

Trade Models, Trade Elasticities, and the Gains from Trade

Ina Simonovska
University of California, Davis and NBER

Michael E. Waugh
New York University and NBER

This Version: September 2014

ABSTRACT

We argue that the welfare gains from trade in new models with micro-level margins exceed those in frameworks without these margins. Theoretically, we show that for fixed trade elasticity, different models predict identical trade flows, but different patterns of micro-level price variation. Thus, given data on trade flows and micro-level prices, different models have different implied trade elasticities and welfare gains. Empirically, models with extensive or variable mark-up margins yield significantly larger welfare gains. The results are robust to incorporating into the estimation moment conditions that use trade-flow and tariff data, which imply a common trade elasticity across models.

JEL Classification: F10, F11, F14, F17

Keywords: elasticity of trade, bilateral, gravity, price dispersion, welfare gains

Email: inasimonovska@ucdavis.edu, mwaugh@stern.nyu.edu.

We thank seminar participants at NBER ITI Summer Institute, Ca' Foscari, Rochester and Brown. A very preliminary version of this paper circulated under the title "Different Trade Models, Different Trade Elasticities?".

1. Introduction

This paper argues both theoretically and empirically that the welfare gains from trade in new trade models with various micro-level margins are larger relative to models without these margins. In an important class of trade models, we show that for fixed trade elasticity, different models have different implications for micro-level price variation, even though their predictions for aggregate trade are identical. These facts imply that, given data on aggregate trade and micro-level prices, different assumptions about the underlying model result in different trade elasticities and consequently different welfare gains from trade. Empirically, we quantify these differences and we find significantly larger welfare gains in new models with various micro-level margins versus old models without these margins.

Our theoretical analysis focuses on three canonical models of trade: Anderson (1979) (henceforth Armington), Eaton and Kortum (2002) (henceforth EK), and Bernard, Eaton, Jensen, and Kortum (2003) (henceforth BEJK). Analyzing these three models is insightful because each model adds an additional micro-level margin of adjustment from a reduction in trade costs. In particular, Armington features an intensive margin only (i.e. reductions in trade costs lead to higher cross-border purchases of previously traded goods). EK features an intensive and extensive margin (i.e. reductions in trade costs further lead to additional goods being traded). Finally, BEJK further adds to EK a variable mark-up margin. All three models fit into the class of models examined by Arkolakis, Costinot, and Rodriguez-Clare (2012) where the trade elasticity and the share of expenditure on domestic goods are sufficient statistics to measure the welfare cost of autarky.

To understand the link between trade elasticities and micro-level prices, we examine the theoretical distribution of price gaps of identical goods across countries in a symmetric, two-country version of each model. Assuming a fixed trade elasticity, we show that the price gap distribution in the Armington model stochastically dominates that in the EK model, which further dominates the one in the BEJK model. The ranking of these distributions implies a similar ranking in the expectation of the largest order statistic of price gaps from each model. In turn, this statistic is inversely related to the trade elasticity. Hence, to match an observed order statistic in the data, the Armington model requires the highest trade elasticity, while the BEJK model needs the lowest. Because the analysis constrains all models to generate identical aggregate trade flows, the elasticity ranking result together with the sufficient statistic formula of Arkolakis et al. (2012) implies that the welfare cost of autarky is different across models. In particular, the BEJK model yields the highest, while the Armington model yields the lowest welfare gains from trade.

We quantify the importance of the theoretical results by estimating the trade elasticities in the multi-country asymmetric version of each model. The particular estimation approach that we

use builds both on the theory described above and on our earlier work in Simonovska and Waugh (2014) which focused on the specifics of the EK model. The basic idea behind our estimation strategy is to choose the trade elasticity to match moments between the model and the data about order statistics of bilateral price gaps. While we focus on a broader set of models and moments than in Simonovska and Waugh (2014), the methodological approach is similar in spirit. First, we estimate the parameters of the models necessary to simulate micro-level price data using bilateral trade-flow data, which guarantees that all models have identical aggregate trade predictions. Second, we use these parameter estimates and a given trade elasticity to simulate micro-level prices from each model and to construct the model-implied moments. We then choose the trade elasticity to minimize the distance between the moments in each model and the data.

Using price and trade-flow data for the year 2004 for the 30 largest countries in terms of gross output, we estimate trade elasticities for each of the three models under various specifications including exactly-identified and overidentified specifications using different weighting matrices. Across all specifications, the estimate of the trade elasticity is systematically lower for BEJK relative to EK and EK relative to the Armington model. The difference in magnitudes is substantial. The EK estimate is about 20 percent lower than Armington, implying that the welfare cost of autarky is 20 percent higher in the model that features an extensive margin. In comparison, the estimate of the trade elasticity in the BEJK model is about 33 percent lower relative to EK, implying the welfare gains are estimated to be 50 percent larger in a model that features an extensive and variable mark-up margin relative to the benchmark model that only captures an intensive margin of trade.

We further examine versions of the two canonical endogenous variety models: the monopolistic competition framework of Krugman (1980) and the framework of Melitz (2003) which has an extensive margin of trade with productivity-based firm selection. We apply our estimation procedure on the two frameworks and obtain trade elasticity estimates in the Melitz model that are 30 percent lower than Krugman, resulting in 30 percent higher welfare gains from trade in the former. The result is due to presence of an extensive margin of trade in Melitz, but not in Krugman.

A natural question is why we focus on micro-level price variation as a means to estimating trade elasticities. The focus on micro-level price variation is important for several reasons. First, emphasis on prices is important because all models make concrete and distinct predictions about price variation, e.g. in contrast, moments about the firm/sales size distribution are indeterminate in the Armington and EK model.

Second, alternative approaches to estimating the trade elasticity do not appear to be very informative relative to our “micro” approach. There are estimation approaches that utilize aggregate

data and relationships between model and data that are straightforward to implement and are common across models (see, e.g., the discussion in Arkolakis et al. (2012) and Caliendo and Parro (2014)). These are not invalid approaches and provide additional evidence on the elasticity of trade.

We show, however, that the aggregate moments employed by these alternative approaches are not informative relative to our approach. To demonstrate this point, we carry out a joint “micro”- and “macro”-based estimation of the trade elasticity, which, respectively, combines our moment conditions with those utilized by Caliendo and Parro (2014) (i.e. triple differenced trade flows and tariffs, and a model-independent orthogonality condition). Thus, depending upon aspects of the data, the combined estimation procedure has the ability to deliver an estimate of the trade elasticity that is common across models.

We find that the macro moment conditions have no impact on our estimates (see, e.g. third column of Table 7). The reason is that cross-country tariff data is extremely noisy and has virtually no explanatory power for trade flows. Thus, any loss in fit due to the failure to satisfy the linear best fit between trade flows and tariff-elasticity estimates is trivial; hence, tariff-based moment conditions have little power relative to our “micro” approach.

The final reason to focus on micro-level price variation is that it speaks directly to the economic mechanisms at work in each model. The active extensive margin in the EK or Melitz model or variable markups in BEJK all manifest themselves in different patterns of micro-level price variation. Thus, to learn about the importance of these margins, one should focus on data that these margins can speak to. Moreover, even if aggregate, model-independent approaches were informative about the trade elasticity, these approaches lack the ability to discriminate across models and are uninformative about the mechanisms at work in the data. We make some preliminary steps in this direction by finding that Armington/Krugman and BEJK do not describe well (in a formal statistical sense) the micro-level price patterns found in the data.

A closely related paper to ours is the work by Melitz and Redding (2014), who focus on the subset of endogenous variety models that we examine. Their welfare ranking between the Krugman and the Melitz model is identical to ours. The argument, however, differs. Keeping other parameters fixed, Melitz and Redding (2014) argue that the Melitz model generates higher trade shares than the Krugman model, thus yielding higher welfare gains from trade. Instead, we estimate the trade-elasticity parameters of the two models subject to the restriction that the models generate *identical* trade shares—the ones observed in the data. Hence, our approach is most closely related to Arkolakis et al.’s (2012). Like Arkolakis et al. (2012), we keep trade shares fixed, since all models generate identical predictions along that dimension. Unlike Arkolakis et al. (2012), we do not assume an identical trade elasticity across models; rather, we allow data to determine the appropriate trade elasticity for each model. Our welfare conclusions, then,

differ from Arkolakis et al.'s (2012) because trade elasticities turn out to systematically vary across models due to the differences in the models' predictions about prices.

2. Fixed Variety Models of Trade

2.1. Armington

The simplest model of international trade that yields a gravity equation is the Armington model outlined in Anderson and van Wincoop (2003). The framework features N countries populated by consumers with constant elasticity of substitution (CES) preferences and tradable goods that are differentiated by the country of origin. Perfect competition among producers and product differentiation by origin imply that each good is purchased in each destination at a price that equals the marginal cost of production and delivery of the good there.

We consider a more empirically-relevant version of the Armington model that features product differentiation within and across countries.¹ Throughout the paper, let i denote the source country and n the destination. We allow each country to produce an exogenously given measure of tradable goods equal to $1/N$. We assume that products are differentiated. Within each country n , there is a measure of consumers L_n . 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 $\rho > 1$,²

$$U_n = \left[\sum_{i=1}^N \int_0^{1/N} x_{ni}(j)^{\frac{\rho-1}{\rho}} dj \right]^{\frac{\rho}{\rho-1}}.$$

To produce quantity $x_{ni}(j)$ in country i , a firm employs labor using a linear production function with productivity $T_i^{1/\theta}$, where T_i is a country-specific technology parameter and $\theta = \rho - 1$. The perfectly competitive firm from country i incurs a marginal cost to produce good j of $w_i/T_i^{1/\theta}$, 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 , $\tau_{ni} \leq \tau_{nk}\tau_{ki}$. We maintain this assumption about the nature of iceberg trade costs in all the models that we describe below.

Perfect competition forces the price of good j from country i to destination n to be equal to the

¹This model yields identical aggregate predictions to the simple Armington model, but it accommodates a greater number of goods than countries, which is an observation that holds true in the data.

² $j \in [0, 1/N]$ is an index for an individual good. Since goods are differentiated by country of origin, a good should be denoted by j, i ; however, we suppress the i notation for brevity as prices and quantities are denoted by origin i and destination n .

marginal cost of production and delivery

$$p_{ni}(j) = \frac{\tau_{ni}w_i}{T_i^{1/\theta}}.$$

Since goods are differentiated, consumers in destination n buy all products from all sources and pay $p_{ni}(j)$ for good j from i . Hence, all tradable goods are traded among the N markets and consumers buy a unit measure of goods.

2.2. Eaton and Kortum (2002)

We now outline the environment of the multi-country Ricardian model of trade introduced by Eaton and Kortum (2002)—EK. As in the Armington model, there is a continuum of tradable goods indexed by $j \in [0, 1]$. Preferences are represented by the following utility function

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

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

$$F_{\mathcal{EK},i}(z_i) = \exp(-T_i z_i^{-\theta}).$$

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 $w_i/z_i(j)$. 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

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

Hence, the EK framework introduces endogenous tradability into the Armington model outlined above. In particular, countries export only a subset of the unit measure of tradable goods for which they are the most efficient suppliers.

2.3. Bernard, Eaton, Jensen, and Kortum (2003)

Bernard et al. (2003)—BEJK—introduce Bertrand competition into EK’s model. The most important implication from this extension is that individual good prices differ from the EK model.

Let $c_{kni}(j) \equiv \tau_{ni}w_i/z_{ki}(j)$ be the cost that the k -th most efficient producer of good j in country i faces in order to deliver a unit of the good to destination n . With Bertrand competition, as with perfect competition, the low-cost supplier of each good serves the market. For good j in market n , this supplier has the following cost $c_{1n}(j) = \min_i \{c_{1ni}(j)\}$. This supplier is constrained not to charge more than the second-lowest cost of supplying the market, which is $c_{2n} = \min \{c_{2ni^*}(j), \min_{i \neq i^*} \{c_{1ni}(j)\}\}$, where i^* satisfies $c_{1ni^*}(j) = c_{1n}(j)$. Hence, the price of good j in market n is

$$p_n(j) = \min \{c_{2n}(j), \bar{m}c_{1n}(j)\},$$

where $\bar{m} = \rho/(\rho - 1)$ is the Dixit-Stiglitz constant mark-up.

Finally, for each country i , productivity, $z_{ki}(j)$ for $k = 1, 2$ is drawn from

$$F_{BEJK,i}(z_1, z_2) = [1 + T_i(z_2^{-\theta} - z_1^{-\theta})] \exp(-T_i z_2^{-\theta}).$$

Hence, the BEJK model features a key additional component relative to the EK framework: the existence of variable mark-ups. In particular, the most efficient suppliers in this model enjoy the highest mark-ups.

2.4. Trade Flows, Aggregate Prices, and Welfare

The models described above produce identical aggregate outcomes, even though they feature different micro-level behavior. In particular, under the parametric assumption made above, the models yield the same expressions for trade flows, price indices (up to a constant scalar), and welfare gains from trade. Proposition 1 summarizes the result.

Proposition 1 *Given the functional forms for productivity $F_{EK,i}(\cdot)$, $F_{BEJK,i}(\cdot)$ for all $i = 1, \dots, N$,*

a. The share of expenditures that n spends on goods from i , X_{ni}/X_n , predicted by the models is

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

b. The CES exact price index for destination n , P_n , predicted by the models is

$$P_n \propto \Phi_n^{-\frac{1}{\theta}}, \quad \text{where} \quad \Phi_n = \sum_{k=1}^N T_k (\tau_{nk} w_k)^{-\theta}. \quad (2)$$

c. The percentage compensation that a representative consumer in n requires to move between two trading equilibria predicted by the models is

$$\frac{P'_n}{P_n} - 1 = 1 - \left(\frac{X'_{nn}/X'_n}{X_{nn}/X_n} \right)^{\frac{1}{\theta}}. \quad (3)$$

We prove a. and b. for the Armington model in the Appendix. The results for the EK and BEJK models are derived in the respective papers. The proof of part c. can be found in Arkolakis et al. (2012). Across the models, the welfare gains from trade are essentially captured by changes in the CES price index that a representative consumer faces. Using the objects from the Proposition above, it is easy to relate the price indices to trade shares and the parameter θ . In particular, expressions (1) and (2) allow us to relate trade shares to trade costs and the price indices of each trading partner via the following equation

$$\frac{X_{ni}/X_n}{X_{ii}/X_i} = \frac{\Phi_i}{\Phi_n} \tau_{ni}^{-\theta} = \left(\frac{P_i \tau_{ni}}{P_n} \right)^{-\theta}, \quad (4)$$

where $\frac{X_{ii}}{X_i}$ is country i 's expenditure share on goods from country i , or its home trade share. The welfare equation follows trivially from this expression.

2.5. The Elasticity of Trade

The key parameter determining trade flows (equation (4)) and welfare (equation (3)) is θ . To see the parameter's importance for trade flows, take logs of equation (4) yielding

$$\log \left(\frac{X_{ni}/X_n}{X_{ii}/X_i} \right) = -\theta [\log(\tau_{ni}) - \log(P_i) + \log(P_n)]. \quad (5)$$

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 (3) implies that θ

represents the inverse of the elasticity of welfare with respect to domestic expenditure shares

$$\log(P_n) = -\frac{1}{\theta} \log\left(\frac{X_{nn}}{X_n}\right). \quad (6)$$

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.

The challenge with estimating the trade elasticity is that one must *separately* disentangle θ from trade costs, which are not observed. To overcome this challenge, we theoretically show that micro-level cross-country price variation or price gaps identifies the trade elasticity. More importantly, different models represent different mappings between price gaps and the trade elasticity, which implies that in order to match the observed price variation in the data, different models require different trade elasticities.

3. Price Gaps and Trade Elasticities

This section shows that different trade models have different implications for the distribution of relative price gaps. This implies that, given data on trade flows and micro-level prices, different models have different trade elasticities and thus different welfare gains from trade.

To illustrate the argument theoretically, we focus on a symmetric, two country world with the countries denoted as home (h) and foreign (f). In a symmetric two-country world, the only two parameters that govern trade flows (as well as prices) in the models above are θ and τ .³ Throughout this section, we assume that θ and τ are constrained to fit the bilateral trade share. This implies that if $\theta > \theta'$, then $\tau < \tau'$.

Given the constraint that all models must imply the same bilateral trade flows, we then make the following argument. First, we describe properties of the distribution of relative prices at the micro-level. Second, we order these distributions (across models and for different θ s) by first-order stochastic dominance and, in turn, by the largest order statistic or maximal price gap in a finite sample. This then allows us to order the implied θ s across models given the observed largest order statistic from a sample of data.

³To see this, from the first equality in expression (4) and the definition of Φ_n in (2), notice that, given trade share data and a value for θ , $\log \tau$ can be computed immediately in a symmetric two-country world.

3.1. Ordering Price Distributions and the Largest Order Statistic Across Models

In this section, we focus on deviations from the law of one price at the micro level, i.e. for an individual variety. Specifically, we focus on the logged price gap between the home and the foreign country:

Definition 1 For any micro-level good ℓ , define the random variable that equals the logged price gap between the home and the foreign country as $\log \tilde{P}_{h,\ell} \equiv \log P_{h,\ell} - \log P_{f,\ell}$.

Throughout, we drop the subscript ℓ for brevity, and we only use it when it is necessary to distinguish prices across goods.

The distribution of price gaps is an important object of interest. Given a model \mathcal{M} , denote the cumulative density function of log price gaps, conditional on the price gap being positive as:

$$\forall \text{ models } \mathcal{M}, \quad G_{\mathcal{M}}(\log \tilde{p}_h, \theta) = \begin{cases} 0 & : \log \tilde{p}_h < 0 \\ \in [0, 1] & : 0 \leq \log \tilde{p}_h \leq \log \tau \\ 1 & : \log \tau < \log \tilde{p}_h. \end{cases} \quad (7)$$

Note that we index the density by parameter θ . This indexing is sufficient because other parameters are either the same across models (e.g. the technology parameters), or are determined by the value θ , e.g. θ determines τ so aggregate trade is held constant.

An important property of this density is that all price gaps lie below $\log \tau$. The reason is the following: suppose that $\log P_{h,\ell} - \log P_{f,\ell} > \log \tau$, then an arbitrage opportunity exists as an agent could import good ℓ from the foreign country at a lower price. Thus, this inequality places an upper bound on the possible observable price gaps. This property is important because the support of the density depends on θ . Different θ s imply different τ s because aggregate trade constrains the values that these parameters can take. Thus, this property allows us to make conclusions as to how parts of the density shift in a model with parameter θ relative to the model with parameter θ' .

How the density shifts as θ changes at any price gap in certain models is analytically intractable (though numerically verifiable). Thus, we assert that the following regularity conditions holds on the price gap distribution.

Regularity Condition 1 (No Crossing.) Let the density $G_{\mathcal{M}}(\log \tilde{p}_h, \theta)$ satisfy the following no-crossing property. If $\theta' > \theta$, then for $\log \tilde{p}_h < \log \tau'$:

$$\text{Prob}_{\mathcal{M}}(\log \tilde{P}_h < \log \tilde{p}_h, \theta) \leq \text{Prob}_{\mathcal{M}}(\log \tilde{P}_h < \log \tilde{p}_h, \theta'). \quad (8)$$

This property says the following: Fix a model \mathcal{M} . If we lower the elasticity from θ' to θ and constrain aggregate trade to stay the same, the probability of seeing a particular price gap cannot increase. So as the upper end of the support increases from $\log \tau'$ to $\log \tau$, it can not induce a “cross” in the density.

This property clearly holds in the Armington model. In the Armington model, all goods are traded and thus the probability mass of the price gaps lies completely at $\log \tau'$. If we lower the elasticity from θ' to θ , then all the probability mass shifts rightward to $\log \tau$. For the EK model, this holds as well (see online appendix of Simonovska and Waugh (2014)). For the BEJK model, analytically verifying this is difficult because a closed-form expression for the price gap distribution does not exist; we have, however, verified that this holds numerically for all relevant parameter values.

Below, we define the data features we will focus on. Specifically, the maximal price gap from a sample of price gaps:

Definition 2 Consider a finite, random sample of positive price gaps $\log \tilde{P}_{h,1}, \log \tilde{P}_{h,2}, \dots, \log \tilde{P}_{h,L}$ which are ordered from smallest to largest; The maximal price gap in the sample of L goods is $\log \tilde{P}_{h,L:L}$. Given a model \mathcal{M} , parameter θ , and density $G_{\mathcal{M}}(\log \tilde{p}_h, \theta)$, the expected maximal price gap or largest order statistic is $\mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{\theta, \mathcal{M}\} \right)$, where \mathcal{M} indexes the model and the dependence on the parameter θ is noted.

Focusing on the maximal price gap has several appealing features. First, it has some history of thought as it has been used in Eaton and Kortum (2002) and Simonovska and Waugh (2014) in the estimation of θ . Second, order statistics encode much information about the underlying density. For example, one can show that a sequence of largest order statistics (that is for $L = 1, 2, \dots$) completely characterizes the density G (see, e.g., Arnold, Balakrishnan, and Nagaraja (1992)). We will never have access to an infinite sequence of largest order statistics. However, by utilizing recurrence relationships between largest order statistics of different sample sizes (again see Arnold et al. (1992)), one can construct testable restrictions for each model.

Lemma 1 connects the expected maximal price gap with a stochastic dominance relationship in the density.

Lemma 1 If the price density, $G_{\mathcal{M}}(\log \tilde{p}_h, \theta)$, first-order stochastically dominates the price density $G_{\mathcal{M}'}(\log \tilde{p}_h, \theta')$, then $\mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{\theta, \mathcal{M}\} \right) > \mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{\theta', \mathcal{M}'\} \right)$.

Proof: If a random variable first-order stochastically dominates another, the same ranking holds for the distribution of the variables’ maxima, that is $G_{\mathcal{M}}(\log \tilde{p}_h, \theta)^L >_{fosd} G_{\mathcal{M}'}(\log \tilde{p}_h, \theta')^L$. This property implies that the expected maximal price gap under model \mathcal{M} is larger than in model \mathcal{M}' . ■

The result here is that if one density dominates another density, this implies that the maximal price gap must be larger in the model that has the density which dominates. Application of Lemma 1 and the regularity condition above allows us to rank densities according to θ for a given model \mathcal{M} .

Lemma 2 *For a given model, \mathcal{M} , the parameter θ indexes the price density by first-order stochastic dominance, and strictly indexes the expected maximal price gap. That is, if $\theta' > \theta$ then $G_{\mathcal{M}}(\log \tilde{p}_h, \theta) >_{f_{osd}} G_{\mathcal{M}}(\log \tilde{p}_h, \theta')$ and $\mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{\theta, \mathcal{M}\} \right) > \mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{\theta', \mathcal{M}\} \right)$.*

Proof: First, recall that, when $\log \tau$ and θ are constrained to fit bilateral trade shares, $\theta' > \theta$ implies that $\tau' < \tau$. Then, $\tau' < \tau$ and Regularity Condition 1 implies that, for $\log \tilde{p}_h < \log \tau'$, the cumulative probability in the model with θ is weakly less than the cumulative probability in the model with θ' . For $\log \tilde{p}_h = \log \tau'$, the cumulative probability in the model with θ' is equal to one, and because there is strictly positive probability mass in the region $\log \tilde{p}_h > \log \tau'$ for the model with θ , this implies a strict inequality in probability mass at $\log \tilde{p}_h = \log \tau'$. This implies first-order stochastic dominance. From Lemma 1, the ordering of the order statistics follows. ■

Lemma 2 allows us, for a given model, \mathcal{M} , an ability to strictly rank expected maximal price gaps by θ . One implication of this result is that the expected maximal price gap identifies the θ parameter. We exploit this result in Proposition 3. The second implication is that it assists us in showing that if all models have the same expected maximal price gap, then they must have different θ s. We state the result formally below:

Proposition 2 *Consider the models $G_{\mathcal{ARM}}(\log \tilde{p}_h, \theta_{\mathcal{ARM}})$, $G_{\mathcal{EK}}(\log \tilde{p}_h, \theta_{\mathcal{EK}})$, $G_{\mathcal{BEJK}}(\log \tilde{p}_h, \theta_{\mathcal{BEJK}})$. If the expected maximal price gaps are the same in all models:*

$$\mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{\theta, \mathcal{M}\} \right) = \mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{\theta', \mathcal{M}'\} \right), \quad \forall \text{ models } \mathcal{M},$$

then

$$\theta_{\mathcal{ARM}} > \theta_{\mathcal{EK}} \text{ and } \theta_{\mathcal{ARM}} > \theta_{\mathcal{BEJK}}.$$

Moreover, if $G_{\mathcal{EK}}(\log \tilde{p}_h, \theta) >_{f_{osd}} G_{\mathcal{BEJK}}(\log \tilde{p}_h, \theta)$, then

$$\theta_{\mathcal{ARM}} > \theta_{\mathcal{EK}} > \theta_{\mathcal{BEJK}}.$$

Proof: The proof proceeds by contradiction in the following two cases with the focus on the EK model. The same argument applies with respect to the BEJK model.

- Case #1. Suppose not and that $\theta_{\mathcal{ARM}} = \theta_{\mathcal{EK}} = \theta$. First, note that for the same θ ,

$$G_{\mathcal{ARM}}(\log \tilde{p}_h, \theta) >_{fbsd} G_{\mathcal{EK}}(\log \tilde{p}_h, \theta), \quad (9)$$

because for any $\log \tilde{p}_h < \log \tau$, the probability mass in the Armington model is zero as all goods are traded and hence all positive price differences are equal to the trade friction. This is strictly less than in the EK model as for any $\log \tilde{p}_h < \log \tau$, the probability mass is greater than zero because there are non-traded goods. This argument implies the distribution of price gaps in the Armington model first-order stochastically dominates the distribution in the EK model. Application of Lemma 1 then implies the expected maximal price gaps should be ranked, which is a contradiction.

- Case #2. Suppose not and that $\theta_{\mathcal{EK}} > \theta_{\mathcal{ARM}}$. First, notice that this implies that $G_{\mathcal{ARM}}(\log \tilde{p}_h, \theta_{\mathcal{ARM}}) >_{fbsd} G_{\mathcal{ARM}}(\log \tilde{p}_h, \theta_{\mathcal{EK}})$ from Lemma 2. This then implies that,

$$\mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{ \theta_{\mathcal{ARM}}, \mathcal{ARM} \} \right) > \mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{ \theta_{\mathcal{EK}}, \mathcal{ARM} \} \right) > \mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{ \theta_{\mathcal{EK}}, \mathcal{EK} \} \right)$$

and the last inequality follows from the observation made in Case #1 above, that for the same θ , the Armington price gap distribution strictly dominates the EK price gap distribution. This implication contradicts the assumption that the expected maximal price gap is the same in all models.

Finally note that the same exact arguments apply for the BEJK model, where some goods are non-traded as in the EK model. Thus, if the Armington model, EK model, and BEJK model, all have the same expected maximal price gap, then the elasticity in the Armington model must be strictly greater than the elasticity in EK or BEJK. ■

Figure 1 illustrates the intuition behind this result. The top panel plots the cumulative distribution function of log price gaps for the Armington (solid blue line), the EK (dashed red line) and the BEJK (solid black line) model—for the same θ .

In the Armington model, notice that all the mass lies on $\log \tau$. This is because all goods are traded, i.e. there is no active extensive margin in the Armington model. If the home country imports the good from the foreign country, then the price gap exactly reflects the trade friction. The distribution of price gaps in the EK model is different as it places a positive mass in the region between zero and the trade friction and hence the density in EK is stochastically dominated by that of Armington. This is because there are endogenously non-traded goods, i.e. there is an active extensive margin. In the EK model, there are instances where consumers in *both* countries find it cheaper to consume a variety from the local producer rather than importing the good. If a good is non-traded, the price difference across locations reflects differences in costs which are strictly less than the trade friction. This implies a positive mass in the region





Figure 1: Price Gap Distribution: Armington, EK, BEJK.

between zero and the trade friction.

Proposition 2 says that because the price gap density in Armington dominates the density in EK, the only way for the expected maximal price gap to be the same is if the EK model has a lower θ . The bottom panel of Figure 1 illustrates this by plotting the densities when they have the same expected largest order-statistic. Here, the EK density is spread to the left with a larger upper bound to ensure that the largest order statistic is the same as in the Armington model.

What about the relationship between EK and BEJK? Generically, we can say the following: if the density in EK strictly stochastically dominates the density in BEJK and the largest order statistics are the same, then by logic of Proposition 2 this implies that $\theta_{EK} > \theta_{BEJK}$. A stronger statement is currently not possible as it is difficult to theoretically construct the BEJK price gap distribution because of the presence of variable markups.⁴

The top panel in Figure 1 shows that for the same θ , the price gap density in EK dominates the density in BEJK. More generally we have verified that this relationship holds numerically for all relevant parameter values. And our estimates of θ in a multi-country setting with asymmetries always result in a ranking of $\theta_{EK} > \theta_{BEJK}$. Furthermore, the bottom panel of Figure 1 verifies that if the largest order statistic is the same in EK and BEJK, then the BEJK model must have a lower theta. This can be seen as the density in BEJK is spread even further to the left with a larger upper bound to ensure that the largest order statistic equals that in the EK and Armington models.

Variable markups and how they correlate with costs in the BEJK model explain why the price gap density in EK strictly dominates the density in BEJK. In the EK model, price gaps correspond to cost gaps and these cost gaps are identical in the two models. In BEJK, cost gaps do not correspond with price gaps as producers are able to price at a markup over marginal cost depending on other latent competitors.⁵ Thus, the price gap in BEJK reflects both markups, cost differences (if the good is non-traded), and trade frictions (if the good is traded). In BEJK, markups are negatively correlated with marginal costs. This results in the BEJK distribution being more compressed than the EK distribution.⁶

Consider the following example to illustrate this point. Take a producer in the foreign country who has a very high productivity and is the lowest cost producer in his country of origin as well as the home country. This producer will likely charge a high markup (call it m') in the foreign country because she has a very low cost relative to her latent competitors as that is

⁴In the Appendix, we show that, in the case in which $\rho = 1$ and the Dixit-Stiglitz mark-up is infinite, the probability that the log price gap reaches the boundaries is strictly lower in the BEJK than in the EK model. This implies that the mass has to be distributed between these two end-points—meaning, in the non-traded good region. However, the exact shape of the distribution cannot be easily characterized.

⁵See de Blas and Russ (2010) for discussion of the importance of the number of latent competitors in BEJK.

⁶A related property is that it is straightforward to show that the EK price distribution (not price gaps) is a mean preserving spread of the BEJK distribution.

her country of origin. In contrast, she will likely charge a relatively lower markup (call it m'') in the home country because her cost advantage is eroded from the trade friction faced when exporting to the home country. This implies that the price gap, $\log \tilde{p}_h$, will be *less than* the trade friction even though the good is traded, i.e. the price gap equals $m'' - m' + \log \tau < \log \tau$. This observation implies that the BEJK distribution will have less mass around the log of the trade friction because of producers differentially marking up their products across locations.

3.2. Estimating Trade Elasticities and the Gains from Trade

We now connect Proposition 2 with an econometrician's inference about θ given a sample of data (i.e. aggregate trade flows and a sample of micro-level prices) and discuss how this inference depends upon assumptions about the underlying model. Proposition 3 suggests a simple method of moments estimator and states that the estimated θ will (i) depend on the econometrician's assumptions about the underlying data generating process and (ii) will be ordered across models.

Proposition 3 *Consider the following estimator of θ that chooses $\hat{\theta}$ to minimize the distance between the observed largest order statistic $\log \tilde{P}_{h,L:L}$ and the expected largest order statistic $\mathbb{E}(\log \tilde{P}_{h,L:L}; \{\theta, \mathcal{M}\})$. That is,*

$$\hat{\theta}_{\mathcal{M}} = \arg \min_{\theta} h(\{\theta, \mathcal{M}\})' h(\{\theta, \mathcal{M}\})$$

$$\text{where } h(\{\theta, \mathcal{M}\}) = \left(\log \tilde{P}_{h,L:L} - \mathbb{E} \left(\log \tilde{P}_{h,L:L}; \{\theta, \mathcal{M}\} \right) \right).$$

Then:

1. The estimate $\hat{\theta}_{AR\mathcal{M}}$ is strictly greater than the estimate $\hat{\theta}_{EK}$ and $\hat{\theta}_{BEJK}$.
2. The estimate of the welfare cost of autarky in the Armington model is strictly less than the estimate of the welfare cost of autarky in EK and BEJK.
3. If the density in EK stochastically dominates the density in BEJK, then the estimate $\hat{\theta}_{EK} > \hat{\theta}_{BEJK}$ and the welfare gains from trade in EK are strictly less than the gains in BEJK.

Proof: The first statement follows from Lemma 2 which implies that the order statistic uniquely determines the θ for a given model and Proposition 2 which implies that across models, the θ must differ. The second statement follows from the fact that equation (6) is the same across models and that, by construction, all models have the same predictions for aggregate trade flows, but different θ . ■

A couple of comments are in order about this result. First, the estimator is quite simple. However, our estimation strategy in Section 4 builds off the intuition from this result. Moreover,

we include alternative moment conditions to enrich the estimation, to be able to test over-identifying restrictions, and to exploit empirical regularities we see in the data. We discuss these moment conditions in Section 4.

Second, compare Proposition 3 to the discussion in Section 6 in Arkolakis et al. (2012). In this section, they propose an alternative estimator of the trade elasticity that delivers a common estimate of θ across models. Their estimator uses aggregate trade flows and some measure of trade frictions (say tariffs) and the common connection between aggregate trade flows (i.e. equation (5)) that all these models share. The final piece of their argument is to *assume* that the orthogonality condition that the estimate of θ is based on is model independent. In other words, their analog to the expectation of the function $h(\{\theta, \mathcal{M}\})$ is assumed not to depend on the model \mathcal{M} . Thus, the estimate of θ as well as the gains from trade will be the same across models.

The key distinction between our argument and that of Arkolakis et al. (2012) is that the orthogonality condition we focus on is model dependent. Moreover, the structure of each model implies a specific ranking as to how the estimate of θ should relate to each other across models. Ultimately, however, how different margins of trade affect estimates of the welfare gains from trade is an empirical question which we turn to next.

4. Estimating the Elasticity

This section describes how we estimate the trade elasticity given a data set that features micro-level prices and bilateral trade flows across countries. The basic idea is captured in Proposition 3: we use moment conditions that compare statistics from a sample of micro-level prices with their expected values from each model, where the latter are functions of θ and depend on assumptions about the model that generated the data. Because the expected values do not have closed form expressions, we use simulation methods to approximate the expected value and we discuss how we are able to simulate prices from the different models. Finally, we provide monte-carlo evidence that our estimation procedure works well.

4.1. Estimation

We will focus on three sets of results (i) an exactly identified case, (ii) an overidentified case with an identity weighting matrix, and (iii) an overidentified case with an optimal weighting matrix. Below we discuss each of these cases in turn.

Exactly Identified Case. In the exactly identified case we focus on properties of the logarithm of the maximal price gap adjusted by average log prices in the two locations. We define the

value d_{ni} as

$$d_{ni} = \log \hat{\tau}_{ni} + \log \hat{P}_i - \log \hat{P}_n, \quad (10)$$

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

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

and L is the number of micro-level prices in the sample. Like the statistics discussed in Section 3, we focus initially on the maximal price gap.

One difference from the previous section is that equation (10) adjusts the maximal price gap by average price differences in the two locations. There are empirical and theoretical reasons for this adjustment. Empirically, adjusting for price levels in each country helps us control for country specific factors that may be present in the price data but not in our model.⁷ Theoretically, if one wants to use the maximal price gap as a proxy for trade frictions and connect this value with trade flows, then expression (5) suggests adjusting for the difference in Φ s across countries as well. It turns out that the difference in the means of logged prices equals the difference in the logs of the parameters Φ defined in expression (2).⁸

Given the d_{ni} s from data, we define the following function

$$h(\theta, d_{ni}, \mathcal{M}) = d_{ni} - \frac{1}{S} \sum_{s=1}^S d_{ni}(\theta, u_s, \mathcal{M}), \quad (11)$$

which compares the observed value d_{ni} to its expected value. Because we do not have a closed form expression for the expected value of d_{ni} , $\mathbb{E}(d_{ni}; \{\theta, \mathcal{M}\})$, we approximate this value via simulation. The next section provides the exact details behind a simulation, but the basic idea is the following: given a model \mathcal{M} and a value of θ , we simulate prices, construct s synthetic data sets, and then construct simulated values for $d_{ni}(\theta, u_s, \mathcal{M})$. The expected value of d_{ni} is the average across simulations S .

⁷For example, suppose there are country-specific factors such as sales taxes, which are multiples of the goods' prices. Then, the maximum log price difference between countries n and i will reflect the log differences in the country-specific sales taxes. The tax wedge washes out once we add the difference in the means of logged prices between countries i and n .

⁸In Simonovska and Waugh (2014) we formally prove the result for the EK model. A similar result can be obtained for the Armington and BEJK models where price levels are also proportional to the price parameters as demonstrated in Proposition 1 above.

In the exactly identified case, our estimation is based on the following orthogonality restriction

$$\mathbb{E}(h(\theta_o, d_{ni}, \mathcal{M})) = 0. \quad (12)$$

That is, at the true value θ_o , in expectation the difference between the observed d_{ni} and its expected value should be equal to zero. The sample average of $h(\theta, d_{ni}, \mathcal{M})$ is

$$\bar{h}(\theta, \mathcal{M}) = \left(\frac{1}{N^2 - N} \right) \sum_n \sum_i h(\theta, d_{ni}, \mathcal{M}) \quad (13)$$

where $N^2 - N$ is the number of d_{ni} s we can construct. Our estimate of θ is the value that minimizes the quadratic form of $\bar{h}(\theta, \mathcal{M})$. Or mathematically:

$$\hat{\theta}(\mathcal{M}) = \arg \min_{\theta} \bar{h}(\theta, \mathcal{M})' \bar{h}(\theta, \mathcal{M}). \quad (14)$$

Overidentified Case. The overidentified case focuses on two additional moments: (i) the price gap in the 85th percentile adjusted by average prices and (ii) the covariance of d_{ni} with the logarithm of bilateral distance.

The focus on the price gap in the 85th percentile is a way to incorporate more information about the underlying distribution of price gaps and, hence, information about the underlying model. For example, as discussed in Section 3, the two country Armington model places a restriction on the data that the max should be the same as the price gap in the 85th percentile. The price gap in the 85th percentile is not the only statistic that has this property. We explored other percentiles (e.g. 90 or 75) and found virtually no effect on our results.

The covariance between the d_{ni} statistic and the logarithm of distance was chosen for two reasons. First, as we show in Section 5.1, there is a strong correlation between maximal price gaps and distance. Given the strong role that distance plays in explaining trade flows, we feel this is a natural statistic to target. Second, this also helps guard against estimating the elasticity off of measurement error in the data by focusing on the systematic relationship between price gaps and distance.

Given these moments, the function $\mathbf{h}(\theta, \mathbf{d}_{ni}, \mathcal{M})$ is now a 3×1 vector. The first element of \mathbf{d}_{ni} is d_{ni} as defined in expression (10) above, the second is the price gap in the 85th percentile, and the third is the covariance between d_{ni} and the logarithm of bilateral distance. We are slightly abusing notation here by using the bold-face values \mathbf{d}_{ni} to denote a vector of data. Our orthogonality restriction is

$$\mathbb{E}(\mathbf{h}(\theta_o, \mathbf{d}_{ni}, \mathcal{M})) = 0. \quad (15)$$

The sample average of (each element of) $\mathbf{h}(\theta, \mathbf{d}_{ni}, \mathcal{M})$ is

$$\bar{\mathbf{h}}(\theta, \mathcal{M}) = \left(\frac{1}{N^2 - N} \right) \sum_n \sum_i \mathbf{h}(\theta, \mathbf{d}_{ni}, \mathcal{M}), \quad (16)$$

and we chose θ to minimize the quadratic form

$$\hat{\theta}(\mathcal{M}) = \arg \min_{\theta} \bar{\mathbf{h}}(\theta, \mathcal{M})' \mathbf{W}^{-1} \bar{\mathbf{h}}(\theta, \mathcal{M}), \quad (17)$$

where \mathbf{W} is a positive definite weighting matrix.

For the weighting matrix, we will present two sets of results: one using the identity matrix and another using the optimal weighting matrix suggested by Gouriéroux and Monfort (1996) and described in Adda and Cooper (2003). Because the optimal weighting matrix depends on our estimate of θ , we use a continuous-updating estimator of W which continually updates the weighting matrix within the minimization routine (see e.g. the discussion of this estimator in Hansen, Heaton, and Yaron (1996)).

We compute standard errors using a parametric bootstrap technique (see e.g. Davison and Hinkley (1997)). We add error terms to the trade data (given our parametric assumption on the error term in (19)) and re-estimate the country-specific parameters from the simulated data necessary to simulate prices. Given these parameters, we simulate a sample of prices given our estimate of θ and the assumed underlying model, and then implement our estimation routine on the simulated data. Thus, this procedure takes into account uncertainty in the technology parameters and scaled trade costs estimated from trade data and sampling variability in the micro-level prices. We repeat this procedure 100 times to construct 90-10 percentile confidence intervals.

In the instances where we use an optimal weighting matrix, we also perform tests of overidentifying restrictions. Specifically, we construct the test statistic

$$(N^2 - N) \bar{\mathbf{h}}(\hat{\theta}, \mathcal{M})' \mathbf{W}^{-1} \bar{\mathbf{h}}(\hat{\theta}, \mathcal{M}) \quad (18)$$

which is the standard “J-statistic” used to test the null-hypothesis that the data generating process is correct. Asymptotically, this test statistic is distributed chi-squared with a certain number of degrees of freedom, though in small samples this may not provide an accurate approximation (see e.g. Hansen et al. (1996)). To avoid these difficulties, we use a parametric bootstrap to construct the finite sample distribution of the test statistic in (18) under the null hypothesis that the data generating process is correct (see Davison and Hinkley (1997)).

BEJK Details. One detail associated with the BEJK model is that the value of ρ , the CES parameter, matters unlike in the EK model. The issue is that ρ determines the monopoly pricing rule

which shapes the distribution of price gaps. To deal with this issue we proceed in two ways. The first approach is to pre-calibrate this parameter to the value of 2.5 (which is consistent with evidence regarding this parameter in Broda and Weinstein (2006)) and estimate only θ . The second approach is to estimate ρ and θ in the BEJK model. Specifically, in the overidentified case, we pick θ and ρ to jointly minimize the quadratic form in (17).

Relation to Previous Work. The statistic d_{ni} that we use to estimate θ is developing a history of thought, so a couple of comments are necessary. First, EK used this statistic to directly estimate θ . If the maximal price difference accurately revealed the trade friction, then by examining (4) one can see how averaging across normalized trade shares relative to the average d_{ni} yields an estimate of θ . Using this method, they arrived at the value 8.28. As we demonstrate in Simonovska and Waugh (2014), the maximal price gap generally underestimates the trade friction, and hence overestimates the value of θ . In Simonovska and Waugh (2014), we use indirect inference to correct this bias by matching the observed EK estimate of θ to the model implied value. Using this method, we arrived at the value around four.

We depart from our previous work for the following reasons. First, the new procedure fits transparently within the GMM framework which allows us to formally test the overidentifying restrictions of various models and allows us to incorporate additional moment conditions that existing work uses to estimate the trade elasticity (see e.g. Caliendo and Parro (2014)). Second, we avoid a difficulty that arose when using both simulated trade flows and prices to construct the model implied EK estimate. One difficulty was that the simulated trade flows depended on the number of goods in the economy and for the simulated trade flows to mimic observed trade flows accurately this had to be a very large number. By just simulating prices, we avoid the dependence on the number of goods in the economy and achieve a very large speed up in computation time. Finally, in the exactly identified case our new approach gives nearly the exact same answer for the EK model as the approach used in Simonovska and Waugh (2014), which is reassuring.

Why these Moments. The discussion above provided a brief, tactical rationale for why we chose the particular moments we did. We have carefully chosen the moments, explored a wide variety of alternatives (such as different percentiles, order statistics from different-sized samples via recurrence relationships, etc.) and found similar results.

There are, however, deeper, strategic issues regarding the choice of moments that are worth pointing out. One may wonder why we do not use other micro moments such as the firm size/sales distribution. The reasoning is simple: Moments about the firm/sales size distribution are indeterminate in the Armington and EK models. In contrast, all models make concrete and distinct predictions about price variation; hence, moments from the price distribution are informative.

One may also wonder why we do not find and then use information regarding the existence of traded versus non-traded goods for the identification of trade elasticities. If we could find a good where the arbitrage relationship bound and under the assumption perfect competition, this good’s price difference could be used to estimate a trade friction, and consequently, a trade elasticity. While this information may help, it does not help one answer the central question: Are trade elasticities (and therefore welfare gains) different across models? For example, if the true data generating process is BEJK, then prices reflect variable mark-ups, so cross-country price differences in *traded* goods are not necessarily equal to trade costs as discussed in Section 3. Furthermore, focusing on traded goods does not inform us about the underlying data generating process. As the arguments in Section 3 suggest and Figure 1 makes clear, the data features that distinguish different models are price gaps *away* from the trade costs. Moreover, the economic mechanisms (margins of trade) behind each of these models shape the pattern of price gaps *away* from the trade costs. This is the exact rationale for focusing on moments from the price gap distribution that are away from the trade cost, such as the 85th percentile of the price gap.

4.2. Simulation Approach

A key step in the estimation procedure is the approximation of $\mathbb{E}(d_{ni}; \{\theta, \mathcal{M}\})$ via simulation. Below we describe in detail how we simulate prices from each model and in turn how we approximate $\mathbb{E}(d_{ni}; \{\theta, \mathcal{M}\})$.

There are basically three steps. The distribution of prices depends on θ and the country-specific technology parameters and bilateral trade costs. We use the gravity equation implied by all models (derived in expression (19) below) to estimate the technology parameters and trade costs (up to a scalar θ) from trade flow data. This ensures that all models will have the same aggregate predictions for trade flows. Second, we show how the estimates from the gravity equation are sufficient to simulate micro-level prices given a specified model. Third, we specify a protocol as to how the data is sampled. These three steps allow us to create synthetic data sets from which we can then approximate $\mathbb{E}(d_{ni}; \{\theta, \mathcal{M}\})$.

Step 1.—We estimate the parameters for the country-specific productivity distributions and trade costs from bilateral trade-flow data. We follow closely the methodologies proposed by Eaton and Kortum (2002) and Waugh (2010). First, we derive the gravity equation from expression (1) by dividing the bilateral trade share by the importing country’s home trade share,

$$\log \left(\frac{X_{ni}/X_n}{X_{nn}/X_n} \right) = S_i - S_n - \theta \log \tau_{ni} + \nu_{ni}, \quad (19)$$

where S_i is defined as $\log [T_i w_i^{-\theta}]$. S_i ’s are recovered as the coefficients on country-specific dummy variables given the restrictions on how trade costs can covary across countries. We as-

sume that ν_{ni} reflects other factors and is orthogonal to the regressors and normally distributed with mean zero and standard deviation σ_ν . Following the arguments of Waugh (2010), trade costs take the following functional form

$$\log(\tau_{ni}) = \rho \log m_{ni} + b_{ni} + ex_i. \quad (20)$$

Here, trade costs are a logarithmic function of distance, where m_{ni} is the distance in miles between country n and i and ρ is the distance elasticity. 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. Waugh (2010) shows that including this term helps trade models match both cross-country variation in aggregate prices and trade flows. As we showed in Simonovska and Waugh (2014), using importer fixed effects (as in EK) or aggregate price data to exactly identify trade costs did not change our estimates for the EK model by economically meaningful amounts. Finally, our results are robust to incorporating bilateral colonial, language, and legal origin ties as well as countries' geographical attributes.

The term ex_i is an exporter fixed effect and allows for the trade-cost level to vary depending upon the exporter. We use least squares to estimate equations (19) and (20).

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

Step 2a.—The simulation of the Armington model is simple as there is no micro-level heterogeneity in this framework. $\exp(S_i^{-\theta})$ represents the common marginal cost of production across all producers in country i .

Step 2b.—To simulate micro-level prices in the EK model we follow Simonovska and Waugh (2014). 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 Simonovska and Waugh (2014), we show that u_i is distributed according to:

$$M_i(u_i) = \exp\left(-\tilde{S}_i u_i^{-\theta}\right), \quad \text{with } \tilde{S}_i = \exp(S_i) = T_i w_i^{-\theta}. \quad (21)$$

Thus, having obtained estimates of S_i from the gravity regression, we can simulate the inverse of marginal costs. To construct prices at which goods are traded, we take the inverse marginal costs that are drawn from the country-specific distributions above and are assigned to each good. Then, for each importing country and each good, the lowest-cost supplier across countries is found and realized prices are recorded.

Step 2c.—To see that we can simulate micro-level prices in the BEJK model as a function of θ

only, we draw on an argument in BEJK. To simulate their model, BEJK reformulate the model in terms of efficiency. In particular, given two productivity draws for good j in country i , $z_{1i}(j)$ and $z_{2i}(j)$, they define the following objects

$$\begin{aligned} u_{1i}(j) &= T_i z_{1i}(j)^{-\theta} \\ u_{2i}(j) &= T_i z_{2i}(j)^{-\theta} \end{aligned}$$

BEJK demonstrate that these objects are distributed according to

$$\begin{aligned} Pr[u_{1i} \leq u_1] &= 1 - \exp(-u_1) \\ Pr[u_{2i} \leq u_2 | u_{1i} = u_1] &= 1 - \exp(-u_2 + u_1) \end{aligned}$$

To simulate trade flows and prices from the model, we define the following variables

$$\begin{aligned} v_{1i}(j) &= \left(\frac{u_{1i}(j)}{T_i w_i^{-\theta}} \right) \\ v_{2i}(j) &= \left(\frac{u_{2i}(j)}{T_i w_i^{-\theta}} \right) \end{aligned}$$

Applying the pdf transformation rule, it is easy to demonstrate that $v_{1i}(j)$ is distributed according to

$$Pr[v_{1i} \leq v_1] = 1 - \exp(-\tilde{S}_i v_1). \quad (22)$$

Similarly,

$$Pr[v_{2i} \leq v_2 | v_{1i} = v_1] = 1 - \exp(-\tilde{S}_i v_2 + \tilde{S}_i v_1), \quad (23)$$

where \tilde{S}_i is defined above.

Thus, to simulate the model we draw minimum unit costs from (22), and conditional on these draws, we draw the second lowest unit costs from (23). Hence, having obtained the estimates of S_i from the first-stage gravity regression, we can simulate the inverse of marginal costs.

To construct prices at which goods are traded, we take the inverse marginal costs that are drawn from the country-specific distributions above and are assigned to each good. Then, for each importing country and each good, the two lowest-cost suppliers across countries are found and realized prices are recorded.

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. The three models give a natural common bas-

ket of goods to be priced across countries. In these models, agents in all countries consume all goods that lie within a fixed interval, $[0, 1]$. Thus, we consider this common list in the simulated models and we randomly sample the prices of its goods across countries.

These steps then provide us with an artificial data set of micro-level prices 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.

4.3. Performance on Simulated Data

To provide evidence that the estimation procedure recovers the underlying parameter for each model, we apply the methodology on data simulated by each model under a known θ . For exposition purposes, we let θ be the same across the models, and we set it equal to four.

Table 2: Estimation Results With Artificial Data, Underlying $\theta = 4$

	Model	Estimate of θ	Mean d_{ni}	Cov($d_{ni}, \log m_{ni}$)	Mean d_{ni}^{85th}
Exactly Identified	Armington	3.99 [3.83, 4.26]	1.27	—	—
	EK	3.99 [3.80, 4.18]	0.96	—	—
	BEJK	3.99 [3.84, 4.15]	0.63	—	—
Over Identified	Armington	3.97 [3.82, 4.19]	1.21	0.20	0.24
	EK	3.99 [3.90, 4.10]	0.96	0.10	0.37
	BEJK	4.00 [3.88, 4.10]	0.63	0.04	0.25

Note: Value is the median estimate across simulations. Values within brackets report 90th and 10th percentiles. In each simulation there are 30 countries and 100 simulations are performed. Over Identified uses the optimal weighting matrix; using the identity matrix yielded near identical results.

The first column in Table 2 reports the results. Our estimation routine recovers the value of θ in all three models, both in the exactly identified and overidentified case. Foreshadowing our empirical findings, the final three columns in Table 2 report the average moments across simulations. Consistent with the theoretical results in Section 3, there is a systematic ranking in largest price gaps with Armington being the highest and BEJK being the lowest. Consistent with Proposition 2, this implies that given observed price gaps, the models will be ordered by θ to match this data.

5. Empirical Results

5.1. Data Description

Our sample contains the thirty largest countries in the world (in terms of absorption). We use trade flows and production data for the year 2004 to construct trade shares. The price data used to compute moments on price gaps and aggregate price levels come from basic-heading-level data from the 2005 round of the International Comparison Programme (ICP). The dataset has been employed in a number of empirical studies. For example, Bradford (2003) and Bradford and Lawrence (2004) use the ICP price data in order to measure the degree of fragmentation, or the level of trade barriers, among OECD countries. In addition, the authors provide an excellent description of the data-collection process. Eaton and Kortum (2002) use a similar dataset for the year 1990 in their estimation.⁹

The ICP collects price data on goods with identical characteristics across retail locations in the participating countries during the 2003-2005 period.¹⁰ The basic-heading level represents a narrowly-defined group of goods for which expenditure data are available. The data set contains a total of 129 basic headings, and we reduce the sample to 62 categories based on their correspondence with the trade data employed. Simonovska and Waugh (2014) provide a more detailed description of the ICP data.

The ICP provides a common list of “representative” goods whose prices are to be randomly sampled in each country over a certain period of time. A good is representative of a country if it comprises a significant share of a typical consumer’s bundle there. Thus, the ICP samples the prices of a common basket of goods across countries, where the goods have been pre-selected due to their highly informative content for the purpose of international comparisons. It is precisely this sampling procedure that we mimic in the simulation methodology described above.

Table 3 reports how price gaps in the ICP data covary with typical gravity variables that determine observed trade patterns. The first row regresses the object d_{ni} on the logged distance between country pairs and a dummy variable for contiguity. Clearly, maximum price gaps are rising in distance and are lower for countries that share a border. The second row repeats the exercise using the object that relies of the 85th percentile of price gaps. While the qualitative patterns are identical, the magnitudes are lower. Given the systematic relationship between

⁹For robustness, we re-estimate trade elasticities using an alternative dataset: 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 110 tradable goods, 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. Consistent with our findings in Simonovska and Waugh (2014), we obtain lower estimates of the trade elasticity for all models, but the ranking across models remains unchanged. These results are available upon request.

¹⁰The ICP Methodological Handbook is available at <http://go.worldbank.org/MW520NNFK0>.

Table 3: Price Gaps and Trade Determinants

Dependent Variable	Log Distance	Border	Mean Value
d_{ni}	0.13 [0.11, 0.15]	-0.11 [-0.18, -0.02]	2.5
d_{ni}^{85th}	0.055 [0.047, 0.063]	-0.035 [-0.060, -0.010]	1.4

Note: The first two columns report regression of dependent variables on distance, border indicator and country fixed effects. The last column reports the mean value, exponentiated.

price gaps and gravity variables found in Table 3, we conclude that price gaps do not simply reflect measurement error, but rather they contain valuable information about trade barriers.

5.2. Results with ICP Data

Table 4 presents the results for the three specifications. The first column presents the results from the exactly identified case which most closely corresponds with our theoretical results in Proposition 3. Consistent with the theoretical results, the estimate of θ is systematically lower for the new trade models relative to the Armington model, and further the estimate for BEJK is lower than for EK. The difference in magnitudes is substantial. The EK estimate is about 20 percent lower than Armington, implying that the welfare cost of autarky is 20 percent higher in the EK model. In comparison, our estimate of the trade elasticity in the BEJK model is about 33 percent lower relative to EK, implying the welfare gains are estimated to be 50 percent larger in BEJK relative to Armington.

Table 4: Estimation Results, Fixed Variety Models

Model	Exactly Identified	Overidentified, Identity	Overidentified, Optimal	J-stat P-value
	Estimate of θ	Estimate of θ	Estimate of θ	
Armington	5.24 [5.02, 5.45]	5.06 [4.84, 5.26]	3.94 [3.73, 4.08]	< 0.01
EK	4.17 [4.00, 4.34]	4.13 [3.99, 4.29]	4.05 [3.97, 4.14]	0.25
BEJK, $\rho = 2.5$	3.32 [3.18, 3.48]	3.28 [3.16, 3.43]	3.15 [3.08, 2.96]	< 0.01
BEJK, est. $\hat{\rho}$	---	3.21 [2.90, 3.64]	2.74 [2.62, 2.89]	< 0.01

Note: Values within brackets report 90th and 10th confidence intervals. J-stat P-value reports probability of observing a J-statistic as large as observed. Estimates of ρ for BEJK model are 2.25 and 1.37.

Table 5: Data and Fixed Variety Model Moments

	Moment		
	Mean d_{ni}	Mean d_{ni}^{85}	Cov($d_{ni}, \log m_{ni}$)
Data	0.93	0.37	0.14
Armington	1.23 [0.16]	0.25 [0.02]	0.21 [0.15]
EK	0.95 [0.09]	0.37 [0.01]	0.10 [0.09]
BEJK, $\rho = 2.5$	0.99 [0.09]	0.35 [0.01]	0.09 [0.09]
BEJK, est. $\hat{\rho}$	0.95 [0.09]	0.37 [0.01]	0.07 [0.08]

Note: Reports model implied moments at estimated parameters values from the overidentified estimation with optimal weighting matrix. Values within brackets report the diagonal entry of W , the variance-covariance matrix. $\log m_{ni}$ is the log distance in miles between country n and i .

The last three columns of Table 4 present the results for the overidentified case. Note that with the weighting matrix equal to the identity matrix, the results are effectively identical to the results in the exactly identified case.

When the optimal weighting matrix is used, results do begin to change. Specifically, our estimate for the Armington model declines substantially relative to the previous two columns. This may seem problematic relative to our results in Proposition 3, yet in fact it is highly symptomatic of the fact that the Armington model is misspecified. The first and second row of Table 5 illustrate this point. The top two rows present the data and moments implied by the Armington model. Across all moments, the Armington model has problems: it overpredicts the mean order statistic by a large margin, underpredicts the 85th percentile by a large margin, and overpredicts the covariance term by a large margin.

The difficulties of the Armington model are formalized in the P-value for the J-statistic which tests the null hypothesis that the Armington model is the data generating process. For the Armington model, the null hypothesis is rejected at a level of less than one percent. Given the starkness of the Armington model, this should be a reassuring result.

For the EK model, the results with the optimal weighting matrix are nearly identical to the other results. The third row of Table 5 provides an informal sense of fit. The EK model exactly matches the 85th percentile (which is given the most weight) and it comes very close to matching the mean order statistic and the covariance term. The J-statistic formalizes this and we fail to reject the null hypothesis that the EK model is the data generating process.

Across specifications, the results for the BEJK model are similar. The one instance where the estimate changes is in the joint estimation of the CES parameter ρ and θ with the optimal weighting matrix. In this instance, our estimate of θ becomes smaller and our estimate of ρ is lower than the calibrated value of 2.5.

While the BEJK model is the richer of the three models in terms of market structure and pricing patterns, there is evidence that it is misspecified as the null hypothesis that the BEJK model generated the data is rejected at the 1 percent level. Interestingly, the issue with the BEJK model is that it has a difficult time fitting the covariance between the d_{ni} and distance. This can be seen, for example, in the last row/final column of Table 5; while the BEJK model has the ability to match both the order statistics and 85th percentile; it has trouble with the covariance term.¹¹ The problem is due to variable markups. Variable markups naturally induce “noise” between price gaps (even if the good is traded in which case price gaps reflect the distance cost), which makes it harder for the BEJK model to replicate the strong covariance observed in the data between price gaps and distance.

6. Endogenous Variety Models

The analysis thus far has focused on three canonical models of trade where the set of varieties is fixed. Two other canonical models that feature an endogenous set of varieties produced/consumed are the frameworks of Krugman (1980)—Krugman—and Melitz (2003)—Melitz. While these models can satisfy the same aggregate restrictions as the three models analyzed, we have kept them separate largely because we make additional assumptions about the structure of fixed costs and how we interpret the data when the set of varieties varies across countries.

Below, we outline versions of these two models and we repeat the quantitative exercises from above. We find that the Melitz model yields a lower estimate of the trade elasticity than the Krugman model, thus generating higher welfare gains from trade.

6.1. Krugman (1980)

Preferences are identical to the Armington model, but the measure of goods produced by any country i is endogenous,

$$W_n = \left[\sum_{i=1}^N \int_{J_i} x_{ni}(j)^{\frac{\rho-1}{\rho}} dj \right]^{\frac{\rho}{\rho-1}}.$$

In the above expression, J_i denotes the set of goods produced by country i .

¹¹More formally, we discovered this by removing the covariance moment and only estimating to the two price gaps; in this specification, we failed to reject the BEJK model.

In particular, we assume that a producer in country i has monopoly power over a blueprint to produce a single good j . In order to produce the good, a firm in i uses the following production technology

$$y_i(j) = T_i^{1/\theta} l_i(j) + e_i.$$

Thus a firm incurs a marginal cost of $w_i/T_i^{1/\theta}$ as well as a fixed cost of entry e_i , both in labor units. There is an unbounded pool of potential producers and the mass of firms is pinned down via a zero average profit condition.

We make the assumption that fixed costs are proportional to the mass of workers in country i , L_i . Under the proportionality assumption on fixed costs, the mass of entrants is identical across countries and proportional to $1/(1 + \theta)$.

Conditional on entry, all consumers buy all products from all sources. Moreover, due to the proportionality assumption on entry costs, each country contributes a measure of $1/N$ toward the unit measure of consumed products as in the Armington model. In contrast to Armington, the price $p_{ni}(j)$ for good j from i in destination n is given by the product of the marginal cost of production and delivery and the Dixit-Stiglitz constant mark-up \bar{m} .

6.2. Melitz (2003)

Finally, we outline a variant of the Melitz (2003) model parameterized as in Chaney (2008). In the exposition, we follow closely Arkolakis, Demidova, Klenow, and Rodríguez-Clare (2008).

The preference relation as well as the market structure are identical to the Krugman (1980) model described above. There are two novel features. First, in order to sell its product to market n , a firm from any country i incurs an additional market access cost, a_n , which is proportional to the mass of workers in the destination, L_n . Second, marginal costs of production are not identical across producers within a country. In particular, upon paying the entry cost e_i , a firm from i draws a productivity realization for good j , $z_i(j)$, from the country-specific Pareto distribution with pdf $T_i z^{-\theta}$. As in the Frechet distribution, the parameter θ in the Pareto distribution governs the variability of productivity among firms.

The two additional assumptions imply that there exist cost thresholds that limit the participation of certain firms to certain markets. In particular, if a firm obtains a cost draw above the threshold cost that characterizes the “most accessible” destination, it exits immediately without operating. Thus, because of free entry, in equilibrium expected profits of a firm must be equal to the fixed entry cost. In addition, more productive firms serve more markets and the “toughest” markets attract the most efficient producers from all sources.

As in the Krugman (1980) model, firms charge a constant mark-up over their marginal cost of

production. However, unlike the Krugman (1980) model, the prices of goods offered by a given country differ according to the efficiency of the individual producer. More importantly, even though entry costs are proportional to the mass of workers, the measure of firms that serve each market is no longer constant. Thus, selection among exporters alters the composition of the bundle of consumed goods.

6.3. Trade Flows, Aggregate Prices, and Welfare

In this section, we show that the two models with an endogenous set of varieties produce identical aggregate outcomes to the models with an exogenous set of varieties. In particular, under the parametric assumption made above, all five models yield the same expressions for trade flows, price indices (up to a constant scalar), and welfare gains from trade. Proposition 4 summarizes the result. The proof can be found in the Appendix.

Proposition 4 *Given the functional forms for productivity $F_{\mathcal{EK},i}(\cdot)$, $F_{\mathcal{BEJK},i}(\cdot)$, and $F_{\mathcal{MEL},i}(\cdot)$ for all $i = 1, \dots, N$, and assuming that (i) fixed entry costs are proportional to the mass of workers in the country of origin, $e_i \propto L_i$ for all $i = 1, \dots, N$; (ii) market access costs are proportional to the mass of workers in the destination market, $a_n \propto L_n$ for all $n = 1, \dots, N$;*

a. *The share of expenditures that n spends on goods from i , X_{ni}/X_n , predicted by all five models is*

$$\frac{X_{ni}}{X_n} = \frac{T_i(\tau_{ni}w_i)^{-\theta}}{\sum_{k=1}^N T_k(\tau_{nk}w_k)^{-\theta}}. \quad (24)$$

b. *The CES exact price index for destination n , P_n , predicted by all five models is*

$$P_n \propto \Phi_n^{-\frac{1}{\theta}}, \quad \text{where} \quad \Phi_n = \sum_{k=1}^N T_k(\tau_{nk}w_k)^{-\theta}. \quad (25)$$

c. *The percentage compensation that a representative consumer in n requires to move between two trading equilibria predicted by all five models is*

$$\frac{P'_n}{P_n} - 1 = 1 - \left(\frac{X'_{nn}/X'_n}{X_{nn}/X_n} \right)^{\frac{1}{\theta}}. \quad (26)$$

6.4. Simulation of Prices for Endogenous Variety Models

Since the endogenous variety models yield an identical gravity equation of trade as the exogenous variety models, the first stage of the estimation is identical. What remains is to describe the micro price-level simulation.

Step 2d.—The simulation of the Krugman model is trivial because it does not feature micro-level heterogeneity. In this model, \tilde{S}_i fully characterizes the common marginal cost of production across all producers in country i .

Step 2e.—The simulation of the Melitz (2003) model is more intricate. In the Eaton and Kortum (2002) and Bernard et al. (2003) models—which feature heterogeneity—a good is indexed by $j \in [0, 1]$ or an integer when discretized on the computer. In the Melitz (2003) model, different countries are consuming and producing different goods. To carry out the simulation, we discretize the continuum to a set of potentially produced goods. We proceed to determine the subset of goods produced domestically for each country. Then, we simulate inverse marginal costs for each good. Inverse marginal costs are determined following Eaton, Kortum, and Kramarz (2011) who show how normalized inverse marginal costs can be sampled from a parameter free uniform distribution. Then, we determine the set of exported varieties for all country-pairs and we compute their prices.

As in the exogenous variety models, we must specify how the subsample of prices is recorded. In the fixed variety models, because all goods are consumed everywhere, there is a natural common list from which we randomly sample the prices of its goods across countries. The key issue in the Melitz model is that the set of goods consumed is different across countries. Below we describe how we deal with this issue.

Step 3a.—In the Krugman model, all produced goods appear in all countries, and each country contributes a fraction $1/N$ of them. Thus, like the exogenous variety models, there is a natural common basket of goods prices to be randomly sampled from across countries.

Step 3b.—In the Melitz model, we define the common set of goods whose prices will be sampled to be the set of goods produced by firms with low enough cost draws so that they serve all N destinations (see Appendix for cost threshold derivations). Let M_i represent the measure of goods from i that appear on the common list and let $M = \sum_m M_m$ be the common list. Then, it is straightforward to derive the following expression for the share of country i 's goods on the common list

$$\frac{M_i}{M} = \frac{\min_l \tilde{S}_i \tau_{li}^{-\theta} \left[\sum_k \tilde{S}_k \tau_{lk}^{-\theta} \right]^{-1}}{\sum_m \min_l \tilde{S}_m \tau_{lm}^{-\theta} \left[\sum_k \tilde{S}_k \tau_{lk}^{-\theta} \right]^{-1}}. \quad (27)$$

Notice that this share varies across countries. This variation is precisely the reason behind the different elasticity estimates that we obtain from the two models. In the Krugman model, the maximal bilateral price difference is likely to equal the unobserved trade friction (exactly like in the Armington model). In the Melitz model, only goods sold by firms that enter all markets are able to make the sample of prices that data collectors would consider. Moreover, the firms

that enter all markets are likely to be from source countries with low trade costs to export. This implies that, in the Melitz model, the average maximal price difference across country pairs is likely to be below the average trade friction across country pairs. Thus, one needs a lower trade elasticity in Melitz relative to Krugman to rationalize the same amount of trade observed in the data. The differences in trade elasticities then translate into differences in the measured welfare gains from trade in the two models.

6.5. Results

Table 6 reports the estimates of θ for the two endogenous variety models obtained using the ICP data. The first column reports the results from the exactly identified estimation. The Melitz model yields an estimate of θ that is roughly 30 percent lower than the estimate for the Krugman model. Hence, adding an extensive margin of trade increases the welfare gains from trade by 30 percent.

Table 6: Estimation Results, Endogenous Variety Models

Model	Exactly Identified	Over-Identified, Identity	Over Identified, Optimal	
	Estimate of θ	Estimate of θ	Estimate of θ	J-stat P-value
Krugman	5.24 [5.02, 5.45]	5.06 [4.84, 5.26]	3.94 [3.73, 4.08]	< 0.01
Melitz	3.69 [3.32, 4.21]	3.69 [3.37, 4.09]	3.65 [3.36, 3.86]	0.35

Note: Values within brackets report 90th and 10th confidence intervals. J-stat P-value reports probability of observing a J-statistic as large as observed.

The last three columns report the results from the overidentified estimation. Once again a similar picture emerges as in the estimation of the exogenous variety models. The overidentified estimation that uses the optimal weighting matrix yields an estimate of θ for the model that features an extensive margin—Melitz in this case—that is roughly the same as in the overidentified case with an identity matrix and the exactly identified case. However, when using the optimal weighting matrix, the estimate of θ falls considerably for the model that features only an intensive margin of trade—in this case Krugman. Once again, the J statistic allows us to reject the hypothesis that the model with an intensive margin only has generated the data. In contrast, we cannot reject the null that the Melitz model has generated the data.

Relation to Melitz and Redding (2014). Our estimation yields a lower trade elasticity estimate for the Melitz model than the Krugman model, which translates into higher welfare gains from trade in the former. Recently, Melitz and Redding (2014) also argue that the Melitz model

generates higher welfare gains from trade than the Krugman model. Their argument, however, is very different from ours and it goes as follows. The Krugman model is a special case of the Melitz model where the distribution of firm productivities becomes degenerate. Naturally, θ is the central parameter that governs this distribution. The authors then demonstrate that, if all other parameters of the two models are kept identical, the Melitz model generates higher trade shares and therefore higher welfare gains from trade.

Instead, our exercise estimates the parameters of the two models such that they generate *identical* trade shares—the ones observed in the data. Then, the reason why the welfare gains differ is because different θ 's are necessary in order to match the price moments in the data.

7. Extensions and Robustness

In this section, we present the results from two exercises. First, we modify our estimation approach to incorporate “macro” moment conditions of the sort articulated in Arkolakis et al. (2012). We find that these aggregate moments are not informative about the trade elasticity relative to our “micro” approach. Second, we evaluate the sensitivity of our estimates of trade elasticities to the presence of measurement error in the micro price data. We find that incorporating measurement error only strengthens our results by magnifying the quantitative differences across models.

7.1. Macro and Micro Estimation

How do our estimates of trade elasticities compare to estimates that one would obtain by following a “macro” based approach using aggregate measures of trade frictions? This is a particularly relevant question in light of the results of Arkolakis et al. (2012). As they suggest, a macro based approach that uses tariffs and assumes a model-independent orthogonality condition results in an estimate of θ that is model independent. This implies that the welfare gains from trade are the same across models.

To compare the two approaches, we modify our estimation to incorporate macro moment conditions that use aggregate measures of tariffs. The idea behind this exercise is to let the data speak to which moment conditions are more powerful. In other words, depending upon aspects of the data, the combined estimation procedure has the ability to deliver an estimate of the trade elasticity that is common across models or different.

We build on Caliendo and Parro’s (2014) approach, which represents the state-of-the-art technique to estimate the trade elasticity off of aggregate tariff data. For notational convenience, denote the log of triple differenced trade flows by \tilde{x}_{nih} for country triple n, i, h and the log of triple differenced tariffs by \tilde{t}_{nih} . Finally denote by N_T the number of country triple observations.

Caliendo and Parro's (2014) estimation is based on the orthogonality condition

$$\mathbb{E} \{ \tilde{t}_{nih} (\tilde{x}_{nih} - \theta \tilde{t}_{nih}) \} = 0. \quad (28)$$

That is the covariance of triple differenced tariffs is orthogonal to the residual, i.e. their estimation is just OLS. To incorporate this orthogonality condition in our context, we stack moment conditions in the following way. Define

$$h_{nih}(\theta, \tilde{x}_{nih}, \tilde{t}_{nih}) = \tilde{t}_{nih} (\tilde{x}_{nih} - \theta \tilde{t}_{nih}), \quad (29)$$

then averaging across triples yields

$$\bar{h}_T(\theta) = \left(\frac{1}{N_T} \right) \sum_n \sum_i \sum_h h_{nih}(\theta, \tilde{x}_{nih}, \tilde{t}_{nih}), \quad (30)$$

where N_T is the number of triples. Caliendo and Parro's (2014) estimate of θ is the value that minimizes the quadratic form of $h_T(\theta)$. Or mathematically:

$$\hat{\theta}_T = \arg \min_{\theta} \bar{h}_T(\theta)' \bar{h}_T(\theta). \quad (31)$$

where the notation $\hat{\theta}_T$ denotes that this is the estimate from the tariff based approach.

To combine this with our estimation method, we stack the h functions that contain our micro moments and the tariff-based moments. So define $\tilde{\mathbf{h}}(\theta, \mathcal{M})$ as

$$\tilde{\mathbf{h}}(\theta, \mathcal{M}) = \begin{Bmatrix} \bar{h}_T(\theta) \\ \bar{h}(\theta, \mathcal{M}) \end{Bmatrix} \quad (32)$$

where $\bar{h}(\theta, \mathcal{M})$ is defined in equation (13) from our exactly identified estimation approach. Then the estimate of θ which merges our approach with that of Caliendo and Parro's (2014) is the value that minimizes the quadratic form of $\tilde{\mathbf{h}}(\theta, \mathcal{M})$. Or mathematically:

$$\hat{\theta}(\mathcal{M}) = \arg \min_{\theta} \tilde{\mathbf{h}}(\theta, \mathcal{M})' \mathbf{W}^{-1} \tilde{\mathbf{h}}(\theta, \mathcal{M}). \quad (33)$$

where \mathbf{W} is a positive definite weighting matrix.

For the weighting matrix, we will present two sets of results: one using the identity matrix and another using the variance-covariance matrix of the moments. As before, we use the continuous-updating estimator of \mathbf{W} , which continually updates the weighting matrix within the minimization routine.

We compute standard errors using a parametric bootstrap technique. The procedure for sim-

Table 7: Macro and Micro Estimates of θ

Model	Macro	Micro	Macro + Micro	
			Identity Matrix	Optimal Matrix
Armington/Krugman	4.63 [5.21, 4.09]	5.24 [5.02, 5.45]	5.24 [5.01, 5.45]	5.23 [5.01, 5.44]
EK	4.63 [5.21, 4.09]	4.17 [4.00, 4.34]	4.17 [4.03, 4.39]	4.17 [4.04, 4.39]
Melitz	4.63 [5.21, 4.09]	3.69 [3.32, 4.21]	3.70 [3.32, 4.12]	3.70 [3.32, 4.12]
BEJK, $\rho=1.37$	4.63 [5.21, 4.09]	2.74 [2.62, 2.89]	2.74 [2.62, 2.87]	2.74 [2.63, 2.87]

Note: Macro presents estimates from (31). Micro presents estimates from (14). Macro + Micro presents estimates from (33). (Values within brackets report 90th and 10th confidence intervals).

ulating the variability in our micro moments is exactly as described in Section 4. To simulate variability in the macro moments, we assume that the deviation between triple differenced trade flows and triple differenced tariffs is normally distributed. Given our parametric assumption on the residuals, we are able to simulate variability in the macro moments. Together, this approach allows us to construct 90-10 percentile confidence intervals.

Data. To implement this approach we used the trade flow and tariff data from Parro (2013). In this data set, there are 24 countries (20 of them are in our data set) for the year 1990. Because the years differ between our data set and Parro’s (2013) data set, our underlying assumption is that the trade elasticity is invariant across time. Parro’s (2013) data set contains data on tariffs for capital goods and non-capital goods. Since we are interested in an aggregate tariff measure, we construct aggregate tariffs by taking a trade-weighted average of the two sectors. Simple averaging gives similar results.

Results. The first column of Table 7 presents the results using only the tariff based estimation approach. Using this data we find an estimate of θ of around 4.6. This result is very much consistent with the results that Parro (2013) reports. The second column reproduces the results from our exactly identified estimation. For the BEJK model, we set the value of ρ equal to 1.37, its estimated value from the previous section.

The third and fourth column of Table 7 report our estimate of θ when both the micro moment conditions and the macro moment conditions are used. The key result is that the third and fourth column are nearly identical to the second column. Either with an identity or optimal weighting matrix, incorporating macro based moment conditions has virtually zero effect on our estimates.




Figure 2: Trade Flows and Tariffs

The reason for this result is that the tariff based method provides very little information about the exact nature of the trade elasticity. To see this, Figure 2 plots the triple differenced tariff data on the x-axis and triple differenced trade flow data on the y-axis. Plotted on top of the data are the “best fit” lines implied by the different estimation methods and models.

The key thing to notice is that the data are exceptionally noisy and that the pure tariff based approach does not describe the data that well. This implies that large deviations (say from 4.6 to 2.7) from the least square best fit provide negligible losses in mean square error. This means that tariffs provide little valuable information with regards to θ . Thus when the macro approach is combined with our estimation method, the tariff data does not assist in identifying θ .

7.2. Measurement Error in Prices

The second extension we explore is the role of measurement error in our micro price data. Measurement error in our micro price data is a concern for several reasons. First, as we showed in Simonovska and Waugh (2014), abstracting from measurement error will bias our estimates of the trade elasticity downwards. A second concern, is that this bias may be model specific. That is the differences in θ across models could be smaller (or larger) depending on how measurement error interacts with the structure of the model.

To address these concerns we jointly estimate θ and the amount of measurement error in the data. To do so, we assume that the micro-level prices are measured with error of the following

Table 8: Estimates of θ with Measurement Error

Model	Overidentified, Identity		Overidentified, Optimal		
	Estimate of θ	Estimate of σ_ϵ	Estimate of θ	Estimate of σ_ϵ	J-stat P-value
Armington/Krugman	6.79 [6.70, 7.27]	0.186 [0.175, 0.194]	7.80 [7.22, 8.54]	0.220 [0.213, 0.2275]	< 0.01
EK	4.13 [3.96, 4.29]	0.000 [0, 0.085]	4.05 [3.96, 4.29]	0.000 [0, 0.074]	0.20
Melitz	4.79 [4.09, 5.83]	0.151 [0.104, 0.189]	5.07 [4.31, 6.18]	0.176 [0.133, 0.203]	0.33
BEJK, $\rho = 1.37$	2.75 [2.65, 3.05]	0.000 [0, 0.114]	2.70 [2.66, 2.99]	0.000 [0, 0.109]	< 0.01

Note: σ_ϵ is the standard deviation of measurement error. Values within brackets report 90th and 10th confidence intervals. J-stat P-value reports probability of observing a J-statistic as large as observed.

form

$$\log p_n(\ell) = \log p_n^*(\ell) + \epsilon, \quad \text{where } \epsilon \sim \text{N}(0, \sigma_\epsilon), \quad (34)$$

where $p_n(\ell)$ is the observed price of good ℓ in country n , $p_n^*(\ell)$ is the true price, and ϵ is the deviation between the true price and the observed price, i.e. measurement error. We assume that ϵ is distributed normally with mean zero and standard deviation σ_ϵ .

We then jointly estimate the two parameters, θ and σ_ϵ , to minimize the distance between the three model moments and the three data moments in the overidentified case as described in Section 4. We present the results with both the identity and optimal weighting matrix. Standard errors and the weighting matrix are computed as described in Section 4.

Results. Table 8 presents the results. The first two columns present the estimate of θ and measurement error with an identity matrix. The last three columns present the results with the optimal weighting matrix.

The main result from Table 8 is that even in the presence of measurement error, we find that the ordering of models is still preserved with new trade models delivering lower trade elasticities. In fact, the quantitative differences are larger, e.g. EK is now nearly 50 percent lower relative to Armington/Krugman. This difference is largely because we find a much higher value of θ (6.78 versus 5.23) in the Armington/Krugman model relative to our baseline results. What is occurring is that the Armington/Krugman model requires a substantial amount of measurement error with an offsetting reduction in the trade elasticity to best fit the observed amount of price variation. This does not imply, however, that Armington/Krugman plus measurement error is an accurate description of the data; we clearly reject the null hypothesis in this case that the data generating process is true.

There are several other notable features of these results. For both EK and BEJK, our point estimates for θ are effectively the same as our baseline results. While the point estimates for σ_ϵ indicate little presence of measurement error, the confidence interval is wide. Second, like our results in Table 4 and Table 6, we fail to reject the null hypothesis that the EK model or the Melitz model generated the data.

8. Conclusion

We argue that the welfare gains from trade in new trade models with various micro-level margins are larger relative to models without these margins. First, we make a theoretical argument: Assuming a fixed trade elasticity, different models have identical predictions for bilateral trade flows but different predictions for cross-country price variation. Thus, for given data on trade flows and prices, different models have different trade elasticities and thus different welfare gains from trade. Second, we estimate the trade elasticity and the welfare cost of autarky under different assumptions about the underlying model. Relative to an Armington trade model, an active extensive margin increases the welfare cost of autarky by 20 percent; variable markups further increase the cost by 30 percent. In addition, relative to a Krugman model, an active extensive margin as in Melitz increases the welfare cost of autarky by 30 percent. Third, we incorporate into our estimation moment conditions that use trade flow and tariff data, which imply a common trade elasticity; this modification has virtually no effect on our results.

References

- ADDA, J. AND R. W. COOPER (2003): *Dynamic economics: quantitative methods and applications*, MIT press.
- ANDERSON, J. AND E. VAN WINCOOP (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93, 170–192.
- ANDERSON, J. E. (1979): "A Theoretical Foundation for the Gravity Equation," *American Economic Review*, 69, 106–16.
- ARKOLAKIS, C., A. COSTINOT, AND A. RODRIGUEZ-CLARE (2012): "New Trade Models, Same Old Gains?" *American Economic Review*, 102, 94–130.
- ARKOLAKIS, C., S. DEMIDOVA, P. KLENOW, AND A. RODRÍGUEZ-CLARE (2008): "Endogenous Variety and the Gains from Trade," *The American Economic Review*, 98 (2), 444–450.
- ARNOLD, B. C., N. BALAKRISHNAN, AND H. N. NAGARAJA (1992): *A first course in order statistics*, vol. 54, Siam.
- BERNARD, A., J. EATON, J. B. JENSEN, AND S. KORTUM (2003): "Plants and Productivity in International Trade," *American Economic Review*, 93, 1268–1290.
- BRADFORD, S. (2003): "Paying the Price: Final Goods Protection in OECD Countries," *The Review of Economics and Statistics*, 85, 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.
- CALIENDO, L. AND F. PARRO (2014): "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studies* (forthcoming).
- CHANEY, T. (2008): "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, 98, 1707–1721.
- 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.

- DAVISON, A. C. AND D. HINKLEY (1997): *Bootstrap methods and their application*, vol. 1, Cambridge university press.
- DE BLAS, B. AND K. RUSS (2010): "Understanding Markups in the Open Economy under Bertrand Competition," Working Paper 16587, National Bureau of Economic Research.
- EATON, J. AND S. KORTUM (2002): "Technology, Geography, and Trade," *Econometrica*, 70, 1741–1779.
- EATON, J., S. KORTUM, AND F. KRAMARZ (2011): "An Anatomy of International Trade: Evidence From French Firms," *Econometrica*, 79, 1453–1498.
- GOURIÉROUX, C. AND A. MONFORT (1996): *Simulation-based econometric methods*, CORE lectures, Oxford University Press.
- HANSEN, L. P., J. HEATON, AND A. YARON (1996): "Finite-sample properties of some alternative GMM estimators," *Journal of Business & Economic Statistics*, 14, 262–280.
- KRUGMAN, P. (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *American Economic Review*, 70, 950–959.
- MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725.
- MELITZ, M. J. AND S. J. REDDING (2014): "New Trade Models, New Welfare Implications," Working Paper 18919, National Bureau of Economic Research.
- PARRO, F. (2013): "Capital-skill complementarity and the skill premium in a quantitative model of trade," *American Economic Journal: Macroeconomics*, 5, 72–117.
- SIMONOVSKA, I. AND M. E. WAUGH (2014): "The elasticity of trade: Estimates and evidence," *Journal of International Economics*, 92, 34 – 50.
- WAUGH, M. E. (2010): "International Trade and Income Differences," *American Economic Review*, 100, 2093–2124.

For Online Publication

A. Proofs

The proof to Proposition 1 follows.

Proof: a. and b. We derive the trade shares and price index for the version of the Armington model that we consider in the present paper.

Maximizing U_n subject to a standard budget constraint yields the following CES demand function

$$x_{ni}(j) = \left(\frac{p_{ni}(j)}{P_n} \right)^{-\rho} \frac{w_n L_n}{P_n}, \quad (35)$$

$$P_n = \left[\sum_{i=1}^N \int_0^{1/N} p_{ni}(j)^{1-\rho} dj \right]^{\frac{1}{1-\rho}}. \quad (36)$$

Perfect competition ensures that pricing is according to marginal cost of production and delivery,

$$p_{ni}(j) = \frac{\tau_{ni} w_i}{T_i^{1/\theta}}. \quad (37)$$

Substituting (37) into (36), integrating, and letting $\theta = \rho - 1$ yields $P_n \propto \Phi_n^{-1/\theta}$, with proportionality constant $1/N$.

Finally, using (35) and (37), and integrating, the share of expenditure that n spends on goods from i is given by

$$\frac{X_{ni}}{\sum_{k=1}^N X_{nk}} = \frac{\left(\tau_{ni} w_i T_i^{-1/\theta} \right)^{1-\rho}}{\sum_{k=1}^N \left(\tau_{nk} w_k T_k^{-1/\theta} \right)^{1-\rho}} \quad (38)$$

Letting $\theta = \rho - 1$ yields the expression in the text. ■

The proof to Proposition 4 follows.

Proof: We only need to show the results for the endogenous variety models.

a. and b. We derive the trade shares and price indices for the versions of the Krugman (1980) and Melitz (2003) models that we consider in the present paper.

For the Krugman (1980) model, maximizing W_n subject to a standard budget constraint yields

the following CES demand function

$$x_{ni}(j) = \left(\frac{p_{ni}(j)}{P_n} \right)^{-\rho} \frac{w_n L_n}{P_n}, \quad (39)$$

$$P_n = \left[\sum_{i=1}^N \int_{J_i} p_{ni}(j)^{1-\rho} dj \right]^{\frac{1}{1-\rho}}. \quad (40)$$

Moreover, maximizing firm profits and accounting for the demand in (39) yields

$$p_{ni}(j) = \frac{\bar{m} \tau_{ni} w_i}{T_i^{1/\theta}}. \quad (41)$$

Free entry implies that the expected profit of a firm is zero, after paying the fixed entry cost. Using the firm's optimal pricing rule and quantity demanded, which must equal to the quantity produced, yields

$$\sum_{n=1}^N \frac{(p_{ni}(j))^{1-\rho} w_n L_n}{P_n^{1-\rho} \rho} - w_i e_i = 0 \quad (42)$$

Let M_i denote the measure of entrants in country i , which corresponds to the measure of set J_i . In order to compute the measure of entrants, combine the free entry condition from (42) with the labor market clearing condition for country i ,

$$M_i \sum_{n=1}^N \frac{(p_{ni}(j))^{1-\rho}}{P_n^{1-\rho}} w_n L_n + e_i = L_i, \quad (43)$$

to obtain $M_i = L_i / (\rho e_i)$. Using the definition $\theta = \rho - 1$ and assuming that $L_i \propto e_i$ yields $M_i \propto 1/(1 + \theta)$. Relying on this result, use the equilibrium measure of entrants M_i and the optimal pricing rule (41) into (40) to obtain that $P_n \propto \Phi_n^{-1/\theta}$.

Finally, the share of expenditure that n spends on goods from i is given by

$$\frac{X_{ni}}{\sum_{k=1}^N X_{nk}} = \frac{M_i \frac{(p_{ni}(j))^{1-\rho}}{P_n^{1-\rho}} \frac{w_n L_n}{\rho}}{\sum_{k=1}^N M_k \frac{(p_{nk}(j))^{1-\rho}}{P_n^{1-\rho}} \frac{w_n L_n}{\rho}} \quad (44)$$

Using the definition $\theta = \rho - 1$, the equilibrium measure of entrants M_i under the proportionality assumption, the optimal pricing rule (41) and the exact price index (40) in (44) yields the expression in the text.

For the Melitz (2003) model, following the same steps as for the Krugman (1980) model obtains the demand function in (39) and the price index in (40). Let the marginal cost of a firm with

efficiency draw $z_i(j)$ of producing good j in i and delivering it to n be

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

We will express all equilibrium objects in terms of cost draws in the remainder of the Appendix. As in the Krugman (1980) model, the optimal pricing rule is given by

$$p_{ni}(j) = \bar{m} c_{ni}(j) \tag{45}$$

Using the firm's optimal pricing rule and quantity demanded, which must equal to the quantity produced, yields the firm's profit function

$$\pi_{ni}(c_{ni}(j)) = w_n L_n \left(\frac{\rho}{\rho - 1} \right)^{-\rho} \frac{1}{\rho - 1} \left(\frac{c_{ni}(j)}{P_n} \right)^{1-\rho} - w_n a_n$$

The cost cutoff for a firm from i to sell good j to n , $\bar{c}_{ni}(j)$, satisfies $\pi_{ni}(\bar{c}_{ni}(j)) = 0$ and is given by

$$\bar{c}_{ni}(j) = \left[\frac{a_n}{L_n} \frac{P_n^{1-\rho}}{\left(\frac{\rho}{\rho-1} \right)^{-\rho} \frac{1}{\rho-1}} \right]^{\frac{1}{1-\rho}}$$

Since $\bar{c}_{ni}(j) = \bar{c}_{nn}(j) \forall i$, denote cost cutoffs by the destination to which they apply, $\bar{c}_n(j)$.

To determine the measure of entrants, repeat the procedure applied to the Krugman (1980) model to obtain $M_i = L_i/e_i(\rho - 1)/(\rho\theta)$. Only those firms with cost draws below \bar{c}_i serve the domestic market and the remainder immediately exit. Using the Pareto distribution, the measure of domestic firms that sell j in country i is $M_{ii}(j) = M_i T_i w_i^{-\theta} (\bar{c}_i(j))^\theta$. Similarly, the measure of firms from i that sell j to n is $M_{ni}(j) = M_i T_i (w_i \tau_{ni})^{-\theta} (\bar{c}_n(j))^\theta$.

Furthermore, assuming that $L_i \propto e_i$ yields $M_i \propto (\rho - 1)/(\rho\theta)$. Relying on this result, use the equilibrium measure of firms that serve destination n from any source i , M_{ni} , the optimal pricing rule (45), and the Pareto distribution into (40) to obtain that $P_n \propto \Phi_n^{-1/\theta}$.

Finally, to derive the trade share, use the optimal pricing rule (45), quantity demanded, which must equal to the quantity produced, the equilibrium measure of firms that serve each destination, and the Pareto distribution.

c. We refer the reader to Arkolakis et al. (2012) for derivations concerning welfare in all the models. ■

The proof of the ranking between the price gap distributions of the EK and BEJK models at the boundary points follows.

Proof: We need to show that $P_{\mathcal{EK}}(\tilde{P}_h = \log \tau) > P_{\mathcal{BEJK}}(\tilde{P}_h = \log \tau)$ and $P_{\mathcal{EK}}(\tilde{P}_h = -\log \tau) > P_{\mathcal{BEJK}}(\tilde{P}_h = -\log \tau)$. Let $\tilde{P}_{h,\mathcal{EK}}$ and $\tilde{P}_{h,\mathcal{BEJK}}$ denote the logged price gaps in the EK and BEJK models respectively. Following the notation in the text, c_{khf} , $k = 1, 2$ denotes the k th best supplier from F to H. We show that $\tilde{P}_{h,\mathcal{EK}} = \log \tau$ holds true in higher number of instances than $\tilde{P}_{h,\mathcal{BEJK}} = \log \tau$, and similarly for the lower bound, $-\log \tau$.

Broadly, there are two cases to consider: (1) traded goods (either exported by the home or foreign country) and (2) non-traded goods. For each case, there are sub cases that correspond to different price gaps in the BEJK model as pricing is based on the next best competitor. In the EK model, price gaps always reflect cost gaps due to perfect competition.

1. Traded goods

(a) Foreign (F) exports to Home (H). Then F is the lowest cost supplier in both F and H. There are three sub cases:

i. F is the second lowest cost supplier in both F and H. Then,

$$\begin{aligned}\tilde{P}_{h,\mathcal{EK}} &= \log c_{1hf} - \log c_{1ff} = \log \tau \\ \tilde{P}_{h,\mathcal{BEJK}} &= \log c_{2hf} - \log c_{2ff} = \log \tau \\ \Rightarrow \tilde{P}_{h,\mathcal{EK}} &= \tilde{P}_{h,\mathcal{BEJK}} = \log \tau\end{aligned}$$

ii. F is the second lowest cost supplier in F; H is the second lowest cost supplier in H. Then,

$$\begin{aligned}\tilde{P}_{h,\mathcal{EK}} &= \log c_{1hf} - \log c_{1ff} = \log \tau \\ \tilde{P}_{h,\mathcal{BEJK}} &= \log c_{1hh} - \log c_{2ff} < \log \tau \\ \Rightarrow \tilde{P}_{h,\mathcal{EK}} &= \log \tau > \tilde{P}_{h,\mathcal{BEJK}}\end{aligned}$$

iii. H is the second lowest cost supplier in both F and H. This violates no-arbitrage and cannot occur.

(b) Home (H) exports to Foreign (F). Then H is the lowest cost supplier in both F and H. There are three sub cases:

i. H is the second lowest cost supplier in both F and H. Then,

$$\begin{aligned}\tilde{P}_{h,\mathcal{EK}} &= \log c_{1hh} - \log c_{1fh} = -\log \tau \\ \tilde{P}_{h,\mathcal{BEJK}} &= \log c_{2hh} - \log c_{2fh} = -\log \tau \\ \Rightarrow \tilde{P}_{h,\mathcal{EK}} &= \tilde{P}_{h,\mathcal{BEJK}} = -\log \tau\end{aligned}$$

- ii. F is the second lowest cost supplier in F; H is the second lowest cost supplier in H. Then,

$$\begin{aligned}\tilde{P}_{h,\varepsilon\mathcal{K}} &= \log c_{1hh} - \log c_{1fh} = -\log \tau \\ \tilde{P}_{h,\mathcal{BEJK}} &= \log c_{2hh} - \log c_{1ff} > -\log \tau \\ \Rightarrow \tilde{P}_{h,\varepsilon\mathcal{K}} &= -\log \tau < \tilde{P}_{h,\mathcal{BEJK}}\end{aligned}$$

- iii. F is the second lowest cost supplier in both F and H. This violates no-arbitrage and cannot occur.

All four sub cases occur with positive probability. Given the strict inequality in sub cases a.ii and b.ii, it must be that $P_{\varepsilon\mathcal{K}}(\tilde{P}_h = \log \tau) > P_{\mathcal{BEJK}}(\tilde{P}_h = \log \tau)$ and $P_{\varepsilon\mathcal{K}}(\tilde{P}_h = -\log \tau) > P_{\mathcal{BEJK}}(\tilde{P}_h = -\log \tau)$. There remain four sub cases that correspond to non-traded goods. In each sub case, $P_{h,\varepsilon\mathcal{K}}, P_{h,\mathcal{BEJK}} \in (-\log \tau, \log \tau)$, but the two random variables cannot be ordered theoretically. Hence, while the mass that lies in the interior is larger in the BEJK versus the EK model, it cannot be shown analytically how this mass is distributed. It is for this reason that it cannot be shown analytically that $G_{\varepsilon\mathcal{K}}$ stochastically dominates $G_{\mathcal{BEJK}}$. We present the remaining four sub cases for completeness.

2. Non-traded goods. F is the lowest cost supplier in F and H is the lowest cost supplier in H. There are four sub cases:

- (a) F is the second lowest cost supplier in both F and H. Then,

$$\begin{aligned}\tilde{P}_{h,\varepsilon\mathcal{K}} &= \log c_{1hh} - \log c_{1ff} \in (-\log \tau, \log \tau) \\ \tilde{P}_{h,\mathcal{BEJK}} &= \log c_{1hf} - \log c_{2ff} = \log \tau + \log c_{1ff} - \log c_{2ff} \in (-\log \tau, \log \tau)\end{aligned}$$

- (b) H is the second lowest cost supplier in both F and H. Then,

$$\begin{aligned}\tilde{P}_{h,\varepsilon\mathcal{K}} &= \log c_{1hh} - \log c_{1ff} \in (-\log \tau, \log \tau) \\ \tilde{P}_{h,\mathcal{BEJK}} &= \log c_{2hh} - \log c_{1fh} = \log c_{1hh} - \log \tau - \log c_{1hh} \in (-\log \tau, \log \tau)\end{aligned}$$

- (c) F is the second lowest cost supplier in F; H is the second lowest cost supplier in H. Then,

$$\begin{aligned}\tilde{P}_{h,\varepsilon\mathcal{K}} &= \log c_{1hh} - \log c_{1ff} \in (-\log \tau, \log \tau) \\ \tilde{P}_{h,\mathcal{BEJK}} &= \log c_{2hh} - \log c_{2ff} \in (-\log \tau, \log \tau)\end{aligned}$$

(d) H is the second lowest cost supplier in F; F is the second lowest cost supplier in H.
Then,

$$\begin{aligned}\tilde{P}_{h,\varepsilon\mathcal{K}} &= \log c_{1hh} - \log c_{1ff} \in (-\log \tau, \log \tau) \\ \tilde{P}_{h,\mathcal{B}\varepsilon\mathcal{J}\mathcal{K}} &= \log c_{1hf} - \log c_{1fh} = \log c_{1ff} - \log c_{1hh} \in (-\log \tau, \log \tau)\end{aligned}$$

■