particular, we intend to repeat our exercise by first using the prices of items collected in the mid-price store, which appears to be cheaper on average, and then the prices found in supermarkets.

The second data set we plan to explore is the data set of Waugh (2009). He employed an earlier round of the ICP data that included developing and developed countries to arrive at an estimate of θ using the same approach as Eaton and Kortum (2002). Hence his estimate is subject to the same critique we have outlined here.

TO BE COMPLETED ...

6.4 Discussion

Our estimation results compare favorably with alternative estimates of θ which *do not* use the max over price data to approximate trade costs. For example, estimates of θ using firm level data as in Bernard, Eaton, Jensen, and Kortum (2003) and Eaton, Kortum, and Kramarz (2008) are in the range of 3.6 to 4.8—exactly in the range of values we find. Eaton and Kortum (2002) provide an alternative estimate of θ using wage data and find a value of 3.6. Burstein and Vogel (2009) estimate θ matching moments regarding the skill intensity of trade and find a value of 4.

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

Finally, note that the estimates of θ when only OECD countries are considered (Eaton and Kortum’s (2002) data) are similar to our baseline with a large number of developed and developing countries. This evidence is suggestive that θ does not vary systematically across countries depending upon the level of development of the country.

7 Robustness

In this section, we discuss a number of potential issues of concern.

7.1 The Number of Goods

The estimation routine requires us to take a stand on the actual number of goods in the economy. This is a potential issue because if there were only 50 goods in the economy and we had 50 prices of each of these goods, than there would be no problem with existing estimation approaches. Clearly, there are a large number of goods in an economy. However, what the exact number of goods is is an impossible number to discipline. Instead, we argue that our estimates are not sensitive to the particular number of goods chosen as long as the number of goods is reasonably large.

Table 8: Results With EK (2002) Data, Second Order Statistic, Different # of Goods

Estimation Approach	Estimate of θ	Standard Error
SMM (300,000 Goods)	4.52	???
SMM (50,000 Goods)	4.55	???
SMM (Baseline)	4.50	???
Method of Moments	8.28	0.18
Least Squares	8.03	0.18

Table 8 presents the results as we both increase and decrease the number of goods in the simulated economy and estimated θ using the data of Eaton and Kortum (2002). The top row presents the results with 300,000 goods and the estimate is 4.52, of the same magnitude as 4.50 in the baseline case. Decreasing the number of goods to 50,000 generates similar results as well with the estimate of 4.55. We view these results as suggesting the number of assumed number of goods—as long as it is large—does not quantitatively affect our estimates.

7.2 Country-Specific Taxes and Distribution Costs

The price data used in our estimation is collected at the retail level. As such, it necessarily reflects local (distribution) costs and sales taxes. It turns out that these market frictions do not affect our estimates of the elasticity parameter, for as long as they are country- but not good-specific. To see this, suppose consumers in destination n must pay a marginal sales tax $\tau_n - 1$

on each product. Alternatively, $\tau_n - 1$ can also be thought of as a destination-specific marginal retail cost. Under these assumptions, the price (inclusive of taxes) a consumer in destination n pays for product j , $p_n^T(j)$, within the context of the model of Eaton and Kortum (2002) becomes:

$$p_n^T(j) = \tau_n \min_{k=1, \dots, N} \{p_{nk}(j)\}.$$

Substituting the pricing rule into the productivity distribution allows us to obtain the following price index for each destination n :

$$p_n^T = \tau_n \gamma \left[\sum_{k=1}^N T_k (\tau_{nk} w_k)^{-\theta} \right]^{-\frac{1}{\theta}}. \quad (20)$$

The expression for a trade share remains unchanged as all products sold in destination n are taxed uniformly:

$$\frac{X_{ni}}{X_n} = \frac{T_i (\tau_{ni} w_i)^{-\theta}}{\sum_{k=1}^N T_k (\tau_{nk} w_k)^{-\theta}}. \quad (21)$$

Expressions (20) and (21) yield:

$$\frac{X_{ni}/X_n}{X_{ii}/X_i} = \left(\frac{\tau_{ni} p_i^T / \tau_i}{p_n^T / \tau_n} \right)^{-\theta} = \left(\frac{\tau_{ni} p_i}{p_n} \right)^{-\theta}, \quad (22)$$

which is equivalent to expression (3).

So, from the model's perspective, sales taxes should not affect estimates of the key parameter.

In order to estimate the parameter θ , however, we must first arrive at a measure of trade frictions. If the price data we observe include sales taxes, the measured trade friction exporters from i face in order to serve destination n also reflects these taxes:

$$\hat{\tau}_{ni} = \frac{\tau_n}{\tau_i} \max_{\ell} \left\{ \frac{p_n(\ell)}{p_i(\ell)} \right\}. \quad (23)$$

Using $\hat{\tau}_{ni}$ in (22) would necessarily change the estimate of θ , should the pre-tax price indices, p_i , be used. However, if we use the observed price indices, which include taxes, together with $\hat{\tau}_{ni}$, expression (22) becomes:

$$\frac{X_{ni}/X_n}{X_{ii}/X_i} = \left(\frac{\hat{\tau}_{ni} p_i^T}{p_n^T} \right)^{-\theta} = \left(\frac{\tau_n \tau_{ni} \tau_i p_i}{\tau_i \tau_n p_n} \right)^{-\theta},$$

which reduces to (3).

Sales taxes that appear in observed price data are completely offset by the estimated trade barriers using these data, thus yielding identical estimates of the elasticity parameter as in the benchmark model. Hence, the presence of local taxes or distribution costs do not bias our estimates of the elasticity of trade.

7.3 Mark-ups

7.3.1 Data Approach

The price data used in our estimation likely reflects retailers' mark-ups. We address these issues in the subsection below within the context of a trade model. However, there is a potential robustness exercise that we are able to perform even with the use of a different data source. We obtain price data provided by the EIU Worldwide Cost of Living Survey, which features a large subset of the original 123 countries we consider. More importantly, the data comprises of 228 tradable price observations per country, among which we observe 105 products whose prices are recorded once in a supermarket and once in a mid-price store in each country. We can use this additional dimension of the data to check whether our estimates are potentially biased by the presence of retail mark-ups. In particular, we intend to repeat our exercise by first using the prices of items collected in the mid-price store, which appears to be cheaper on average, and then the prices found in supermarkets.

7.3.2 Model Approach

Since prices of identical goods appear to vary across the types of stores in which they are collected, one would naturally argue that different retailers charge different mark-ups for identical goods. For example, mid-price stores may charge lower prices for identical goods relative to supermarkets, because they target lower-income consumers. Similarly, firms might exercise such price discriminating rules across countries with different incomes.

In order to study how prices of identical goods behave in such environments, it is useful to consider the model proposed by Simonovska (2009). This model proposes the following non-homothetic specification of consumer preferences within a monopolistic competition framework featuring heterogeneous productivity firms:

$$U = \int_{j \in J} \log(c(j) + \bar{c}) dj,$$

where J is the set of available varieties and $\bar{c} > 0$ is a country-neutral constant.

In this framework, a firm from country i with a productivity draw $z_i(j)$ sells its product j in destination n at the following price:

$$p_{ni}(j) = \left(\frac{z_i(j)}{z_{ni}^*} \right)^{\frac{1}{2}} \frac{\tau_{ni} w_i}{z_i(j)},$$

where z_{ni}^* denotes the minimum productivity a firm from country i must have in order to serve destination n , an object that reflects local characteristics of each market such as size, income and trade frictions. First, notice that two firms with different productivities charge different mark-ups for their products. This allows us to interpret the retailers in the data discussed above as two firms that sell intrinsically differentiated products: for example, the same bottle of water sold in an expensive supermarket can be thought of as a higher-quality good because consumers derive higher utility from their shopping experience there. Hence, it is natural to treat identical products sold in different types of stores as differentiated varieties and to carry out the exercise proposed above.

Moreover, in this model, each firm charges variable mark-ups across destinations whose productivity cutoffs differ, $z_{ni}^* \neq z_{ki}^*$ for $k \neq n$. These good- and country-specific mark-ups would then be reflected in the cross-country price data we observe. However, consider the relative price of an identical product j , produced by a firm with productivity draw $z_i(j)$, sold domestically in market i , and also exported (after incurring marginal shipping cost τ_{ni}) to destination n :

$$\frac{p_{ni}(j)}{p_{ii}(j)} = \left(\frac{z_{ii}^*}{z_{ni}^*} \right)^{\frac{1}{2}} \tau_{ni}.$$

The relative price no longer accounts for the firm's productivity and only reflects market-specific conditions, in addition to trade barriers. However, as shown in the previous section, market-specific conditions do not affect the estimates of the elasticity parameters. Hence, variable mark-ups do not bias the estimates of θ .¹³

8 Why This Matters: The Welfare Gains From Trade

The elasticity parameter θ is key in measuring the welfare gains from trade across all models outlined in this paper. Arkolakis, Costinot, and Rodríguez-Clare (2009) argue that the percentage change in real income necessary to compensate a representative consumer for going

¹³It is important to note however that frameworks that model consumer preferences to be non-homothetic do not give rise to a CES price index. Thus, if one wants to derive a structural relationship between trade flows, prices and trade frictions, one must use a geometric price index, rather than the exact price index arising from the model.

to autarky—or the welfare cost of autarky—is uniquely measured by the share of domestic expenditure in a country and the elasticity of trade parameter.

To understand the argument, recall that all models outlined above rely on a CES representative consumer specification. Hence, welfare gains from trade are essentially captured by changes in the CES price index a representative consumer faces. Unfortunately, data necessary to construct pre- and post-trade CES price indices is unavailable, as we emphasize throughout the text. However, the models generate the following relationship between (unobservable) changes in price indices and (observable) changes in domestic expenditure shares as well as the elasticity parameter:

$$\frac{P'_n}{P_n} - 1 = 1 - \left(\frac{X'_{nn}/X'_n}{X_{nn}/X_n} \right)^{\frac{1}{\theta}}, \quad (24)$$

where the left-hand side can be interpreted as the percentage compensation a representative consumer requires to move from a trade to an autarky equilibrium. Notice that trade liberalization episodes, which imply a relative decrease in the domestic expenditure share of a country, necessarily generate welfare gains by lowering the price index in the particular country.

It is fairly easy to demonstrate that (24) implies that θ represents the inverse of the elasticity of welfare with respect to domestic expenditure shares:

$$\log(P_n) = -\frac{1}{\theta} \log \left(\frac{X_{nn}}{X_n} \right) \quad (25)$$

Hence, decreasing (increasing) the domestic (imported) expenditure share by 1% generates $(1/\theta)/100$ percent increase in consumer welfare. Using the estimates for θ arising from the simple procedure and the improved simulated method of moments procedure, roughly 8 and 4, respectively, the welfare gains from trade would be mis-measured by a hundred percent. Namely, an estimate for θ of 8 would generate 0.125% welfare increase for a percent rise in the import share, while an estimate of 4 suggests a 0.25% welfare gain from trade, twice as high as the original calculation. These striking differences illustrate the importance to arrive at ever better estimates of the elasticity of trade.

9 Conclusion

The methodology in our paper has broader implications than merely arriving at a better estimate of the elasticity of trade. Results from Arkolakis, Costinot, and Rodríguez-Clare (2009) suggest that heterogenous firm and production models provide no value added for aggregate

outcomes over models which abstract from heterogeneity. Our methodological approach suggests otherwise. In this paper, we exploited the structure of the Eaton and Kortum (2002) model to provide a better estimate of the elasticity of trade which is the key parameter to measuring the welfare gains from trade. Our approach would not have been possible in models without heterogeneous outcomes. Thus while the Eaton and Kortum (2002) model may provide no new additional gains from trade, its structure allowed us to provide a better elasticity of trade than a simple Armington model would have allowed. The ability to use both measurement and theory in ways that alternative models would not allow is an important component of the value added that new heterogeneous firm and production models of international trade provide.

References

- ANDERSON, J., AND E. VAN WINCOOP (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93, 170–192.
- (2004): "Trade Costs," *Journal of Economic Literature*, 42(3), 691–751.
- ARKOLAKIS, C., A. COSTINOT, AND A. RODRÍGUEZ-CLARE (2009): "New Trade Models, Same Old Gains?," *unpublished paper*, MIT.
- BAIER, S. L., AND J. H. BERGSTRAND (2001): "The growth of world trade: tariffs, transport costs, and income similarity," *Journal of International Economics*, 53(1), 1–27.
- BERNARD, A., J. EATON, J. B. JENSEN, AND S. KORTUM (2003): "Plants and Productivity in International Trade," *American Economic Review*, 93(4), 1268–1290.
- BURSTEIN, A., AND J. VOGEL (2009): "Globalization, Technology, and the Skill Premium," Discussion paper.
- CHANEY, T. (2008): "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, 98(4), 1707–1721.
- DAVIDSON, R., AND J. MACKINNON (2004): *Econometric Theory and Methods*. New York: Oxford University Press.
- DONALDSON, D. (2009): "Railroads of the Raj: Estimating the Economic Impact of Transportation Infrastructure," *unpublished paper*, MIT.
- EATON, J., AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70(5), 1741–1779.
- EATON, J., S. KORTUM, AND F. KRAMARZ (2008): "An Anatomy of International Trade: Evidence from French Firms," NBER Working Paper 14610,.
- 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.
- KRUGMAN, P. (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70(5), 950–959.
- MACKINNON, J. G., AND A. A. SMITH (1998): "Approximate bias correction in econometrics," *Journal of Econometrics*, 85(2), 205–230.

- MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725.
- MUENDLER, M.-A. (2009): "Converter from SITC to ISIC," *University of California - San Diego, unpublished mimeo*.
- SIMONOVSKA, I. (2009): "Income Differences and Prices of Tradables," *University of California - Davis, unpublished mimeo*.
- SMITH, A. (2008): "Indirect inference," *The New Palgrave Dictionary of Economics, 2nd Edition (forthcoming)*.
- WAUGH, M. E. (2009): "International Trade and Income Differences," Staff Report 435, Federal Reserve Bank of Minneapolis.
- YI, K.-M. (2003): "Can Vertical Specialization Explain the Growth of World Trade?," *Journal of Political Economy*, 111(1), 52–102.

10 Data Appendix

10.1 Trade Shares

To construct trade shares, we used bilateral trade flows and production data in the following way:

$$\frac{X_{ni}}{X_n} = \frac{\text{Imports}_{ni}}{\text{Gross Mfg. Production}_n - \text{Total Exports}_n + \text{Imports}_n},$$
$$\frac{X_{nn}}{X_n} = 1 - \sum_{k \neq n}^N \frac{X_{nk}}{X_n}.$$

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

To construct $\frac{X_{ni}}{X_n}$, the numerator is the aggregate value of manufactured goods that country n imports from country i . Bilateral trade flow data are from UN Comtrade for the year 2004. We obtain all bilateral trade flows for our sample of 123 countries at the four-digit SITC level. We then used concordance tables between four-digit SITC and three-digit ISIC codes provided by the UN and further modified by Muendler (2009).¹⁴ We restrict our analysis to manufacturing bilateral trade flows only, namely, those that correspond with manufactures as defined in ISIC Rev.#2.

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

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

¹⁴The trade data often report bilateral trade flows from two sources. For example, the exports of country A to country B can appear in the UN Comtrade data as exports reported by country A or as imports reported by country B. In this case, we take the report of bilateral trade flows between countries A and B that yields a higher total volume of trade across the sum of all SITC four-digit categories.

where β_x is a 1×3 vector of coefficients corresponding to C_x , an $N \times 3$ matrix which contains $[\log(x), (\log(x))^2, (\log(x))^3]$ for the sub-sample of N countries for which gross output data are available.

10.2 Prices

The ICP price data we employ in our estimation procedure is reported at the basic-heading level. Here we discuss briefly how these prices are collected. An issue we discuss is that the prices in the data are aggregates over even more detailed products. In our estimation routine we abstracted from this issue. However, we should emphasize that a key advantage of our simulated method of moments procedure is that these aggregation problems can be explicitly addressed.

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

According to the ICP Handbook, the price of the basic heading "Rice" is constructed using a transitive Jevons index of prices of different varieties of rice. To illustrate this point, suppose the world economy consists of 3 countries, A, B, C and 10 types of rice, 1-10. Further suppose that consumers in country 1 have access to all 10 types of rice; those in country 2 only have access to types 1-5 of rice; and those in country 3 have access to types 4-6 of rice. Although all types of rice are not found in all 3 countries, it is sufficient that each pair of countries shares at least one type of rice.

The ICP obtains unit prices for all available types of rice in all three countries and records a price of 0 if the type of rice is not available in a particular country. The relative price of rice between countries 1 and 2, based on goods available in these two countries, $p_{AB}^{A,B}$, is a geometric average of the relative prices of rice of types 1 – 5

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

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

by the ICP is:

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

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

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