Quantifying the Sensitivity of Quantitative Trade Models*

Habib Chenguiti Ansari
Birkbeck

Dave Donaldson
MIT

Edward Wiles
Harvard

Abstract

Modern quantitative trade models emphasize cross-country heterogeneity in productivity across goods as the source of trade and its benefits. In workhorse versions, this heterogeneity is assumed to follow specific functional forms, such as i.i.d. Frechet or Pareto. We provide a systematic exploration of how sensitive the gains from trade estimated from these models are to relaxing this assumption—first, with a new formula for sharp bounds on gains from trade across all possible distributions, and then with a numerical characterization for the case where distributions must be close to the workhorse form. Importantly, for each candidate distribution we ensure that its implied trade model is consistent with standard estimation methods such as the ability to match all bilateral trade data points and an exogeneity assumption. Our estimated bounds are wide, even for small departures from the canonical model: the maximal gain is about four times larger and the minimal gain is zero. Imposing high-level restrictions on the distributions considered—for example, maintaining Frechet or Pareto marginals while relaxing independence, or vice versa—does little to shrink these bounds.

*This version: September 2026. Author contacts: mhansari@mit.edu, ddonald@mit.edu, edwardwiles@fas.harvard.edu. We are grateful to Isaiah Andrews, Costas Arkolakis, Kirill Borusyak, Tim Christensen, Arnaud Costinot, Federico Huneeus, Samuel Kortum, Thierry Mayer, Eduardo Morales, Stephen Redding, Andres Rodriguez-Clare, Felix Tintelnot, Daniel Trefler, Jonathan Vogel, David Weinstein, and Daniel Xu for helpful comments and to Noah Siderhurst for outstanding research assistance.

1 Introduction

Major advances in the field of international economics have come recently from the development of quantitative trade models: general equilibrium models that are designed to match empirical patterns about factual scenarios and make predictions about economic responses to counterfactual ones.[1] A core ingredient is the extent to which productive capabilities in the model are heterogeneous across goods and countries. Indeed, the exploitation of this heterogeneity—where countries import the goods that cannot be made cheaply at home—is what gives rise to trade and its benefits in the first place. However, quantifying this heterogeneity poses a foundational challenge: how can one hope to measure a country's ability to produce the goods it doesn't produce?

To make progress, researchers have posited specific functional forms to describe production heterogeneity around the world. Two choices that result in so-called gravity models—the independent and identically distributed (i.i.d.) Frechet distribution used by Eaton and Kortum (2002) for competitive models and the i.i.d. Pareto distribution used by Chaney (2008) for monopolistically competitive models—have become enormously influential for workhorse quantitative modeling. But how would economists' predictions about counterfactuals, such as the size of the gains from trade, change if alternative forms of production heterogeneity were considered instead? The goal of this paper is to conduct a comprehensive quantification of the sensitivity of estimated gains from trade in quantitative trade models as we allow for departures from these workhorse assumptions.

Our approach is disciplined by three important constraints. First, in order to focus on the role of production heterogeneity, we hold constant other model components, such as preferences and factor market structure. Second, in line with standard practice for quantitative modeling, we ensure that all bilateral trade flow data points are matched exactly as we vary countries' productivity distributions—ensuring that empirical facts such as the extent of trade openness, the fit of the gravity equation, and the correlation between trade flows and tariffs are always replicated too. Finally, we impose a standard identification assumption that observed trade costs are

[1] Two recent chapters in the *Handbook of International Economics*, Costinot and Rodríguez-Clare (2014) and Caliendo and Parro (2022), provide surveys of this practice. Analogous progress on quantitative spatial models of intra-national trade is also surveyed in two recent chapters of the *Handbook of Regional and Urban Economics*, Allen and Arkolakis (2025) and Redding (2025).

orthogonal to unobserved ones.

Then, subject to these constraints, we derive new formulae for the maximal and minimal bounds on countries' gains from trade that are possible across all underlying productivity distributions. We also compute similar bounds for the case where admissible distributions must be "similar" to (i.e. have low statistical divergence from) the gravity model versions. Our main finding is that the estimated gains from trade in quantitative models are highly sensitive to alternative distributions of production heterogeneity.

Beginning with the distribution-agnostic case, our new analytical result states that (provided a mild data condition holds, as it always does in our sample) any country's gains from trade are bounded in a very simple manner. The maximal gain is given by a simple formula involving one data input (the country's openness) and one preference parameter (the elasticity of substitution across goods). For France, a representative case, this upper bound is 3.6 times larger than the typical gravity model benchmark. There are therefore data-consistent versions of underlying productivity distributions for which gains from trade can be very large. On the other hand, our lower bound is zero, so it is also possible that there are no gains from trade at all.

While these wide bounds allow for distributions that may be arbitrarily different from those of the gravity model workhorses of Eaton and Kortum (2002) and Chaney (2008) our next result suggests that similarly wide bounds are possible even for distributions that are restricted to be "close" to these benchmarks. Our approach here draws on recent advances in distribution-robust econometrics, beginning with Hansen and Sargent (2001). This allows us to quantify bounds on the gains from trade when distributions must lie within a fixed level of divergence (a measure of the similarity of any two probability distributions) from the gravity case. Focusing on divergence levels that are analogous to other robustness checks that are routinely performed on workhorse models, we find that France's gains from trade can range between zero and 3.4 times the gravity model value. So they can essentially span most of the range (i.e. zero to 3.6 times) that exists when all distributions are allowed for.

It is of course infeasible to evaluate every one of the infinite number of distributions within any given neighborhood that is close to the gravity model. And doing so is harder still when every candidate distribution must be tested for whether it

can generate a worldwide equilibrium whose trade flows match all data points and that it is consistent with standard gravity model estimation requirements. To circumvent this challenge we deploy the approach in Christensen and Connault (2023), who consider an optimization problem over all distributions of a random variable within a given divergence neighborhood, where the objective and constraints are expectations of (known functions of) the random variable. Following this logic, we write our problem—maximize or minimize gains from trade subject to the constraints of model equilibrium, matching all trade data, and an orthogonality assumption—as one with an inner loop that searches over distributions (given all unknown model parameters and market-clearing equilibrium variables in the factual and counterfactual equilibria) and an outer loop that searches over the parameters and equilibrium variables. The inner loop is infinite-dimensional but benefits from a dual formulation proposed by Christensen and Connault (2023) that is both finite-dimensional and convex. The outer loop is high-dimensional and typically non-convex but benefits from a dimensionality-reduction procedure that we develop. Putting these results together, the program becomes feasible.

Given that we find so little robustness of quantitative models, even for small amounts of departure from the gravity model, it is natural to restrict the admissible sets of productivity distributions further. Our final set of results does this by imposing high-level requirements that are guided by three core features of the benchmark i.i.d. distributions used in gravity models: (i) countries take draws independently from one another; (ii) countries all have a common marginal distribution; and (iii) the particular marginal distribution that is common is the Frechet form in Eaton and Kortum (2002) and the Pareto form in Chaney (2008). We show how to extend our baseline analysis—varying the model's production heterogeneity while matching all trade flows and standard orthogonality conditions—while also imposing constraints (i) through (iii), one at a time or in combination, in order to learn which particular features of the benchmark distribution do the most to constrain the model's predictions.

Across a wide range of restrictions imposed, the broad message that emerges is that the bounds on gains from trade remain wide despite requiring that productivity distributions satisfy reasonable high-level restrictions. We begin by establishing that for the case of monopolistically competitive models these restrictions have no direct bite at all because competition is never head-to-head on the same good. Turning

to the competitive case, where such head-to-head competition is central, we first extend our previous distribution-agnostic upper bound formula to cases where the productivity distribution must satisfy requirements such as independence, common (and Frechet) marginals, or even that countries draw i.i.d from some common marginal distribution. Each case requires a different data condition to hold, but these conditions are verified in our sample.

We then move on to evaluate neighborhoods that are close to the gravity model. In this setting, the bounds on gains from trade at fixed neighborhood size do shrink in the face of high-level restrictions, but the effects are modest. For example, when we impose conditions that approximate independence our upper bound shrinks from 3.4 times the baseline to 3.0 times the baseline (i.e. by 11%). And this is similar for the case where we impose zero correlation (18% shrinkage), common but flexible marginals (5%), common and Eaton and Kortum (2002) marginals (5%), and the i.i.d. condition of both independence and common but flexible marginals (16%). Perhaps surprisingly, this set of results implies that one can depart from Ricardian models with i.i.d. Frechet marginals in either of two obvious directions—dropping independence and keeping the identical Frechet marginal, or keeping independence and dropping Frechet marginals—and achieve similarly modest reductions in sensitivity in either case.

Related literature. Our analysis builds on previous explorations of how quantitative trade models behave under varying assumptions about production heterogeneity.

In particular, our approach is heavily influenced by that of Arkolakis et al. (2012), who proposed that the sensitivity of conclusions obtained across models should be conducted while conditioning on two requirements: (i) that the models are calibrated to match the same underlying trade data; and (ii) that the models are estimated from a gravity equation regression that assumes the same underlying orthogonality condition between observed and unobserved trade costs.[2] Indeed, those authors show that, once these two conditions hold, counterfactuals are strikingly robust across the four types of production heterogeneity they study—the Frechet distribution of Eaton and Kortum (2002) and the degenerate distribution of

[2]In this regard, our analysis also builds on the vast literature on gravity models of trade synthesized by Costinot and Rodríguez-Clare (2014) and Head and Mayer (2014).

Armington (1969) in the competitive case, and the Pareto distribution of Chaney (2008) and the degenerate distribution of Krugman (1980) in the monopolistically competitive case. We maintain these same two requirements but then broaden the set of productivity distributions considered to include all such distributions in order to fully assess the robustness uncovered by Arkolakis et al. (2012).

Our assessment is therefore complementary to other work that has evaluated productivity distributions that depart from the four members of the Arkolakis et al. (2012) class. On the competitive side, for example, Lind and Ramondo (2023) emphasize productivity distributions in which countries' draws exhibit rich forms of dependence. We embrace this direction by considering arbitrary distributions, as well as a restricted class that allows for any form of dependency but where countries' marginal distributions are the same as in the workhorse model. Adao et al. (2017) demonstrate how the implications of underlying productivity heterogeneity can be summarized by "reduced" import demand systems and that those demand systems can be identified from a cross-section of bilateral trade observations as long as the underlying system is assumed to be identical across countries.[3] Since one microfoundation for this case occurs when all countries draw productivities independently from the same marginal distribution, the bounds we compute in the special case where we impose this i.i.d. restriction explores what is possible along the direction proposed by Adao et al. (2017).

Turning to monopolistic competition, a wide range of studies have introduced models where countries' productivities continue to be drawn independently, as in the gravity model, but from marginal distributions that depart from the case of common Pareto. For example, Helpman et al. (2004) examine country-specific Pareto distributions; Helpman et al. (2008), Melitz and Redding (2015), and Feenstra (2018) consider the truncated Pareto version; Head et al. (2014), Bas et al. (2017), and Fernandes et al. (2023) analyze log-normal distributions; Nigai (2017) proposes a mixture of log-normal and Pareto; and Erhardt et al. (2021) and Adao et al. (2026) allow for nonparametric distributions estimated separately by country.[4] The

[3] An advantage of the reduced demand system approach in Adao et al. (2017) is that it summarizes the joint role played by technology, relative factor endowments, and tastes for generating trade. Our analysis focuses on a single factor setting and holds preferences constant across all models considered so it only relates to the technology side of their analysis.

[4] Several of these studies also follow Eaton et al. (2011) in allowing for additional forms of production-side heterogeneity beyond that in Melitz (2003), such as heterogeneity in fixed costs of exporting in Fernandes et al. (2023) and Adao et al. (2026). Our analysis sticks to the Melitz (2003)

primary contribution of our analysis is to generalize this exploration to encompass all possible distributions. We also highlight the (unavoidable) extrapolation that occurs when mapping from firm-level micro-data—an option in the monopolistic competition model case, since each product corresponds to a unique firm—to the welfare cost of autarky, since that cost depends only on the shape of a country's productivity distribution over a range of support for which no firm is active in the trading equilibrium. Our analysis therefore bounds all possible versions of this extrapolation.

2 Quantitative trade model environments

We begin by describing our core economic environment: a multi-country, single-factor trade model with many goods and arbitrary distributions of productivities (across goods and countries) around the world. We then describe how one can conduct counterfactual exercises (such as computing the gains from trade for any country) in such a model. Finally, we describe the version of this model that obtains when the global productivity distribution takes a particular form, as in workhorse versions of the model. Our presentation of these elements begins with a competitive version of this model (i.e. a Ricardian model of international trade) but then turns to the monopolistically competitive version (i.e. a Melitz (2003) model of international trade).

2.1 A competitive model of trade

Setup. There are D countries, potentially trading any of the goods along a continuum indexed by $\omega \in [0,1]$. Each country d produces goods using a factor of production that is owned by a representative household there and is immobile across countries.

Demand. The representative household in each country d has CES preferences across the continuum of goods:

---

case.

$$U_d = \left(\int_0^1 q_d(\omega)^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}},$$

with σ capturing the elasticity of substitution and $q_d(\omega)$ denoting the quantity of good ω consumed in country d. As is standard, we work with the case in which goods are substitutes (i.e. $\sigma > 1$).[5]

We let $p_{od}(\omega)$ denote the price of good ω in country d, *if* it were to be shipped from country o to country d. The determination of these prices involves the supply side of the economy, introduced below. But regardless of such determination, utility-maximization implies that the expenditure in d on good ω from country o will be

$$X_{od}(\omega) = E_d \gamma_d^{-1} p_{od}(\omega)^{1-\sigma}, \tag{1}$$

where E_d is total expenditure in d and where

$$\gamma_d \equiv \sum_o \int_0^1 p_{od}(\omega)^{1-\sigma} d\omega, \tag{2}$$

is a transformation of the standard CES price index. In particular, the real income of the representative household in d is given by $W_d = E_d \gamma_d^{\frac{1}{\sigma-1}}$.

Supply. The core of the Ricardian model lies on the supply side. We assume that producers in country o are endowed with a technology with productivity $z_o(\omega)$ for producing any good ω. The marginal cost of a producer of ω in o when selling in country d is given by

$$MC_{od}(\omega) \equiv \frac{w_o(1+t_{od})}{A_{od} z_o(\omega)}, \tag{3}$$

where w_o is the cost of a unit of the factor of production in o, t_{od} denotes the ad-valorem tariff charged in d on imports from o, and A_{od} is a bilateral productivity shifter.[6] The productivity level $z_o(\omega)$ augments production by shifting downwards the marginal cost of producing ω in o.

[5] This restriction is needed to ensure that the gains from trade, discussed below, always remain finite as we consider alternative technologies in which d may be unable to produce some goods. It is also natural in the monopolistic competition version of the model that we introduce in 2.5.

[6] The shifter A_{od} incorporates (inversely) an "iceberg" transport cost that acts like a tariff but generates no tax revenue.

Because the good ω is homogeneous, perfect competition implies that prices in location d are determined by the marginal cost of the lowest-cost supplier to that location. In particular, the price of the version of good ω that hails from o and is sold in d must satisfy

$$p_{od}(\omega) = \begin{cases} MC_{od}(\omega) & \text{if } o = \arg\min_{o' \in D} MC_{o'd}(\omega), \\ \infty & \text{otherwise.} \end{cases} \tag{4}$$

Naturally, consumers in country d only buy from the minimum-price supplier of any good; that is, $p_{od}(\omega) = \infty$ implies $X_{od}(\omega) = 0$ in equation (1).

Equilibrium. We let L_o denote the (exogenous) supply of the factor of production in o, meaning that total factor income in o is given by $w_o L_o$. For simplicity, we omit the contributions of tariff revenue and trade imbalances and hence have $E_o = w_o L_o$ for all o.[7] A competitive global equilibrium in this model is then one in which the factor market clears,

$$w_o L_o = \sum_d \int_0^1 X_{od}(\omega) d\omega \tag{5}$$

in every country o.

2.2 Stochastic representation

The key ingredient of this competitive model is the joint distribution of countries' productivities, $z_o(\omega)$ for all $o \in D$. For example, if productivities were to be distributed in such a way that they are perfectly correlated across countries then no country would have a comparative advantage in any good and there would be no trade (or gains from trade). On the other hand, if countries' distributions were relatively uncorrelated, or even negatively correlated, then a rationale for international specialization and hence trade emerges in the competitive equilibrium.

To emphasize the role of such distributions we follow the approach in Eaton and Kortum (2002) and treat the productivity $z_o(\omega)$ as a random variable that is

[7] This assumption simplifies the formulas presented below, which makes our core focus—the role of productivity heterogeneity—clearer. In practice, for France, our focal country of interest, both tariff revenues and the trade deficit are negligible fractions (0.07% and 0.67%, respectively) of total expenditure.

drawn from a multi-country stochastic process. However, in a departure from such work we allow for an *arbitrary* joint cumulative distribution function F. That is, F denotes a function $F(\{z\})$ that equals the probability that the D-dimensional vector of random variables $\{z_o(\omega)\}$ lies below the D-dimensional vector $\{z\}$. Because F is the joint distribution across countries this notation allows for two important possibilities: that countries have different marginal distributions (e.g. that country o tends to get more dispersed draws than country o' does) and that countries' draws display forms of dependence (e.g. that the draws of country o are negatively correlated with those of country o').

Using this stochastic representation we let the random variable $p_{od}(\omega)$ be defined as a function of $z_o(\omega)$ using equations (3) and (4). Similarly, the random variable $X_{od}(\omega)$ is defined, via equation (1), as a function of $p_{od}(\omega)$. Finally, we follow the law-of-large-numbers convention of using the expectation operator $\mathbb{E}_F[\cdot]$ over the random variable $\{z_o(\omega)\}$ instead of integrals across the continuum of goods. Doing so, we can re-write the endogenous variable γ_d from equation (2) as

$$\gamma_d = \sum_o \mathbb{E}_F\left[p_{od}(\omega)^{1-\sigma}\right], \tag{6}$$

and the market clearing condition (5) as

$$w_o L_o = \sum_d \mathbb{E}_F\left[X_{od}(\omega)\right]. \tag{7}$$

For future reference, we denote the total value of bilateral trade as $X_{od} \equiv \mathbb{E}_F\left[X_{od}(\omega)\right]$ and the total bilateral import share in expenditure as $\lambda_{od} \equiv X_{od}/w_d L_d$.

2.3 Counterfactuals and gains from trade

Our goal is to study how economic conditions in any given country d respond to counterfactual changes in the exogenous primitives of this model. In particular, we consider a change from the initial set of tariffs $\{t\}$ to the counterfactual set $\{t'\}$, and do so while holding constant all other exogenous model features—endowments $\{L\}$, bilateral productivity shifters $\{A\}$, preferences σ, and the multi-country productivity distribution F.

The counterfactual equilibrium is one in which all endogenous variables have

changed correspondingly. This means that each producer's counterfactual marginal cost is $MC'_{od}(\omega) = w'_o(1+t'_{od})A_{od}^{-1}[z_o(\omega)]^{-1}$ and the counterfactual price $p'_{od}(\omega)$ follows from (4) but evaluated at $MC'_{od}(\omega)$. Similarly, counterfactual expenditures are

$$X'_{od}(\omega) = w'_d L_d (\gamma'_d)^{-1} p'_{od}(\omega)^{1-\sigma}, \tag{8}$$

$$\text{with } \gamma'_d = \sum_o \mathbb{E}_F \left[p'_{od}(\omega)^{1-\sigma} \right], \tag{9}$$

and equilibrium factor prices clear markets in the counterfactual equilibrium

$$w'_o L_o = \sum_d \mathbb{E}_F \left[X'_{od}(\omega) \right]. \tag{10}$$

Gains from trade. A key counterfactual exercise of interest concerns the gains from trade for country d.[8] This is calculated by changing tariffs from their factual values of $\{t\}$ to the counterfactual values $\{t'\}$ that would place d in autarky (i.e. $t_{od} = \infty$ for all $o \neq d$) and then solving for the proportional change in real income. Following literature conventions, we define country d's gains from trade GT_d as the change in its real income (factual minus autarky counterfactual) divided by its factual value, or

$$GT_d \equiv 1 - \left(\frac{w'_d L_d}{w_d L_d} \right) \left(\frac{\gamma_d}{\gamma'_d} \right)^{\frac{1}{1-\sigma}} = 1 - \left(\frac{\gamma_d}{\gamma'_d} \right)^{\frac{1}{1-\sigma}}, \tag{11}$$

where the second equality follows from the choice of d's factor price as the numeraire.

Following Costinot and Rodríguez-Clare (2018), a useful device for evaluating gains from trade is via the elasticity of demand for imports. This proceeds as follows. First, let $X_{od}(\kappa)$ denote how country d's purchases from o would respond if the price (relative to the numeraire) of all of d's imports were to rise proportionally by the factor $\kappa \geqslant 1$. Then the import demand function, defined as $\lambda_{Md}(\kappa) \equiv \frac{1}{L_d} \sum_{o \neq d} X_{od}(\kappa)$, measures how the share of d's expenditure on imports would

[8] Though our focus is on gains from trade, additional sets of counterfactual calculations—such as those that examine finite tariff changes, endowment changes, and non-welfare outcomes such as trade flows and relative wages—are also feasible using the methods presented below.

change as a function of κ.[9] Using this notation, $GT_d = 1 - \exp(-\mathcal{A}_d)$, where $\mathcal{A}_d$ is the area under the $\lambda_{Md}(\kappa)$ function between $\kappa = 1$ (i.e. the baseline) and $\kappa \to \infty$ (i.e. autarky, where all imports are prohibitively costly):

$$\mathcal{A}_d \equiv \int_1^{\infty} \lambda_{Md}(\kappa) d \log \kappa.$$

Finally, let $\epsilon_d(\kappa) \equiv -d \log(\sum_{o \neq d} X_{od}(\kappa)/X_{dd}(\kappa))/d \log \kappa$ denote the elasticity of imports, relative to domestic purchases, with respect to the price of imports. Using this elasticity, the area under the demand curve can be expressed as

$$\mathcal{A}_d = \int_1^{\infty} \left[1 + \frac{\lambda_{dd}}{1-\lambda_{dd}} \exp\left(\int_1^{\kappa} \epsilon_d(k) d \log k\right)\right]^{-1} d \log \kappa, \tag{12}$$

where $\lambda_{dd} \equiv X_{dd}/w_d L_d$ is the share of d's expenditure at baseline that it sources from itself.

Equation (12) highlights how the area $\mathcal{A}_d$ under country d's import demand curve, which governs its gains from trade, depends on two components. The first is the country's baseline import share λ_{dd}, which is directly observable in standard data sources and hence revealed without knowledge of the productivity distribution F. The second is the schedule of elasticity values $\epsilon_d(\cdot)$ with which an increase in import prices affects the country's relative expenditure on imports and domestic goods at each point on the path to autarky. As discussed further below, this schedule depends intimately on the productivity distribution F.

2.4 The Eaton and Kortum (2002) competitive model

Our presentation of the competitive Ricardian model so far has held for any given distribution of global productivities F.[10] However, one particularly tractable version of F, proposed in the seminal model of Eaton and Kortum (2002), is that where

[9] Formally, Adao et al. (2017) show how competitive models can be represented by a reduced factor demand system in which the total expenditure by d on o's goods takes the form $X_{od}(w_d, \{W_{od}\}_{o \neq d})$, where $W_{od} \equiv w_o(1 + t_{od})$ is the tariff-inclusive price of o's foreign factor services in d. The function $X_{od}(\kappa)$ then corresponds to $X_{od}(w_d, \{W_{od}\}_{o \neq d})$ at $w_d = 1$ and where all elements of the vector $\{W_{od}\}_{o \neq d}$ are scaled up proportionally from their baseline values by the amount $\kappa \geqslant 1$.

[10] The only (implicit) assumption has been that all relevant expectations are finite. To facilitate interpretation, we will restrict attention to distributions for which $\mathbb{E}_F[z_o(\omega)^{\sigma-1}] < \infty$ for all o, which is necessary and sufficient for the price index (and all other relevant expectations) to be finite.

countries' productivity levels are drawn independently from a Frechet distribution with scale parameter equal to one and shape parameter equal to $\theta^* > 0$, i.e. $z_o(\omega) \sim Frechet(1, \theta^*)$ for all o.[11] We refer to this distribution as F^* in what follows (for the subset of our analysis referring to the competitive model) and treat it as a key reference point with which to compare the competitive model based on other distributions F.

We note that F^* imposes three separate types of restrictions relative to the unrestricted F: (i) countries' distributions are independent from one another, so the joint global distribution is the product of country-specific marginal distributions; (ii) those country-specific marginal distributions are all common; and (iii) the common marginal distribution takes the particular form of $Frechet(1, \theta^*)$.

2.5 Monopolistic competition and the Melitz (2003) model

Having described the general Ricardian model of trade under perfect competition, along with its Eaton and Kortum (2002) special case, we turn now to the case of monopolistic competition a la Melitz (2003). This is similar in many respects to the Ricardian model, so our discussion is brief.

Model setup. This model continues to feature D countries in which the representative household has CES preferences and hence demand follows equation (1). Production continues to be supplied by producers with marginal costs $MC_{od}(\omega)$ given by equation (3), on the basis of (variable) productivity draws $z_o(\omega)$ from an *arbitrary* global distribution that we continue to denote by F.[12]

However, there are two differences relative to the Ricardian model. First, market conduct is now monopolistically (rather than perfectly) competitive, and there are fixed costs of producing. Monopolistically competitive conduct implies, in the case

[11] That is, $F^*(\{z\}) = \prod_o F_o^*(z)$, where $F_o^*(z) = Pr(z_o(\omega) \leqslant z) = e^{-z^{-\theta^*}}$. We note that some studies, including Eaton and Kortum (2002), define $F_o^*(z)$ to include a scale parameter $T_o > 0$ that is o-specific, with $F_o^*(z) = e^{-T_o z^{-\theta^*}}$, rather than $T_o = 1$ as in our convention. Because we include the free parameters A_{od} in (3), and these parameters enter the model symmetrically to $z_o(\omega)$, our choice of $T_o = 1$ is without loss.

[12] We model tariffs as being applied to marginal cost for two reasons. First, this is standard in the literature (e.g., Costinot and Rodríguez-Clare, 2014). Second, following Caliendo et al. (2015), this ensures that in the workhorse gravity model case (i.e. where F is i.i.d. Pareto) we recover the canonical result that entry does not respond to tariffs.

of CES demand and in the limit where firms are too small to affect the price index, that if a producer from o sells good ω in d then the price $p_{od}(\omega)$ will be marked up over the producer's marginal cost by the amount $\sigma/(\sigma-1)$. In addition, a bilateral fixed cost $w_o f_{od}$ is incurred whenever a producer from o sells in d, so production will only take place where variable profits exceed such fixed costs.[13] Collecting these observations we have

$$p_{od}(\omega) = \begin{cases} \frac{\sigma}{\sigma-1} MC_{od}(\omega) & \text{if } X_{od}(\omega) > \sigma f_{od} w_o, \\ \infty & \text{otherwise.} \end{cases} \tag{13}$$

The inequality $X_{od}(\omega) > \sigma f_{od} w_o$ in (13) is equivalent to a condition that productivity $z_o(\omega)$ exceeds an endogenous pair-specific cutoff that we denote $\hat{z}_{od}$.[14] We impose two restrictions on the parameters that give rise to the baseline equilibrium. First, we require that $\hat{z}_{od} \geqslant \hat{z}_{oo}$ for all $d \neq o$, meaning that selection into exporting is active in the baseline. Second, we require that all bilateral cutoffs in the baseline equilibrium are above the lower bound of the support of F, which means that selection into production is active. These restrictions ensure that the model's trading equilibrium replicates standard selection patterns in Melitz (2003).[15]

The second difference brought about by monopolistic competition is that the equilibrium mass N_o of potential producers in o who choose to take a productivity draw $z_o(\omega)$ from F is endogenous. We assume that taking such a draw costs $w_o f_{Eo}$ and that a free entry condition mandates that the global profits a producer from o will obtain, in expectation, after drawing its productivity will be equal to the cost $w_o f_{Eo}$. This implies:

$$\sum_d \mathbb{E}_F \left[\max\{0, \frac{1}{\sigma} X_{od}(\omega) - w_o f_{od}\} \right] = w_o f_{Eo}. \tag{14}$$

Since aggregate profits are zero, household income in o continues to be given by $Y_o = w_o L_o$.

[13] As is standard, we assume that, in equilibrium, firms produce differentiated products. This means that there is only one producer per good ω in any destination market d.

[14] Specifically, $\hat{z}_{od}$ is defined as the value of z such that $X_{od}(\hat{z}_{od}) = \sigma f_{od} w_o$, where $X_{od}(z)$ denotes sales in d by of a firm in o with productivity z.

[15] Removing these restrictions would have no bearing on the distribution-agnostic GT_d bounds in Proposition 1(b) and serve only to widen the bounds we obtain within a finite neighborhood of the gravity model.

The market clearing and price index expressions of (6) and (7) are similar to the competitive model, though now augmented to include N_o,

$$\gamma_d = \sum_o N_o \mathbb{E}_F \left[p_{od}(\omega)^{1-\sigma} \right], \tag{15}$$

$$w_o L_o = \sum_d N_o \mathbb{E}_F \left[X_{od}(\omega) \right]. \tag{16}$$

Finally, we note that, in this model, total bilateral trade is given by $X_{od} = N_o \mathbb{E}_F \left[X_{od}(\omega) \right]$.

Summarizing, equilibrium in the monopolistically competitive model corresponds to a vector of global wages $\{w\}$, price indices $\{\gamma\}$, and masses of potential producers $\{N\}$ such that equations (1), (3) and (13)-(16) hold for all countries.

Gains from trade counterfactuals. Simulations of scenarios involving counterfactual tariffs $\{t'\}$, such as those that implement autarky for a given country of interest d, can be computed using a similar procedure to that discussed for the Ricardian model above. That is, we hold constant all exogenous primitives—endowments $\{L\}$, bilateral productivity shifters $\{A\}$, bilateral fixed costs of production $\{f\}$, entry costs $\{f_E\}$, preferences σ, and the global distribution of productivities F—and solve for the new equilibrium at counterfactual $\{t'\}$ in which all endogenous variables—wages $\{w\}$, price indices $\{\gamma\}$ and masses of potential firms $\{N\}$—adjust to the new equilibrium. We then compute the gains from trade for any country d following again the formula in equation (11), where $\{t'\}$ implements autarky for country d.

Reference distribution from Chaney (2008). As with the Ricardian model, our presentation of the monopolistic competition model a la Melitz (2003) has allowed for the global distribution of productivity levels F to take any form. However, since the pioneering work of Chaney (2008), it has become common to deploy one particularly tractable distribution: that in which each country o draws its productivity $z_o(\omega)$ independently from a Pareto distribution with scale parameter equal to one and a shape parameter equal to $\theta^* > 0$.[16] We therefore take this to be the reference distribution, denoted F^*, when studying the monopolistic competition model. As with the Ricardian model, relative to the unrestricted distribution F, the

[16] That is, $F^*(\{z\}) = \prod_o F_o^*(z)$, where $F_o^*(z) = Pr(z_o(\omega) \leqslant z) = 1 - z^{-\theta^*}$. Just like the competitive model discussed in footnote 11, there is no loss of generality in setting the scale parameter (and hence lower support) of this distribution to one.

reference distribution F^* imposes three separate types of restrictions: independence, common marginals, and particular (i.e. $Pareto(1, \theta^*)$) marginals.

3 Estimating the gains from trade

Calculating the gains from trade GT_d for any country d requires estimates of all of the model's exogenous features. These features include the observed inputs (tariffs $\{t\}$ and endowments $\{L\}$). They also include the model's unobserved parameters, such as the demand elasticity σ and bilateral technology shifters $\{A\}$) that apply to both competitive and monopolistically competitive models, plus the additional parameters for bilateral fixed costs $\{f\}$ and entry costs $\{f_E\}$ that apply in the case of the monopolistic competition model. Finally, and perhaps most fundamentally of all, the model's set of exogenous features includes the global distribution of productivities F that governs comparative advantage, trade possibilities, and hence the gains from trade.

3.1 Gravity models

Both of the special cases described in Section 2—the Eaton and Kortum (2002) version of the competitive model and the Chaney (2008) version of the monopolistic competition model—are distinctive because they work with a particular version F^* of the global productivity distribution F such that bilateral trade flows between countries take the so-called gravity form. Following Arkolakis et al. (2012) we therefore refer to these as gravity models.

In the case of gravity models, estimation of GT_d follows three well-known steps. Indeed, as shown by Arkolakis et al. (2012), when these three steps are executed on the same dataset the resulting value of GT_d will be identical for the Eaton and Kortum (2002) and Chaney (2008) models. We therefore discuss the general approach that applies to both of these models simultaneously here.

Step 1: Estimate θ^* from an orthogonality assumption. The first step is to estimate the single unknown parameter inside F^*, referred to in both models as θ^*. This begins with the assumption that the model predictions for each of the D^2 bilateral

trade flows X_{od} correspond to those in the data. Denoting values in the data by $\widetilde{X}_{od}$, this amounts to imposing the data-matching constraints

$$X_{od} = \widetilde{X}_{od} \quad \text{for all } o \text{ and } d, \tag{17}$$

with $X_{od} = \mathbb{E}_{F^*}[X_{od}(\omega)]$ in the competitive case and $X_{od} = N_o \mathbb{E}_{F^*}[X_{od}(\omega)]$ in the monopolistically competitive case.

Then, estimation of θ^* relies on the fact that gravity models predict that aggregate bilateral trade flows X_{od} will follow the gravity equation:

$$\ln X_{od} = \alpha_o + \alpha_d - \theta^* \ln(1 + t_{od}) + \varepsilon_{od}. \tag{18}$$

Here, α_o and α_d collect model terms—including endogenous ones such as the wage of the origin and the income of the destination—that are common to a given origin and destination, respectively; and ε_{od} collects all exogenous model terms that are bilateral, other than the tariffs t_{od}. For example, in the Eaton and Kortum (2002) case, we have $\alpha_o = -\theta^* \ln w_o$, $\alpha_d = \ln E_d - \ln \gamma_d$, and $\varepsilon_{od} = \theta^* \ln(A_{od})$. The Chaney (2008) case is similar but with the inclusion of the number of varieties and bilateral fixed costs: $\alpha_o = \ln N_o + \left(1 - \frac{\sigma}{\sigma - 1}\theta^*\right) \ln w_o$ and $\varepsilon_{od} = \theta^* \ln A_{od} + \left(1 - \frac{\theta^*}{\sigma - 1}\right) \ln f_{od}$.

A standard approach to estimating the parameter θ^* therefore involves using data on observed bilateral trade flows $\widetilde{X}_{od}$ and ad-valorem bilateral tariffs t_{od} in order to estimate equation (18) by OLS. One way to partial-out the endogenous terms α_o and α_d is to double-difference equation (18) with respect to sample means. Letting, for example, $\Delta_o y_{od} \equiv y_{od} - \frac{1}{D}\sum_d y_{od}$ for any variable y_{od}, this procedure amounts to an OLS regression of $\Delta_o \Delta_d \ln \widetilde{X}_{od}$ on $\Delta_o \Delta_d \ln(1 + t_{od})$.[17] The coefficient from this regression will provide a consistent estimate of $-\theta^*$ when the following orthogonality condition holds:

$$\sum_o \sum_d \Delta_o \Delta_d \ln(1 + t_{od}) \times \Delta_o \Delta_d \varepsilon_{od} = 0. \tag{19}$$

In the case of the Eaton and Kortum (2002) model the orthogonality condition (19) requires the assumption that tariffs are orthogonal to bilateral productivity shifters in the following sense

[17] In cases with no missing observations this is algebraically identical to an OLS regression of $\ln X_{od}$ on $\ln(1 + t_{od})$ and origin and destination fixed effects.

$$\sum_o \sum_d \Delta_o \Delta_d \ln(1 + t_{od}) \times \Delta_o \Delta_d \ln A_{od} = 0. \tag{20}$$

Similarly, consistent estimation for the Chaney (2008) case follows under the assumption that bilateral fixed costs are also orthogonal to tariffs:

$$\sum_o \sum_d \Delta_o \Delta_d \ln(1 + t_{od}) \times \Delta_o \Delta_d \ln f_{od} = 0. \tag{21}$$

Step 2: Estimate other model parameters by matching all trade flows. Armed with an estimate of θ^* from step 1, the second step uses this to obtain estimates of the remaining model parameters of interest. Given the data-matching constraint in (17), each of the D^2 bilateral residuals in equation (18)—i.e. $\ln \widetilde{X}_{od} + \theta^* \ln(1 + t_{od})$—provides an unbiased estimate of each of the D^2 terms ε_{od}, up to two irrelevant o-specific and d-specific scale factors. These estimated values of ε_{od} then provide (along with θ^*) unbiased estimates of the remaining model parameters that are necessary for counterfactual simulations (such as computing GT_d). For example, in the Eaton and Kortum (2002) model, since $\varepsilon_{od} = \theta^* \ln(A_{od})$, the translation from each bilateral residual to estimates of each bilateral productivity parameter A_{od} is immediate. We refer to the estimated versions, in this case where F^* is assumed, as A^*_{od} below.

Similarly, standard arguments in the Chaney (2008) case imply that the same combination of A_{od} and f_{od} that is identified from ε_{od} is sufficient for all equilibrium calculations. Finally, a remaining unknown model parameter in both cases is σ but in both models this turns out to only affect the scale of consumer welfare, so it cancels out of GT_d.

Step 3: Solve for counterfactual changes in the estimated model. The final step is to solve for GT_d, as defined in equation (11), according to the estimated model.[18] In the case of the Eaton and Kortum (2002) version of the gravity model, this amounts to solving the system of equations (6) and (7) for the factual equilibrium and (9) and (10) for the counterfactual one, each time with the expectations $\mathbb{E}_F[\cdot]$ evaluated at $F = F^*$ (and where F^* corresponds to independent draws from a $Frechet(1, \theta^*)$ distribution at the estimated value of θ^* obtained from step 1). Similarly, in the Chaney

[18] The second and third steps of this procedure are often combined into an algebraically equivalent single step via the "exact-hat" algebra procedure used by Dekle et al. (2008).

(2008) case, step 3 requires solving equations (14)-(16) for the factual equilibrium, and their counterfactual equilibrium analogs, while setting the distribution F^* to be independent draws from $Pareto(1, \theta^*)$ at the estimated value of θ^*. Either way, as discussed previously, the Eaton and Kortum (2002) and Chaney (2008) models are empirically isomorphic—so whenever steps 1 and 2 have been executed on the same trade and tariff data, these two models return the same value of GT_d.

Indeed, as Arkolakis et al. (2012) have shown, under F^* the result of step 3 for gravity models is given simply by the formula

$$GT_d^{\text{ACR}} \equiv 1 - \widetilde{\lambda}_{dd}^{\frac{1}{\theta^*}}, \tag{22}$$

where $\widetilde{\lambda}_{dd} \equiv \widetilde{X}_{dd} / \sum_o \widetilde{X}_{od}$. This is because gravity models feature a constant import demand elasticity $\epsilon_d(\cdot)$ of θ^*. Substituting this into the area under the import demand curve in equation (12) yields the formula in (22).

3.2 Beyond gravity models: distribution-agnostic bounds

We have seen how estimation of GT_d works in the case of the Eaton and Kortum (2002) version of the Ricardian model and the Chaney (2008) version of the monopolistic competition model. That is, in either case, we have seen how to estimate GT_d for one particular version, denoted F^*, of the global productivity distribution F. This raises two natural questions: How could the three-step estimation procedure of Section 3.1 be repeated for alternative distributions $F \neq F^*$? And how would GT_d change as these alternative distributions are considered?

One direction would be to repeat the analogs of steps 1-3 for particular alternative distributions. For example, one could imagine working with the Ricardian case and positing that the global distribution F corresponds to independent draws from a log-normal marginal distribution, with location parameter zero and scale parameter s, that is common to all countries. This would no longer be a gravity model. But the assumption that log tariffs $\ln(1 + t_{od})$ are mean independent of the log bilateral productivity shifter $\ln A_{od}$ could nevertheless be incorporated, via a simulated method of moments procedure, to estimate s and all of the values of A_{od} (up to scale) as residuals. Having done that, the model could be solved—evaluating the expectations $\mathbb{E}_F[\cdot]$ by simulation, since they are not available in closed form

for this choice of F—and GT_d could be calculated for this particular version of the global productivity distribution.

A concern with this sort of exercise is that it (unavoidably) starts with a particular choice of F, and that choice may be difficult to motivate. In this section we therefore characterize the bounds on gains from trade that could obtain across any choice of F, and refer to these as distribution-agnostic bounds. We begin with the Ricardian case.

Bounds on GT_d in the competitive model. We seek to estimate sharp bounds on the size of country d's gains from trade GT_d that are consistent with the same three-step estimation method used for gravity models in Section 3.1, but now done for all productivity distributions F rather than for the particular F^* assumed by Eaton and Kortum (2002).

To do so, we proceed as follows. Consider the upper bound on GT_d, which we denote by $\overline{GT}_d$. This is given by the solution to the following program, where the counterfactual tariffs $\{t'\}$ are those that would place d in autarky (i.e. $t_{od} = \infty$ for all $o \neq d$):

$$\overline{GT}_d \equiv \max_{F,\{A\},\{w\},\{w'\},\{\gamma\},\{\gamma'\}} \quad 1 - \left(\frac{\gamma_d}{\gamma'_d}\right)^{\frac{1}{1-\sigma}} \tag{23}$$

$$\text{subject to: } \begin{cases} \text{equilibrium } \{w\} \text{ and } \{\gamma\} \text{ hold at } \{t\}, \\ \text{equilibrium } \{w'\} \text{ and } \{\gamma'\} \text{ hold at } \{t'\}, \\ \mathbb{E}_F[X_{od}(\omega)] = \widetilde{X}_{od} \text{ for all } o \text{ and } d, \\ \text{and } \sum_o \sum_d \Delta_o \Delta_d \ln(1 + t_{od}) \times \Delta_o \Delta_d \ln A_{od} = 0. \end{cases}$$

That is, estimating $\overline{GT}_d$ amounts to finding the version of the Ricardian model whose productivity distribution F is chosen to maximize the drop in welfare for country d when it moves from the initial equilibrium to autarky as per equation (11), while ensuring that: (i) factual wages $\{w\}$ and price indices $\{\gamma\}$ satisfy the factual equilibrium equations (1), (3), (4), (6) and (7); (ii) counterfactual wages $\{w'\}$ and price indices $\{\gamma'\}$ satisfy the counterfactual equilibrium equations (8)-(10); (iii) each model-implied bilateral trade flow in the initial equilibrium matches that in the data $\widetilde{X}_{od}$; and (iv) model-consistent bilateral productivity shifters $\{A\}$ are orthogonal to tariffs in the sense of equation (20). The analogous lower-bound, which we denote

by $\underline{GT}_d$, involves an identical program to that in (23) but with the min operator instead. The following proposition—whose proof is in Appendix A along with all others—describes these upper and lower bounds.

Proposition 1(a) (Bounds for the competitive model). *Consider all possible productivity distributions F. Then:*

(i) ***Agnostic Bounds.*** *The smallest and largest possible values for a given destination d's gains from trade in the competitive model, as defined in program (23), are given by*[19]

$$\underline{GT}_d = 0 \qquad \text{and} \qquad \overline{GT}_d = 1 - \widetilde{\lambda}_{dd}^{\frac{1}{\sigma-1}}. \tag{24}$$

(ii) ***Sharpness.*** *Suppose that the baseline data on trade shares satisfy, for country d of interest, the condition: $\widetilde{\lambda}_{dd}/\widetilde{\lambda}_{dn} \geqslant \widetilde{\lambda}_{od}/\widetilde{\lambda}_{on}$ for all $o, n \neq d$. Then, there exists a distribution F that attains the agnostic bounds in (24).*

Before turning to the intuition behind this proposition, we first comment on the data condition in the premise of part (ii). This condition is easily verified in any given dataset and is indeed satisfied for all countries in our application below. This is the case because the diagonals of the trade share matrix (i.e. λ_{dd}, which describe countries' purchases from themselves) tend to be much larger than other elements in the matrix because most countries are not very open to trade.

The intuition behind part (i) of Proposition 1(a) draws on our earlier discussion about what determines the area under the import demand curve, following equation (12). We begin with the lower bound in (24). When the import demand elasticity $\epsilon_d(\kappa) = -d\log(\sum_{o\neq d} X_{od}(\kappa)/X_{dd}(\kappa))/d\log\kappa$ achieves its maximal value of infinity at every point, one can verify that the area $\mathcal{A}_d$ is zero and hence so too are the gains from trade. Naturally, this corresponds to a case in which imports are perfect substitutes for domestic goods, which achieves the lower bound in which there would be no gains at all (i.e. $\underline{GT}_d = 0$). This lower-bound import demand curve—and hence minimal GT_d—case is illustrated with the orange line in Figure 1.

[19] We show in Appendix A.1.1 that analogous bounds obtain for a version of the Ricardian model with multiple sectors and roundabout input-output linkages (as in Caliendo and Parro (2015)). The upper bound becomes $\overline{GT}_d^{\text{MSIO}} = 1 - \prod_{s=1}^{S} \widetilde{\lambda}_{dd,s}^{\beta_{d,s}/(\sigma_s-1)}$, where $\lambda_{dd,s}$ is the share of d's expenditure on sector s that is domestically sourced and $\beta_{d,s}$ is the standard exponent combining the Cobb-Douglas preference parameter and the Leontief inverse of the input-output matrix. Thus, to the extent that these features amplify the gains from trade when F is Frechet, they also amplify the distribution-agnostic bounds.

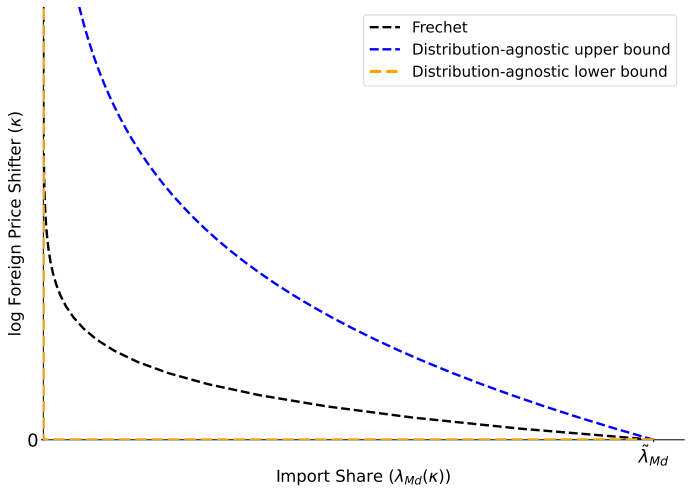


Figure 1: Import demand curves for the competitive model

Turning to the upper bound in (24), define $p_{Md}(\omega) \equiv \min_{o \neq d} MC_{od}(\omega)$ as the best foreign price for good ω in destination d, and write $I(\omega;\kappa) \equiv \mathbb{1}\{\kappa p_d^M(\omega) < p_{dd}(\omega)\}$ as an indicator that equals 1 if d imports good ω. Then, substituting the stochastic representation of demand $\mathbb{E}_F[X_{od}(\omega)]$ into the definition of $\epsilon_d(\kappa)$ yields[20]

$$\epsilon_d(\kappa) = \underbrace{\sigma - 1}_{\text{intensive margin}} + \underbrace{\frac{1}{1-\lambda_{dd}(\kappa)} \frac{d \log \mathbb{E}_F[p_{dd}(\omega)^{1-\sigma}(1 - I(\omega;\kappa))]}{d \log \kappa}}_{\text{extensive margin}}. \tag{25}$$

The first term in this equation captures the intensive margin: increasing the price of any good that d imports affects expenditure with an elasticity of $\sigma - 1$ under CES preferences. The second term captures the extensive margin: increasing the price of imports causes d to begin sourcing goods domestically that it previously imported. Since the extensive margin can never be negative, the lowest possible value that $\epsilon_d(\kappa)$ can take, at any κ, is therefore $\sigma - 1$. Substituting this into the area formula (12) verifies that the largest possible gains from trade—when there is no extensive margin of substitution between goods that are imported and produced domestically—must therefore be $\overline{GT}_d = 1 - \tilde{\lambda}_{dd}^{\frac{1}{\sigma-1}}$. Figure 1 illustrates this case with

[20] Since $\epsilon_d(\kappa) = -d \log \mathbb{E}_F[(\kappa p_{Md}(\omega))^{1-\sigma} I(\omega;\kappa)] / d \log \kappa + d \log \mathbb{E}_F[p_{dd}(\omega)^{1-\sigma}(1 - I(\omega;\kappa))] / d \log \kappa$, the intensive margin follows immediately from the shifter in the first term. Then, because κ enters the foreign and domestic terms symmetrically through $I(\omega;\kappa)$ and $(1 - I(\omega;\kappa))$, and foreign and domestically prices are equal in the neighborhood of a sourcing switch, the two extensive margin components can be combined to obtain the single term in (25).

the blue line.

In between the two extremes of $\underline{GT}_d$ and $\overline{GT}_d$ lies the gravity model case discussed in Section 3.1, in which the import demand elasticity is constant. For the Eaton and Kortum (2002) case of $F = F^*$, Arkolakis et al. (2012) show that the extensive margin elasticity in (25) is equal to $\theta^* - \sigma + 1$ and hence that $\epsilon_d(\kappa) = \theta^*$ for all κ. Following equation (32) this results in gains of GT_d^{ACR}. The black line in Figure 1 illustrates this case.

We turn now to the intuition behind the proof of part (ii) of Proposition 1(a). Like all other results about sharpness below, our proof of this statement is constructive—it provides an example of two distributions, denoted F^+ and F^-, that approach the bounds arbitrarily closely.[21] Beginning with F^+, the above discussion highlights that GT_d is maximized if and only if the extensive margin is inactive as κ goes all the way from 1 to ∞—that is, d must never switch from foreign to domestic sourcing for goods that it imports at baseline. Intuitively, this happens if and only if d's productivity when producing such goods is arbitrarily close to zero.[22] Our construction F^+ designs a global productivity distribution that has this property.

However, the key challenge that remains is to ensure that F^+ is capable of matching the constraints in program (23)—generating an equilibrium that matches all baseline trade flow values and the orthogonality condition (20). We do so as follows. First, we begin with the distribution F^*. As discussed in Section 3.1, this distribution is able to match all required constraints when the bilateral productivity shifters A_{od} take the values A_{od}^*. We will hold $A_{od} = A_{od}^*$ for all o and d in what follows, which ensures that the orthogonality condition holds. Second, we build F^+ from F^* by altering the copula of F^*—introducing negative dependence between the

[21] Since, as described in Section 3.3, we are primarily interested in continuous distributions, in some cases it is not possible to achieve the bounds with equality. We therefore say that a given bound is attained if, for any $\epsilon > 0$, there exists a distribution F_ϵ such that $|GT_d^{F_\epsilon} - GT_d^{\text{bound}}| < \epsilon$. For upper bounds, this means $GT_d^{F_\epsilon} > \overline{GT}_d - \epsilon$. For lower bounds, this means $GT_d^{F_\epsilon} < \underline{GT}_d + \epsilon$.

[22] Seen another way, one can combine (i) the definition of λ_{dd}, (ii) the fact that autarky requires $\gamma'_d = \mathbb{E}_F[MC_{dd}(\omega)^{1-\sigma}]$, and (iii) the choice of numeraire to write equation (11) as

$$GT_d \;=\; 1 - \lambda_{dd}^{\frac{1}{\sigma-1}} \left(1 + \frac{\mathbb{E}_F\left[z_d(\omega)^{\sigma-1} \mathbb{1}\{d \neq \arg\min_o MC_{od}(\omega)\}\right]}{\mathbb{E}_F\left[z_d(\omega)^{\sigma-1} \mathbb{1}\{d = \arg\min_o MC_{od}(\omega)\}\right]} \right)^{\frac{1}{\sigma-1}}. \tag{26}$$

The ratio of expectations is minimized when the numerator is zero, in which case the upper bound $\overline{GT}_d$ is immediate. This occurs when the expectation of $z_d(\omega)^{\sigma-1}$, an increasing transformation of d's productivity, conditional on the event of d not being the lowest-cost supplier to d, is zero.

productivity draws that d receives and those that its foreign partners receive—such that d's draws approach zero for any draws that would be low enough to be unable to compete with imports at baseline. This ensures that gains are maximized. Third, we observe that whenever two models (such as those with F^* and F^+) match the same trade flows at baseline they will agree on equilibrium baseline wages; we therefore denote these by $\{w^*\}$. Finally, we prove that, provided the data condition in Proposition 1(a) holds, all bilateral trade flows will be unaffected by this change in the copula and hence that all trade values will be matched by F^+.[23]

Turning to the lower bound distribution F^-, this construction also follows the intuition in Figure 1. We know that $\overline{GT}_d$ is achieved whenever the import demand curve is perfectly elastic at all points on the path to autarky. This occurs in the opposite type of scenario from the upper bound case: it requires that country d can produce all imported goods at exactly the same tariff-inclusive prices as occur at baseline, and hence that imports merely reflect indifference.[24] Our construction of F^- follows the same recipe as F^+, but this time by adjusting the copula of F^* so as to create a form of positive dependence that generates this indifference.

We conclude with a final remark about the constructions F^+ and F^- described above. While one might conjecture that the orthogonality and data-matching constraints may place an aggregate or average constraint on the worldwide gains from trade, this turns out to be incorrect. In particular, the steps involved in these constructions could be executed simultaneously for all countries. This means, for example, that country d's gain can be maximized without precluding the possibility that other countries also achieve their own maximal levels.

[23]This succeeds because the only difference between F^* and F^+ occurs among the goods that d imports. Hence the only threat to successful data matching that could arise would be that d actually exports some of these goods that it does not sell to itself. The data condition rules this out under F^* (at $\{A^*\}$ and $\{w^*\}$) and hence also under F^+ (at the same $\{A^*\}$ and $\{w^*\}$).

[24]Again, this can also be seen directly by combining the definitions of γ_d in (6) and GT_d in (11), along with the fact that $\gamma'_d = \mathbb{E}_F[p_{dd}(\omega)^{1-\sigma}]$ in autarky, to obtain

$$GT_d = 1 - \left(\frac{\mathbb{E}_F[MC_{dd}(\omega)^{1-\sigma}]}{\mathbb{E}_F[(\min_o MC_{od}(\omega))^{1-\sigma}]} \right)^{\frac{1}{\sigma-1}}, \tag{27}$$

which takes the value of zero when the two expectations—that for country d's own marginal cost and that of the cheapest global supplier including itself—are equal.

Bounds on GT_d in the monopolistic competition model. Having characterized the competitive model case, the problem of obtaining distribution-agnostic bounds for the monopolistic competition model proceeds similarly. That is, we seek to bound GT_d for the Melitz (2003) model across all productivity distributions F rather than the particular version F^* that was used in the benchmark gravity model case of Chaney (2008). Yet we do so while replicating the standard three-step estimation procedure, which ensures that all bilateral trade flow data points are matched by the model at any of the distributions F under consideration, as well as that bilateral productivity shifters and fixed costs are orthogonal to tariffs in the sense of equations (19) and (21).

The program for maximizing GT_d in the monopolistic competition model is similar to that in (23), but now written as

$$\overline{GT}_d \equiv \max_{F,\{A\},\{f\},\{f_E\},\{w\},\{w'\},\{\gamma\},\{\gamma'\},\{N\},\{N'\}} 1 - \left(\frac{\gamma_d}{\gamma'_d}\right)^{\frac{1}{1-\sigma}} \tag{28}$$

$$\text{subject to: } \begin{cases} \text{equilibrium } \{w\}, \{\gamma\} \text{ and } \{N\} \text{ hold at } \{t\}, \\ \text{equilibrium } \{w'\}, \{\gamma'\} \text{ and } \{N'\} \text{ hold at } \{t'\}, \\ N_o \mathbb{E}_F[X_{od}(\omega)] = \widetilde{X}_{od} \text{ for all } o \text{ and } d, \\ \sum_o \sum_d \Delta_o \Delta_d \ln A_{od} \times \Delta_o \Delta_d \ln(1 + t_{od}) = 0, \\ \sum_o \sum_d \Delta_o \Delta_d \ln f_{od} \times \Delta_o \Delta_d \ln(1 + t_{od}) = 0, \\ \hat{z}_{od} \geqslant \hat{z}_{oo} \text{ for all } o, d, \\ \text{and } \hat{z}_{oo} \geqslant 1 \text{ for all } o. \end{cases}$$

Relative to the Ricardian model, the new considerations of the monopolistically competitive model naturally require an augmented set of model equations, as laid out in Section 2.5. It also requires a larger set of model parameters, due to the addition of bilateral fixed costs of production $\{f\}$ and entry costs $\{f_E\}$. In addition, the factual and counterfactual equilibrium masses of potentially producing firms, $\{N\}$ and $\{N'\}$, must be solved for at every candidate distribution F under consideration alongside the corresponding market-clearing wages $\{w\}$ and $\{w'\}$ and price indices $\{\gamma\}$ and $\{\gamma'\}$. Finally, we append the orthogonality condition for bilateral fixed costs (21) and the conditions that selection into exporting and production are always active in the baseline.

The following proposition describes how the bounds $\underline{GT}_d$ and $\overline{GT}_d$ in the monopolistic competition model turn out to be identical to those in the competitive model (though they are attainable under a different data condition).

Proposition 1(b) (Bounds for the monopolistic competition model). *Consider all possible productivity distributions F. Then:*

(i) ***Agnostic Bounds.*** *The smallest and largest possible values for a given destination d's gains from trade in the monopolistic competition model, as defined in program (28), are given by*

$$\underline{GT}_d = 0 \qquad \textit{and} \qquad \overline{GT}_d = 1 - \widetilde{\lambda}_{dd}^{\frac{1}{\sigma-1}}. \tag{29}$$

(ii) ***Sharpness.*** *Whenever there are at least three countries (i.e. $D \geqslant 3$) there exist distributions F^- and F^+ that attain the agnostic bounds in (29).*

To see the intuition for part (i), begin with the lower-bound. This follows simply from noting that—as observed by Melitz and Redding (2013) and Head et al. (2014), for example—it is always the case in the monopolistic competition model that gains from trade are a function of only the ratio of counterfactual and factual cutoff productivity levels for selling domestically in country d, or $GT_d = 1 - \hat{z}'_{dd}/\hat{z}_{dd}$. We then establish that the autarky cutoff $\hat{z}'_{dd}$ can never exceed the baseline cutoff $\hat{z}_{dd}$, and hence that losses from trade are impossible in the model described in program (28), i.e. $\underline{GT}_d = 0$.

Establishing the upper bound in (29) proceeds from routine derivations that allow us to write

$$GT_d = 1 - \widetilde{\lambda}_{dd}^{\frac{1}{\sigma-1}} \left[\frac{1 + RLP_d}{1 - EFCS_d} \right]^{\frac{1}{\sigma-1}}.$$

Here, RLP_d refers to a measure of "relative latent productivity" that (as discussed formally in Appendix A) summarizes how the productivity distribution F among d's "latent" firms (those who only produce in autarky) compares to that for the firms in d that are active at baseline. Trade selects positively, so this term cannot be negative. Similarly, $EFCS_d$ refers to the "export fixed cost share," defined as the ratio of (a) all fixed costs of exporting paid by firms in d at baseline to (b) all fixed and entry costs paid by firms in d at baseline. Naturally, this share is bounded

between zero and one. Put together, the term in square brackets is therefore never below one, hence GT_d cannot exceed $\widetilde{\lambda}_{dd}^{\frac{1}{\sigma-1}}$, as claimed in (29).

What productivity distributions, F^- and F^+, could achieve these bounds? The lack of head-to-head competition means that constructing such distributions is much simpler in the monopolistic competition case than it was for the competitive case above. In particular, both RLP_d and $EFCS_d$ in (3.2) are functions of country d's marginal productivity distribution alone. And if this marginal is Pareto with shape parameter θ_d then—following the arguments in Arkolakis et al. (2012)—the lower bound can be achieved by $\theta_d \to \infty$ and the upper bound by $\theta_d \to \sigma - 1$. In the lower bound case, d's firms become clustered around the same productivity level, so the cutoff $\hat{z}_{dd}$ does not move from baseline to autarky and the lower bound is achieved. And in the upper bound case, the distribution in d is so skewed that latent firms are arbitrarily unproductive relative to active firms (so $RLP_d \to 0$) and very few firms export (so $EFCS_d \to 0$) at baseline, which maximizes GT_d in (3.2).

What remains is to propose marginal distributions for all other countries $o \neq d$, along with a cross-country copula, such that the global distributions F^- and F^+ can always satisfy the orthogonality and data-matching constraints in (28) as we vary θ_d. The proof of part (ii) establishes that, as long as $D \geqslant 3$, values of $\{A\}$ and $\{f\}$ can always be found such that country d's own contribution to the summations in orthogonality constraints (20) and (21) is zero.[25] Hence a particularly simple construction that works (for both F^- and F^+) is one where all countries draw independently from Pareto distributions, with shape parameters of θ^* for all countries other than d. By design, the summation constraints (20) and (21) are satisfied by the model at F^* among all countries, so they continue to be satisfied by the model at F^- and F^+ since these differ from F^* only in terms of d's marginal.

3.3 Beyond gravity models: distributions "close" to gravity

Our analysis has so far developed bounds on the gains from trade that are possible if one allows for any productivity distribution F. As we discuss below, these bounds are many times wider than those obtained under the distribution F^* assumed by gravity models. To explore this gulf further we now describe a procedure for

[25] In practice, we prove a slightly stronger result that the result remains true whenever a condition involving the tariff data is satisfied. This condition always holds when $D \geqslant 3$ but can also hold at $D = 2$.

estimating sharp on GT_d that hold for neighborhoods of distributions F that are "close" to the reference distribution F^*. Varying the size of the neighborhood will permit a systematic investigation of the possibilities as one considers larger and larger departures from the workhorse case at F^*.

Neighborhoods of the F^* distribution. To begin, we use a standard method for comparing two distributions from the statistics literature. In particular, the "ϕ-divergence" (henceforth divergence) $D(F||F^*) \geqslant 0$ of F from F^* is defined as

$$D(F||F^*) \equiv \int_{\{z\}\in\mathbb{R}_D^+} \phi\left(\frac{dF(\{z\})}{dF^*(\{z\})}\right) dF^*(\{z\}), \tag{30}$$

for a suitable convex function $\phi(\cdot)$ with $\phi(1) = 0$.[26] That is, $D(F||F^*)$ applies the penalty function $\phi(\cdot)$ to deviations of F from F^* at each point $\{z\}$ in the support of F^* and then integrates over all such penalties, weighted according to the distribution F^*. Several conventional choices of $\phi(\cdot)$ are available, such as the Kullback-Leibler (KL) form of $\phi(x) \equiv x\ln x - x + 1$ and the χ^2 form of $\phi(x) \equiv (x-1)^2/2$. The particular definition of divergence that we use below is the "hybrid" divergence, recommended by Christensen and Connault (2023) for cases like ours where F^* has fat tails; this uses KL divergence for all $x \leqslant e$ and χ^2 divergence for $x > e$.[27] Based on this definition of divergence, we define the neighborhood $\mathcal{N}_\delta$ of distributions around F^* as the set of all distributions F for which $D(F||F^*) \leqslant \delta$.

There are three attractive features of divergence for this purpose. First, it does not depend on the choice of units for the random variable $z_o(\omega)$. Second, transformations of location and/or scale in F^* do not affect divergence. Finally, as $\delta \to \infty$, any distribution that is absolutely continuous with respect to F^* can be approximated arbitrarily well by distributions contained within $\mathcal{N}_\delta$. Since the Eaton and

[26] Formally, divergence is given by (30) for any F that is absolutely continuous with respect to F^*, and ∞ otherwise. A distribution F is absolutely continuous with respect to F^*, written $F \ll F^*$, if any event to which F^* assigns zero probability is also assigned zero probability by F; i.e., F can be obtained by reweighting the density of F^* within its support. As discussed below, for the F^* in our applications, essentially any continuous distribution with density on the positive real number line satisfies $F \ll F^*$.

[27] The advantage of the hybrid divergence is that it combines the quadratic upper tail of χ^2—whose regularity condition is easily satisfied in our application—with the log x behavior near 0 of KL that ensures that the marginal penalty goes to ∞ as x goes to 0. The χ^2 component is re-normalized slightly to $\phi(x) = (x-e)^2/2e + (x-e) + 1$, which retains the quadratic penalty while ensuring that both components agree on the values of $\phi(x), \phi'(x)$ and $\phi''(x)$ at $x = e$.

Kortum (2002) version of F^* has full support over positive values, and the Chaney (2008) version has full support above a value (of one) that is set without loss since it is the Pareto distribution's scale parameter, this means that all distributions with density on the positive real line can be considered as $\delta \to \infty$.

Gains from trade within finite neighborhoods of F^*. Armed with the previous definition of a neighborhood $\mathcal{N}_\delta$ of distributions F around F^* we define the gains from trade bounds within such a neighborhood as

$$\overline{GT}_d(\delta) \equiv \max_F \overline{GT}_d \quad \text{s.t.} \quad F \in \mathcal{N}_\delta. \tag{31}$$

That is, the finite-neighborhood bound $\overline{GT}_d(\delta)$ is simply the same optimization problem as earlier—program (23) in the competitive case and (28) in the monopolistically competitive case—but subject to the added constraint that the only admissible productivity distributions F must lie within divergence δ of F^*.[28] This means that the analytical bounds in Section 3.2 are approached in the $\delta \to \infty$ limit: $\overline{GT}_d = \overline{GT}_d(\delta = \infty)$ and $\underline{GT}_d = \underline{GT}_d(\delta = \infty)$.

While the program in (31) is straightforward to state, obtaining its solution is far from trivial. For any $\delta > 0$ the set of distributions in $\mathcal{N}_\delta$ to be considered is infinite, so conducting a search in the space of distributions F is clearly infeasible. And it is unclear how one could obtain analytical derivations like those of Section 3.2 for this finite-neighborhood case. Fortunately, however, Christensen and Connault (2023) have recently developed tools for making computational progress on problems of this sort, showing that a particular sub-problem of (31) has a dual that is feasible.[29]

[28] In the case of the monopolistic competition model we impose an additional constraint on the level of the bilateral fixed costs, $\{f\}$. The special property that Pareto distributed random variables retain the same distribution when truncated from below implies that multiplying all of the $\{f\}$ by a common scalar $c > 1$ (and suitably adjusting $\{A\}$) does not affect aggregate quantities under F^*. However, as this shifts all cutoffs rightwards, it shrinks the economically relevant part of the support of F^*, which affects divergence. Indeed, we show in Lemma 2 in Appendix A that the divergence needed to attain a given value of GT_d can be made arbitrarily small by taking $c \to \infty$ and thereby concentrating all of the economic activity in the tail. Thus, for program (31) to be of interest in the monopolistic competition model, it is necessary to remove the degree of freedom that arises from the scale of $\{f\}$. We set f_{dd} to be such that the autarky cutoff is 1, which ensures that the full support of F^* is economically relevant, as it is under the Ricardian model.

[29] The dual formulation is valid under the quadratic divergence tail that we use if, for any constraint $\mathbb{E}[g(z(\omega))]$ in (23) and (28), the constraint function $g(z(\omega))$ satisfies $\mathbb{E}_{F^*}[g(z(\omega))^2] < \infty$. In both models that we study, this condition is equivalent to $\theta^* > 2(\sigma - 1)$, which is easily satisfied

In particular, their result refers to cases that involve maximizing, over all distributions F in $\mathcal{N}_\delta$, an objective function that can be written as (a known function of) an expectation over F subject to any number of constraints that can also be written as (known functions of) expectations over F. In such a case, Christensen and Connault (2023) establish that the dual of this infinite-dimensional program is concave and finite-dimensional, with dimensionality given by the number of constraints, and where all expectations in this dual program are expectations over F^*.[30]

The program in (31) does not immediately fit this desired form where the objective function and all constraints involve only expectations over F. For example, the choice set in $\overline{GT}_d$, as defined in equation (23) for the competitive model case, directly involves model parameters such as the productivity shifters, $\{A\}$. Similarly, the factual and counterfactual equilibrium objects $(\{w\}, \{\gamma\})$ and $(\{w'\}, \{\gamma'\})$ enter the objective function (11) and constraints in nonlinear ways that disrupt the pure expectations form. To circumvent these challenges, Christensen and Connault (2023) propose that one treat all such variables—not just primitives such as $\{A\}$, but also endogenous variables such as $\{w\}$, $\{w'\}$, $\{\gamma\}$ and $\{\gamma'\}$ in our context—as an expanded set of "parameters", which we denote ψ. After doing so, the program in (23) can be split into two stages. First, for any given value of the parameter vector ψ, solve the maximization program over $F \in \mathcal{N}_\delta$ by converting it from an infeasible infinite-dimensional program into a feasible finite-dimensional program. Second, optimize over ψ in an outer loop.

Implementation details. We discuss how we implement this two-step optimization program numerically in detail in Appendix B. In brief, we implement both the inner and outer loop using Artelys Knitro, with expectations computed with quasi-Monte Carlo integration with 500,000 draws constructed using Sobol sequences. The inner loop is convex and converges in few iterations when supplied with exact gradient and Hessian, which are available in closed form and can be constructed

in our application.

[30] We provide an intuitive exposition in Appendix B.1 of how the infinite-dimensional search over F can be replaced with a finite-dimensional convex program. In brief, as any $F \ll F^*$ can be constructed by re-weighting the density of F^*, searching over F is equivalent to searching over weights, $dF(\{z\})/dF^*(\{z\})$. Then, since ϕ imposes a strictly convex penalty as the weights depart from 1, the optimal weights can be solved for in closed-form as a function of the Lagrange multipliers governing the constraints. The dual therefore amounts to searching over the Lagrange multipliers, which is finite-dimensional and convex.

efficiently by exploiting features of the economic models (such as the sparsity of $X_{od}(\omega)$ that arises from perfect competition in the Ricardian model and selection in the monopolistically competitive model). The outer loop is nonconvex and—after profiling out nuisance parameters and imposing appropriate normalizations—we use several gradient-based algorithms initialized at randomized starting points.[31]

3.4 Data

Our analysis focuses on data from 2018. We obtain: aggregate bilateral trade data $\tilde{X}_{od}$ from the OECD ICIO database; tariff data t_{od} from Teti (2024), constructing a measure of the average bilateral tariffs by weighting 6-digit HS-level tariffs by their corresponding flows from the CEPII BACI database; and total labor force data L_o from the World Bank's WDI database. In order to keep the dimensionality manageable, we work with a sample of countries that includes the 19 largest countries by total expenditure and aggregate all remaining countries into a Rest of World group. This results in $D = 20$.[32]

4 Sensitivity of quantitative trade models

We turn now to the following question: how sensitive are the conclusions that standard quantitative models make about the gains from trade to changes in the underlying distribution of productivities? Our focus lies on estimates of GT_d for the case where d is France, a country with import penetration that is representative among the countries in our sample.

[31] For the monopolistically competitive model, we also prove in Proposition 5 that it is without loss to consider a substantially lower-dimensional problem that contains only A_{od} and f_{od} where $o =$ France, and only the trade flow constraints that require matching $\tilde{X}_{od}$ for $o =$ France. Intuitively, the lack of head-to-head competition means that the contribution to France's gains from trade stemming from parameters with $o \neq$ France is fully summarized by baseline trade shares. We can therefore always find a vector of $\{A\}$ and $\{f\}$ *ex post* that rationalizes the (omitted) trade flows $\tilde{X}_{od}$ for $o \neq$ France by inverting the mapping from parameters to trade flows under the extremizing distribution. Finally, we show that there is sufficient flexibility among the omitted $\{A\}$ and $\{f\}$ that this inversion can be conducted in a way that satisfies the orthogonality restrictions.

[32] The full list of countries is Australia, Brazil, Canada, China, France, Germany, India, Indonesia, Italy, Japan, Mexico, the Netherlands, Russia, South Korea, Spain, Switzerland, Turkey, the United Kingdom, and the United States.

4.1 Gains from trade in the gravity model

Workhorse quantitative trade models of the gravity form, such as Eaton and Kortum (2002) and Chaney (2008), predict how the welfare of a country d, such as France, would change if the country's tariffs were to rise to the point of autarky. These predictions about GT_d are made by employing the model estimation procedure detailed in Section 3.1 for a version of the model in which global production productivities F are drawn from the particular distribution F^*—the i.i.d. $Frechet(1, \theta^*)$ distribution in the case of Eaton and Kortum (2002) or the i.i.d. $Pareto(1, \theta^*)$ distribution in the case of Chaney (2008). Since the distribution F^* is, of course, a divergence $\delta = 0$ from itself, the Arkolakis et al. (2012) point estimate of GT_d obtained when assuming $F = F^*$ is the point at which our upper and lower bounds from Section 3.3 coincide, yielding

$$\overline{GT}_d(\delta = 0) = \underline{GT}_d(\delta = 0) = GT_d^{\text{ACR}} = 1 - \widetilde{\lambda}_{dd}^{\frac{1}{\theta^*}}. \tag{32}$$

We replicate standard estimation procedures in order to obtain an estimate of GT_d^{ACR} for France in 2018.[33] First, the OLS estimate of θ^* in equation (18) is 7.49 (with a standard error of 2.65). Second, in the 2018 data of our application, France has a value of $\widetilde{\lambda}_{dd} = 0.835$ (i.e. the country imports 16.5% of its expenditure). Putting these two numbers together, our estimate of France's gains from trade, according to the gravity model of either Eaton and Kortum (2002) or Chaney (2008), is $1 - (0.835)^{\frac{1}{7.49}} = 0.024$. In other words, according to the workhorse gravity model, aggregate welfare in France would drop by 2.4% if the country's trade policy were to retreat into autarky.

The remainder of our analysis is focused on exploring the sensitivity of this estimate to alternative distributions of global productivities.

[33] Following standard practice, we restrict the sample used to estimate this regression to international trade between countries (i.e., we exclude rows where $d = ROW$ or $o = d$). We also exclude one (extreme) outlier—Brazil's exports to South Korea, with effective tariff rate of 52%. These restrictions apply only for the purposes of identifying and estimating θ^*; we still require the economic model to match all D^2 trade flows and for the full set of equilibrium conditions to hold.

4.2 Distribution-agnostic bounds on the gains from trade

We turn now to calculating the range of values for GT_d that are possible if one is completely agnostic about the distribution of productivities in either the competitive or monopolistically competitive models. Propositions 1(a) and 1(b) have established that, in both cases, and under data conditions that hold for all countries in our application, the lower bound is $\underline{GT}_d = 0$ and the upper bound is $\overline{GT}_d = 1 - \widetilde{\lambda}_{dd}^{\frac{1}{\sigma-1}}$.

Like the Arkolakis et al. (2012) formula for $\overline{GT}_d(\delta = 0) = \underline{GT}_d(\delta = 0)$ in the gravity model case, the expression for the maximum possible value $\overline{GT}_d$ depends on the baseline value of λ_{dd}, discussed in Section 4.1 above. But rather than depending on the F^* shape parameter θ^*, the distribution-agnostic version $\overline{GT}_d$ depends on σ, the elasticity of substitution across products. Perhaps surprisingly, the value of σ does not enter the Arkolakis et al. (2012) formula at all, so the standard estimation template for GT_d in gravity models does not deliver (or require) an estimate of σ. But σ does matter for counterfactuals under essentially all other possible distributions $F \neq F^*$. Because aggregate trade data does not identify σ in the models we consider, we follow standard practice in macroeconomics (e.g. Hsieh and Klenow, 2009 and Einav et al., 2022) and set this parameter to $\sigma = 3$, while also exploring larger and smaller values below.[34]

Using this value of σ in the formula $\overline{GT}_d = 1 - \widetilde{\lambda}_{dd}^{\frac{1}{\sigma-1}}$ implies that the maximal value that France's gains from trade can take, across all possible production heterogeneity distributions F, is 8.6% (i.e. $1 - (0.835)^{\frac{1}{3-1}}$). This value is 3.6 times larger than that of 2.4% predicted by the workhorse model under F^*. Combined with the fact that the distribution-agnostic lower bound is $\underline{GT}_d = 0$, we conclude that a very wide range of predictions about GT_d are possible if one were completely agnostic about the underlying production heterogeneity in either the competitive or monopolistically competitive model.

4.3 Interpreting measures of divergence δ

One reaction to the wide bounds $\underline{GT}_d$ and $\overline{GT}_d$ reported above is that the stance of complete agnosticism about productivity distributions is too nihilistic. We therefore

[34] Disaggregated data on firms has been used to estimate σ in the monopolistic competition literature—for example, Eaton et al. (2011) obtain a value of $\sigma = 2.98$ when using data on French manufacturing firms' sales and productivity levels—but this strategy is not applicable in the case of competitive models where firm boundaries are not determined.

re-evaluate these bounds for a smaller set of distributions by narrowing the neighborhood size $\mathcal{N}_\delta$ such that δ is finite (but still greater than zero so that we consider departures from F^*). Before reporting on this exercise in Section 4.4 below we first pause to provide guidance on how a divergence of any given value, such as δ, can be interpreted.

To provide one benchmark, consider the sensitivity analysis conducted for the case of gravity models when the value of the shape parameter θ^* is adjusted, as is common practice among users of quantitative trade models. Begin with the Eaton and Kortum (2002) case, with the global productivity distribution F^* based on $Frechet(1, \theta^*)$. Now consider changing this distribution to an alternative, denoted $F^{\text{alt.}}(\theta)$, that remains an Eaton and Kortum (2002) gravity model but with a different shape parameter, θ. Because this distribution $F^{\text{alt.}}(\theta)$ differs from F^* we can calculate the resulting divergence, $D(F^{\text{alt.}}(\theta)||F^*) \geqslant 0$.

The red line in Figure 2 plots the value of $D(F^{\text{alt.}}(\theta)||F^*)$ against θ for this Frechet case estimated in Section 4.1.[35] That is, F^* corresponds to $D = 20$ countries taking independent draws from $Frechet(1, \theta^* = 7.49)$. Similarly, the blue line reports an analogous exercise for the case of Chaney (2008), where F^* is based on a $Pareto(1, \theta^* = 7.49)$ distribution and $F^{\text{alt.}}(\theta)$ is based on $Pareto(1, \theta)$. Naturally, in both cases the divergence is zero at $\theta = \theta^* = 7.49$, and the divergence is strictly greater than zero for all other values of $\theta \neq \theta^*$.

Our interest lies in the extent to which the divergence grows as θ differs increasingly from θ^*. For example, we see that divergence $D(F^{\text{alt.}}(\theta)||F^*)$ below the value of 1 corresponds to values of θ that span the range $\theta = [6.0, 9.7]$ in the Frechet case and $\theta = [5.5, 10.5]$ in the Pareto case. These windows seem well within the range of conventional sensitivity analysis that is already a key part of quantitative trade modeling.[36] We therefore view a neighborhood such as $\mathcal{N}_\delta$ at $\delta = 1$—that is, the set of all distributions F that lie within $D(F||F^*) \leqslant 1$ of F^*—to be of substantial interest when conducting sensitivity analysis.

[35] The calculations in Figure 2 use the KL divergence rather than the Hybrid divergence used in all further analysis below. This is conservative (because KL divergence is always weakly smaller than Hybrid divergence) and avoids computation of D-dimensional numerical integrals at each value of θ (since KL divergence is available in closed form for the two cases calculated here).

[36] For example, Eaton and Kortum (2002) deploy three procedures for estimating θ^* and arrive at values ranging from 3.60 to 12.86.

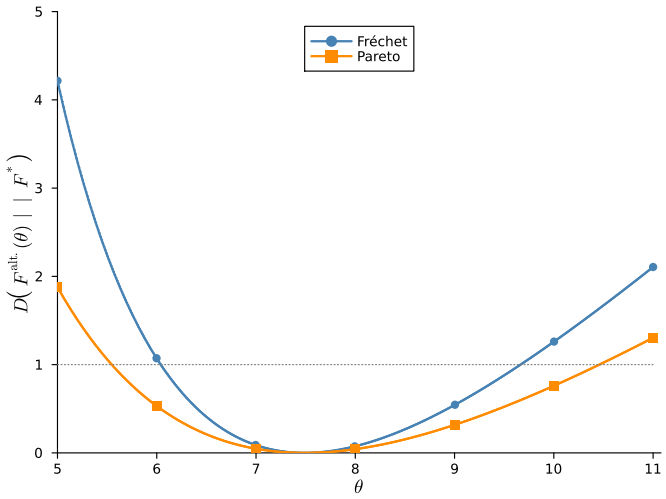


Figure 2: Examples of divergence between distributions

4.4 Gains from trade for distributions "close" to the gravity model

We now estimate bounds on counterfactual predictions about the gains from trade that hold for a varying set of neighborhoods $\mathcal{N}_\delta$ that are defined by their productivity distribution F lying within a divergence of δ from the gravity model distribution F^*.

As described in Section 3.3, for any given value of δ, we calculate the upper and lower bounds on country d's gains from trade, $\overline{GT}_d(\delta)$ and $\underline{GT}_d(\delta)$ respectively, that would be possible across all productivity distributions with divergence less than δ from the gravity model distribution F^*. And we do so while always ensuring that, given any distribution under consideration, it is fed into an otherwise standard equilibrium model that exactly matches all trade flow data points at the factual equilibrium as well as the standard orthogonality conditions of gravity model estimation. That is, we solve program (31) at the chosen value of δ, with $\overline{GT}_d$ drawing on program (23) in the competitive case and program (28) in the monopolistically competitive case; analogously, we obtain the lower bounds $\underline{GT}_d(\delta)$ by using a version of these same programs that minmizes GT_d rather than maximizing it. Throughout, we focus on the case where $d =$ France.

Figure 3 plots our baseline results. For purposes of comparison, the dashed orange line displays the value of GT_d^{ACR}, the extent of gains from trade in the workhorse model (i.e. with $\delta = 0$); as discussed in Section 4.1, we estimate this to be 0.024. In addition, the dashed black lines report the distribution-agnostic upper and

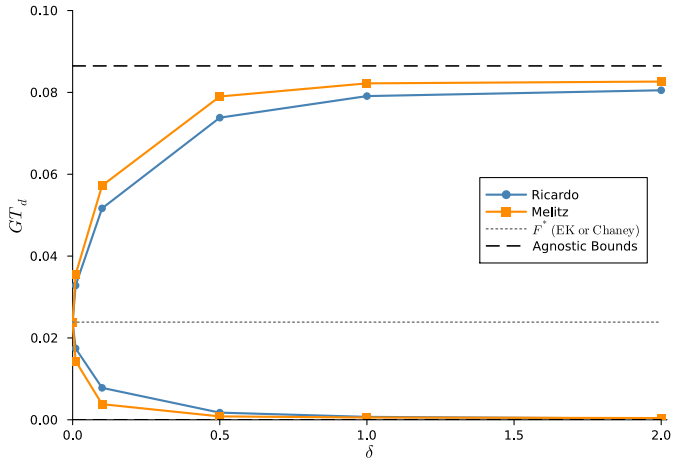


Figure 3: Bounds on gains from trade (for France) across alternative productivity distributions

lower bounds, $\overline{GT}_d$ and $\underline{GT}_d$ respectively, that are possible across all distributions; as discussed in Section 4.2, we estimate these bounds to be $\overline{GT}_d = 0.086$ and $\underline{GT}_d = 0$. Finally, we add to Figure 3 the values of $\overline{GT}_d(\delta)$ and $\underline{GT}_d(\delta)$ as a function of δ for the case of the Ricardian model (in solid blue) and the Melitz (2003) model (in solid red), each computed for five values of δ at each bound. As expected, the finite-δ bounds begin at the standard gravity model estimate of $GT_d = 0.024$ when δ is small and then approach the distribution-agnostic bounds $\overline{GT}_d$ and $\underline{GT}_d$ as δ grows larger.

An object of substantial interest is the *rate* at which the bounds diverge as δ is increased, starting at zero. This encapsulates a notion of sensitivity to alternative productivity distributions, and Figure 3 reveals a striking amount of sensitivity. For example, even at $\delta = 1$—a value of divergence that, as discussed in Section 4.3, we view as corresponding to a conventional exploration of robustness in gravity models—the trade data support Ricardian models with gains from trade that are over three times larger than in the gravity model, as well as Melitz (2003) models that are also over three times larger. Moreover, these bounds are very close (85% of the way for Ricardo and 92% for Melitz (2003)) to the distribution-agnostic maximal upper bound.

Returning to the discussion in Section 3.2, Costinot and Rodríguez-Clare (2018) have stressed that GT_d in the competitive model is pinned down by the area under country d's import demand curve. This observation begs the question: when GT

hits the maximal or minimal values at, say, $\delta = 1$ that are shown in Figure 3, what do the corresponding maximal and minimal import demand curves look like? Figure 4 provides the answer. The curves that achieve the distribution-agnostic ($\delta \to \infty$) upper and lower bounds are illustrated with dashed blue and orange lines, respectively. To this we now add the corresponding solid lines for the $\delta = 1$ case, which naturally lie inside the $\delta \to \infty$ versions.[37] The upper bound curve for $\delta = 1$ tracks that for $\delta \to \infty$ for small increases in the foreign price, κ, before diverging at larger values—which means that the distinction between $\delta = 1$ and $\delta \to \infty$ would not arise for counterfactuals involving modest tariff increases, but it does begin to arise as the import price increase κ gets closer to autarky.[38] However, in line with Figure Figure 3 the difference between the solid and dashed versions of France's import demand curves is relatively minor, so they give rise to similar values for bounds on the gains from trade (i.e. possible areas under the curves).

4.5 Taking stock

Put together, the findings of this section can be summarized as follows. First, as a benchmark, the gravity model, with its productivity distribution F set to the workhorse version of F^*, and when estimated in the standard manner, predicts that French living standards would fall by 2.4% if that country were to cease trading internationally. Second, our new analytical bound formula implies that if one were agnostic about the productivity distribution F to use but followed exactly the same estimation steps as are applied to the gravity model then the cost to France of moving to autarky would lie between 0% and 8.6% of baseline real income. Finally, we find a similar range of values when we restrict the set of possible distributions

[37] To do this, we proceed as follows. An output of the Christensen and Connault (2023) procedure for solving program (31) at any finite δ is a numerical representation of the most favorable distribution F_δ^+ and accompanying parameter vector ψ^+ that solve that program. Armed with these outputs any expectation, under F_δ^+, of any known function of ψ^+ can be solved for numerically. We then plot the solid blue line in Figure 4 by computing the value of import demand $\lambda_{Md}(\kappa) = 1 - \lambda_{dd}(\kappa)$, written as an expectation under F_δ^+ and at the parameters ψ^+, for each point along a grid of values for κ. The solid orange line applies an analogous procedure using the least favorable distribution and parameter vector that solve the minimization version of (31).

[38] This difference reflects the fact that large increases in import prices require the probabilities of low productivity values to be adjusted, and such adjustments induce a large divergence because F^* has relatively little mass at low values. As for the lower bound in orange, this follows the opposite intuition since it requires large movements in the productivity distribution among low productivities even for small import price increases.

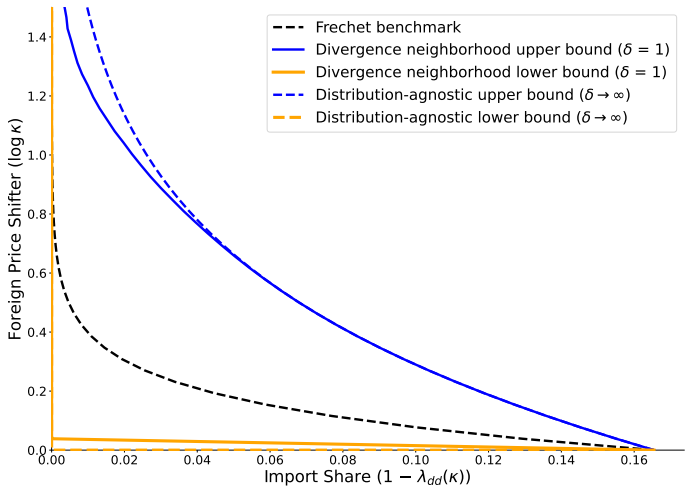


Figure 4: Import demand curve (for France) in the competitive model under the extremizing distributions within divergence neighborhoods

F to only those that lie "close" to F^* (for example, via the restriction that the divergence of F from F^* must be no more than one). All of these conclusions apply to both the competitive version of the gravity model and the monopolistically competitive version.

One implication of these findings is that even the plain-vanilla trade model we study—one without features such as multiple sectors or input-output linkages that tend to increase gains from trade—is entirely consistent with the possibility that gains from trade are many times larger than they are when the gravity assumption of F^* distributions is imposed. But versions of the model with no gains from trade are possible too.

A second implication of our findings is that quantitative trade models appear sensitive to relatively small changes in their underlying productivity distributions. Starting from the workhorse distributional assumption of F^* we have seen how there exist distributions F that are "close" to F^* and yet deliver gains from trade that are much larger or smaller than is possible in the workhorse version. Indeed, almost all of what is possible in general, with complete agnosticism about distributions, is possible even for the set of distributions F that are "close" to F^*. One might conjecture that matching trade data and standard orthogonality conditions would ensure robustness of quantitative trade models to their underlying assumptions about productivity distributions. According to our analysis this is not the case.

5 High-level Restrictions on Productivity Distributions

Our main finding so far is that gains from trade can be many times larger or smaller than typically found in gravity models if one is agnostic about the model's underlying productivity distribution F, and this remains true even when the set of admissible distributions is restricted to lie within a small-divergence neighborhood of the distribution used in the gravity model F^*. In this section we evaluate how much our gains from trade bounds would shrink when the set of distributions under consideration is restricted—both from the starting point of agnosticism and from a given neighborhood around F^*—by requiring that distributions satisfy a range of high-level properties.

In particular, the high-level restrictions that we impose are motivated by three core features of the global productivity distributions F^* used by Eaton and Kortum (2002) in the competitive model and Chaney (2008) in the monopolistically competitive model. The first such feature is that each country takes independent draws—that is, there is no cross-country dependence or correlation across their draws. A second feature is that every country draws from the same marginal distribution. Finally, the particular marginal distribution that all countries draw from is $Frechet(1, \theta^*)$ in the Eaton and Kortum (2002) case and $Pareto(1, \theta^*)$ in the Chaney (2008) case. In this section we impose versions of these high-level characteristics on the distributions F that are admissible when we evaluate bounds on the gains from trade. We work through these restrictions one by one and, where possible, in combination.

5.1 Restrictions: perfect vs. monopolistic competition

One starting observation concerns a broad distinction between the competitive and monopolistically competitive models. The competitive model involves head-to-head competition, in every destination market d, over the worldwide producers of each good ω who are based in each of the D origin countries. When country d moves from its factual trading equilibrium to the counterfactual state of autarky, the supplier of every imported good will shift from the foreign supplier (who had won this head-to-head competition in the factual equilibrium) to a domestic supplier. The gains from trade therefore depend on the size of the foreign supplier's victory. The stochastic properties of this victory depend on country d's marginal

productivity distribution and the correlation of d's draws with the best of the $D-1$ foreign draws, a distribution that itself depends on those $D-1$ foreign countries' marginal distributions and their dependence. In short, the gains in the Ricardian model clearly depend on properties of domestic marginals, foreign marginals, and the dependence among them. We explore these components below.

On the other hand, consider the monopolistically competitive model. In this case, competition is not head-to-head for any good; instead, every firm pays an entry cost in order to be the global monopolist in its own product. When country d moves to autarky it will lose access to all foreign imports and seek to replace them with domestic alternatives. But the lack of head-to-head competition means that what matters for the gains is how well these newly produced domestic varieties can replace the aggregate contribution of the now-unavailable foreign imports, which is purely a function of the domestic distribution. That is, conditional on knowing the marginal distribution of productivities in country d the additional properties of F—such as the $D-1$ foreign countries' marginal distributions or the cross-country dependence of productivity draws—has no direct bearing on GT_d. Given these points, there is little to be learned by imposing high-level restrictions on the Melitz model. All results in the following subsections therefore focus on the Ricardian case alone.

5.2 Imposing independence

The F^* used by the Eaton and Kortum (2002) version of the Ricardian model assumes that, for any good ω, countries o take their productivity draws $z_o(\omega)$ independently from one another. By contrast, the distributions F that we have considered so far have allowed for arbitrary forms of cross-country dependence. We now seek to narrow this gap between F^* and F by restricting the set of distributions F considered such that any admissible F must satisfy cross-country independence just like F^* does. In this case the only distinction between F^* and F that remains is that due to the particular marginal distributions that each country draws from.

Distribution-agnostic bounds. We begin by extending the distribution-agnostic upper-bound argument of Proposition 1(a) to the case where distributions F must obey independence. That is, we characterize the bound $\overline{GT}_d^{\text{indep.}}$ that obtains if

program (23) is solved subject to the constraint that the choice variable F must feature cross-country independence. Indeed, we show that, provided a stronger data condition holds, the bounds with and without independence imposed are the same: $\overline{GT}_d^{\text{indep.}} = \overline{GT}_d$. This is perhaps surprising given that the construction used in the Proposition 1(a) case exploited a strong form of dependence across countries.

Proposition 2 (Bounds for competitive model with independence). *Index the country of interest with j. Let ν_j denote the share of residualized tariff variation in rows where j is the exporter, and let β_j denote the negative of the slope of a regression of $\ln(\lambda_{jd}/(1-\lambda_{jd}))$ on residualized tariffs in rows where j is the exporter. Suppose the following data condition holds:*

$$\nu_j \beta_j < \sigma - 1. \tag{33}$$

Then, there exists a distribution F under which countries draw their productivities independently that attains the agnostic upper bound in (24), i.e.,

$$\overline{GT}_d^{\text{indep.}} = \overline{GT}_d = 1 - \lambda_{dd}^{\frac{1}{\sigma-1}}. \tag{34}$$

Furthermore, there exists such a distribution F that has Frechet marginals.

Like those seen earlier, this proposition presumes a data condition that can be easily checked on any dataset and is verified for all countries in our application. The proof of Proposition 2 provides a construction of a particular F that involves countries drawing independently from a distribution with (non-identical) Frechet marginals. The marginal for country j has a shape parameter θ that approaches $\sigma - 1$ from above, in which case that country's production is dominated by a small mass of highly productive goods, such that in autarky the contribution of other goods is negligible and hence the denominator in (26) becomes arbitrarily large. The marginals for all countries other than j are all common and set to be Frechet with shape parameter increasing from θ^* until the data-matching and orthogonality constraints in program (23) are satisfied. As we show, this is always possible as long as the data condition in the premise of Proposition 2 holds.[39]

[39] Intuitively, on its own, letting j's shape parameter approach $\sigma - 1$ would predict j's trade to be less elastic than the data suggests and therefore rationalizing the data would require setting unobserved trade shifters, $\{A\}$, to values that violate the orthogonality restriction. The construction compensates for this by making trade more elastic among the other countries to maintain overall orthogonality. The data condition ensures that it is possible to do this, by requiring that j's exports are either not too elastic or occupy a small enough share of overall tariff variation.

However, as this construction makes clear, without access to dependence, the data-matching and orthogonality constraints mean that the gains from trade for one country are being maximized in ways that will tend to reduce other countries' gains. This contrasts with the case of Propositions 1(a) and 1(b) where all countries could simultaneously achieve the bound $\overline{GT}_d$.

Bounds for F in a finite neighborhood of F^*. Proposition 2 suggests that when one considers all distributions it is essentially without loss to impose independence when studying country d's gains from trade. However, as the construction provided suggests, this may be a limiting result that may not be very accurate in any neighborhood that is "close" to F^*. We therefore turn to a quantitative exploration of the case where F must both satisfy cross-country independence and lie within the neighborhood $\mathcal{N}_\delta$.

Recall that two jointly distributed random variables are independent if and only if their joint CDF is equal to the product of their marginal CDFs. While the random variable in our application has $D = 20$ dimensions rather than 2, rendering it infeasible to imply this restriction directly to the full joint CDF, the min structure of the Ricardian model suggests a natural way to collapse the dimensionality. In particular, write $p_{Fd}(\omega) \equiv \min_{o \neq d} p_{od}(\omega)$ as the best foreign price for good ω. As seen in the discussion of the distribution-agnostic bounds, the relationship between the domestic price, $p_{dd}(\omega)$, and the best foreign price, $p_{Fd}(\omega)$ is particularly important for gains from trade. We therefore approximate this restriction about CDFs by using a finite number of expectations at L^I specific quantiles of the relevant CDFs. That is, we say that F satisfies approximate best foreign independence, shortened to approximate independence going forward, when

$$\mathbb{E}_F[\mathbb{1}\{z_d(\omega) < \bar{z}_l\}\mathbb{1}\{p_{Fd}(\omega) < \bar{p}_{Fd,m}\}] = \mathbb{E}_F[\mathbb{1}\{z_d(\omega) < \bar{z}_l\}]\,\mathbb{E}_F[\mathbb{1}\{p_{Fd}(\omega) < \bar{p}_{Fd,m}\}]\,, \tag{35}$$

holds for all destinations d and all L^I quantiles l, m, and where $\bar{x}_l$ denotes the value of the lth quantile of a distribution of interest.[40] Since this constraint imposes the necessary and sufficient condition for independence at a finite number of points in the distribution, by letting L^I grow we can approximate increasingly well the case

[40] We use the marginals of F^*, i.e. Frechet$(1, \theta^*)$, for this purpose. This has the advantage of using quantile values that are spaced out in the same proportion to their distribution under F^*.

of true independence.

We impose this form of approximate independence by setting $L^I = 10$ and adding the $D(L^I - 1)^2 = 1620$ constraints embodied in (35) to the program (31) when computing the upper bound $\overline{GT}_d(\delta)$ and the lower bound $\underline{GT}_d(\delta)$.[41] This continues to search over all distributions $F \in \mathcal{N}_\delta$ that deliver Ricardian models that are consistent with trade data and the standard orthogonality condition, but with the space of such distributions now confined to those that feature our form of approximate independence.

Figure 5 displays the results from implementing this independence restriction: blue lines show the unrestricted bounds from Figure 3 and the red lines show those that impose approximate independence. Both the lower and upper bounds are somewhat tighter than the unrestricted case. For example, at the $\delta = 1$ level, $\overline{GT}_d(\delta)$ is 7.3% when F must satisfy our approximate independence restriction and 7.9% when F is unrestricted. This implies that the maximum possible gains from trade in France would be about 11% (i.e. $1 - (0.073 - 0.024)/(0.079 - 0.024)$) less sensitive to the underlying productivity distribution F if it were believed that countries draw their technologies in a way that satisfies our independence restriction. As anticipated, even though Propositions 1(a) and 2 imply that imposing an independence restriction has no "bite" at $\delta \to \infty$, the bite is non-negligible at $\delta = 1$.

These results are also informative for the prior literature that has relaxed the independence assumption in Eaton and Kortum (2002). For example, a special case of the analysis in Lind and Ramondo (2023) maintains the use of Frechet marginals but allows for non-zero amounts of correlation. Figure 5 illustrates an analogous idea but in the opposite direction: start with all distributions within a given δ neighborhood (blue lines) and then impose restrictions (red lines) so as to keep only those that satisfy our independence restriction. We learn that this can reduce the maximal gains from trade, but not by much (only by 11% at $\delta = 1$). The effect on the minimal gains from trade is similar.

[41] As the RHS of (35) is a nonlinear function of expectations, we include an additional $2D(L^I - 1) = 160$ restrictions of the form $\mathbb{E}_F[1\{z_o(\omega) < \bar{z}_l\}] = \nu^z_{o,l}$ and $\mathbb{E}_F[1\{p_{Fd}(\omega) < \bar{p}_{Fd,m}\}] = \nu^p_{d,m}$ with parameters $\nu^z_{d,l}$ and $\nu^p_{d,m}$ appended to the vector ψ. The RHS of (35) therefore becomes $\nu^z_{d,l}\nu^p_{d,m}$.

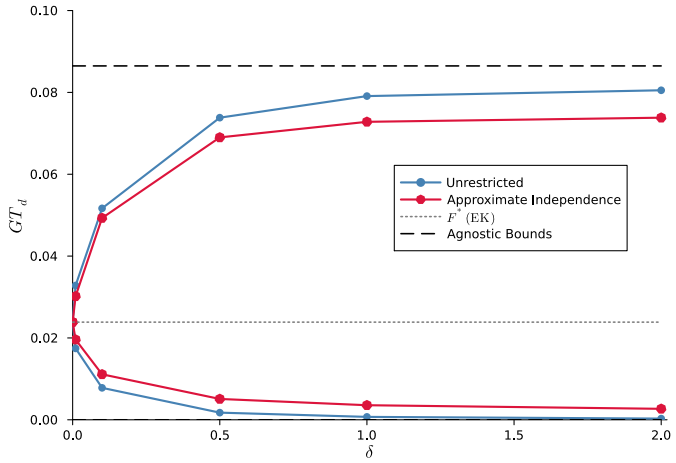


Figure 5: Ricardian gains from trade and cross-country dependence

5.3 Imposing common marginals

The second high-level restriction imposed in the Eaton and Kortum (2002) choice of F^* is that all countries o draw their productivities $z_o(\omega)$ from a common marginal distribution. Of course, the actual distribution F^* involves not only common marginals but specific marginals given by $Frechet(1, \theta^*)$. But the feature of commonality may derive from primitive assumptions on where technologies derive from, or simply from a desire for parsimony. So it is interesting to evaluate how the case of common (but flexible) marginals alone can shape gains from trade in data-consistent Ricardian models.

Distribution-agnostic bounds. We begin again with an extension of Proposition 1(a) to the case where admissible distributions F must feature common marginals. In fact, the following proposition goes further and shows that, even with the additional restriction that countries' common marginals must all be $Frechet(1, \theta)$, the resulting upper bound on the gains from trade $\overline{GT}_d^{\text{CM}}$ will be the same as that of $\overline{GT}_d$ in Proposition 1(a) as long as a stronger data condition holds. It is important to note, however, that the value of θ in this common Frechet marginal will differ, in general, from that of θ^* used in F^*.

Proposition 3 (Bounds for competitive model with common marginal distributions). *Index the country of interest with j. Suppose that there exists a partition of all foreign countries into K mutually-exclusive blocks, with trade flows assigned to blocks based*

on origin, such that

$$\sum_{k:\beta_k>0} \nu_k \beta_k > \theta^* - (\sigma - 1), \tag{36}$$

where ν_k is the share of residualized tariff variation in block k and β_k is the negative of the slope of a regression of trade shares against residualized tariffs in block k.

Then, there exists a distribution F under which countries draw their productivities from a common marginal distribution, with potential dependence across countries, that attains the agnostic upper bound in (24)*, i.e.,*

$$\overline{GT}_d^{\text{CM}} = \overline{GT}_d = 1 - \lambda_{dd}^{\frac{1}{\sigma-1}}. \tag{37}$$

Further, there exists such a distribution with Frechet marginals.

The data condition behind this proposition is more involved than those seen earlier. This is unsurprising, given that stronger restrictions have been placed on the possible distributions that must satisfy all constraints despite achieving the same bound. While we relegate the exact condition to the Appendix, the rough idea is to partition all countries (other than d) into K groups in order to ensure that the sum across groups of certain within-group covariances (between tariffs and certain trade data summaries) is smaller than the overall variance of the tariff data. We find that this condition is satisfied for our application even when only $K = 2$ groups are considered.

The constructive proof of Proposition 3 uses a sequence of common Frechet distributions with shape parameter θ that approaches $\sigma - 1$ from above, but does so while using the dependence structure to ensure that the orthogonality condition continues to hold. In the case of independent distributions (as in F^*) it is impossible for all countries to have a common Frechet marginal distribution with shape parameter different from θ^*, since this would violate the orthogonality condition (19). In particular, $\theta < \theta^*$ would imply that trade should be substantially less elastic than the data suggests. However, creating positive correlation within blocks of foreign countries reduces the role of idiosyncratic productivity dispersion relative to fixed trade shifters, such as tariffs, making within-block trade highly elastic to such shifters and thus rationalizing an overall elasticity of θ^*. The data condition states that it is possible to construct blocks of foreign countries with sufficient negative covariance to offset the reduction in θ from θ^*.

Bounds in the neighborhood of F^*. Turning to the case of finite δ, like in the previous subsection we approximate a restriction about (marginal) CDFs by using a finite number of expectations at L^{CM} quantiles l. All countries' marginal distributions are approximately common if the following condition holds

$$\mathbb{E}_F[\mathbb{1}\{z_o(\omega) < \bar{z}_l\}] = \mathbb{E}_F[\mathbb{1}\{z_{o'}(\omega) < \bar{z}_l\}] \text{ for all } l \leqslant L^{CM} \text{ and } o' \neq o. \tag{38}$$

That is, common marginals across countries requires that the expected mass of the distribution below any given point $\bar{z}_l$ for one country o is the same as that for any other country o'.[42] The gains from trade bounds, such as $\overline{GT}_d(\delta)$, among all distributions F for which all countries share the same marginal distribution can then be approximated by adding the $(D-1)L^{CM} = 1900$ constraints embodied in (38) to the program (31).

The green lines in Figure 6 report the results from this exercise. (Again, the blue lines display the unrestricted baseline case for comparison.) The message from this figure is that imposing the common marginals constraint does not do much to reduce the sensitivity of quantitative Ricardian models—not just at large δ as in Proposition 3, but even at small values of δ such as 1 or 2. In particular, the lower bound of zero is essentially unaffected by this constraint and the upper bound is reduced, but only by about 5% (i.e. $1 - (0.076 - 0.024)/(0.079 - 0.024)$) at a value such as $\delta = 1$.

5.4 Imposing common $Frechet(1, \theta^*)$ marginals

The previous restriction on F requires that all countries share the same marginal distribution, but it remains agnostic about what that marginal distribution is. A tighter restriction of interest is that the common marginal distribution takes a particular form, such as the $Frechet(1, \theta^*)$ distribution, as in the F^* used by Eaton

[42] We again choose the points $\bar{z}_l$ as the L^{CM} quantiles of the $Frechet(1, \theta^*)$ distribution. For any given L^{CM}, the common marginals restriction in (38) requires far fewer restrictions than the pairwise restrictions required for independence in (35), so it is feasible to use a larger number of quantiles. We therefore set $L^{CM} = 50$. As a robustness check, we take the parameter vector that maximizes $GT_d(\delta)$ for $\delta = 1$ with $L^{CM} = 50$ and evaluate the (minimum) divergence at this parameter vector with $L^{CM} = 100$. This gives a divergence of 1.03, suggesting that increasing L^{CM} is unlikely to meaningfully affect the results.

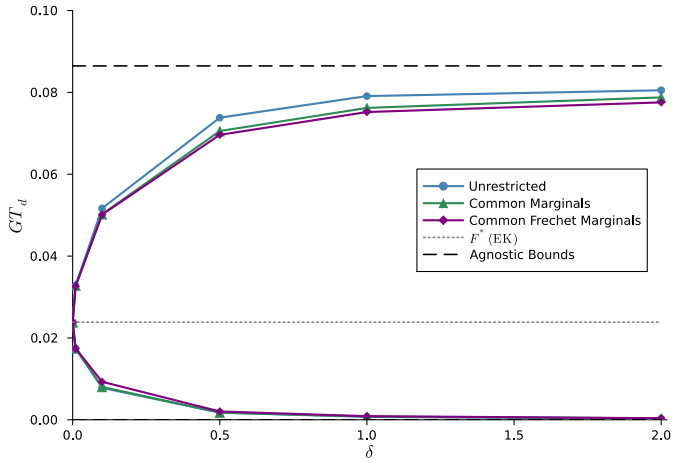


Figure 6: Ricardian gains from trade and countries' marginal distributions

and Kortum (2002).[43] We do this next. It is worth noting that in this case the only difference between F and F^* is that F remains agnostic about the cross-country dependence structure, whereas F^* assumes independence.

To see how this can be done, let $H(\cdot)$ denote the CDF of the inverse of the distribution $Frechet(1, \theta^*)$. Divide the unit-interval support of $H(\cdot)$ into $\overline{H}$ discrete and evenly-spaced points indexed by h. Finally, let $\overline{H}_h$ denote the value of the hth quantile of $H(\cdot)$. Then we can ensure that the marginal of the distribution of productivity draws $z_o(\omega)$ in any country o has approximately the same marginal distribution as $Frechet(1, \theta^*)$, as desired, by matching each of these $\overline{H}$ quantile values, as long as $\overline{H}$ is large enough for this approximation to be accurate.[44] Formally, this amounts to imposing the constraint

$$\mathbb{E}_F[\mathbb{1}\{z_o(\omega) < \overline{H}_h\}] = H(\overline{H}_h) \quad \text{for all } o \text{ and all } h \leqslant \overline{H}. \tag{39}$$

Adding the $D\overline{H} = 2000$ constraints in (39) to the program (31) ensures that bounds on GT_d at any δ only consider distributions F that deviate from F^* in terms of the cross-country dependence of their technology draws and not at all (at least approximately) in terms of the Frechet marginals that all countries share.

We report on this case in Figure 6. The purple lines impose the new restriction

[43] This is a distinct exercise from that in Proposition 2. Here we constrain the common marginal to be $Frechet(1, \theta^*)$, with θ^* set to the estimated value of 7.49; the construction in Proposition 3 instead uses common $Frechet(1, \theta)$ marginals with a flexible θ.

[44] As with the common marginals restriction, we set $\overline{H} = 50$.

of common $Frechet(1, \theta^*)$ marginals. Again, counterfactual sensitivity remains substantial, despite requiring that all countries have the same marginal distributions as they do in Eaton and Kortum (2002). Perhaps surprisingly, the purple lines (where all marginals are the same and Frechet) are very close to the green lines (where all marginals are the same but can be *any* marginal distribution). So, at least when dependence is left flexible, the sensitivity of the Ricardian model's predictions has little to do with the particular marginal distributions employed.

5.5 Imposing both independence and common marginals

The restrictions we have introduced so far have all been imposed one at a time, but it is of interest to also impose subsets of them simultaneously. One such combination is the case where F is *i.i.d*—where countries draw productivities independently and from an identical marginal distribution. Indeed, when F is i.i.d, just as F^* is, the only remaining difference between F and F^* is that the marginals of F are flexible whereas those of F^* are $Frechet(1, \theta^*)$.

Distribution-agnostic bounds. The following proposition provides a sufficient condition under which the maximal gains from trade $\overline{GT}_d$ in Proposition 1(a) can be achieved even for the i.i.d. case.

Proposition 4 (Bounds for competitive model with i.i.d. distribution). *Suppose the baseline tariff and trade data satisfy Condition 1 in Appendix A.4. Consider all possible productivity distributions F under which countries draw their productivities independently from a common marginal distribution. Then the largest possible value for country d's gains from trade in the competitive model, as defined in program (23), is given by*

$$\overline{GT}_d^{\,\text{IID}} = \overline{GT}_d = 1 - \lambda_{dd}^{\frac{1}{\sigma-1}}. \tag{40}$$

The basic idea behind this proposition is to construct a distribution F whose identical marginals approach a distribution comprised of a mixture of two Dirac distributions at the points of zero and one, with an interior weight on each Dirac distribution that is chosen so as to match all countries' data points and satisfy the orthogonality condition. Data Condition 1 in Appendix A.4—which is verified for all countries in our sample—is a sufficient condition that ensures this is possible.

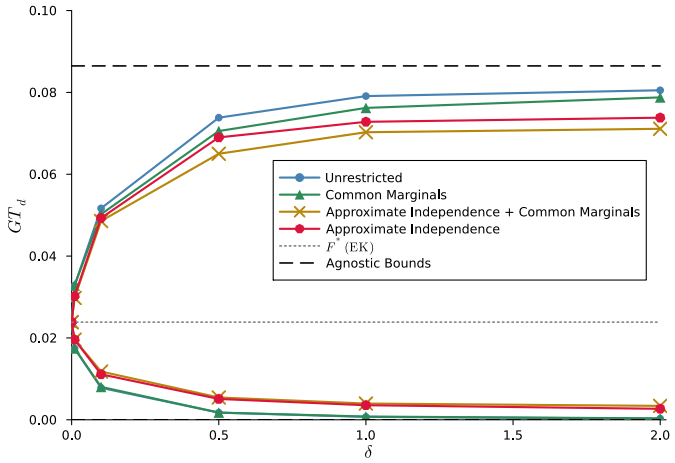


Figure 7: Ricardian gains from trade with independence and common marginals

Bounds in the neighborhood of F^*. As before, even though Proposition 4 suggests that the imposition of the i.i.d. restriction will not shrink our estimated bounds across an infinitely large neighborhood of distributions, the construction offered in the proof of that proposition suggests that the i.i.d. restriction may have substantial bite in finite neighborhoods of F^*. We therefore estimate the bounds for such a case in a manner analogous to that in Sections 5.2 and 5.3. That is, we add to the program (31) the stacked set of $D(L^I - 1)^2 = 1620$ constraints in (35) that ensure approximate independence alongside the $(D-1)L^{CM} = 1900$ constraints in (38) that ensure approximate common marginals.

The results of this exercise are shown in Figure 7. In particular, the tan lines report on the case where all permitted distributions F satisfy both our independence restriction and common marginals. For comparison, the blue lines again display the unrestricted case and the green lines display that with common marginals (but not independence). Once again, the bounds conveyed by the tan lines are still quite wide. Going from the unrestricted case to the doubly-restricted case of both independence and common marginals achieves, at $\delta = 1$ for example, a reduction in sensitivity of the upper bound by about 16% (i.e. $1 - (0.070 - 0.024)/(0.079 - 0.024)$). What is most surprising is that the combination of independence and common marginals (tan in this figure) achieves little more robustness in $\overline{GT}_d(\delta)$ than that of either independence (Figure 5) or common marginals (orange line in this figure) on their own.

This exercise is of particular interest for connecting with the empirical estimation

of Ricardian models that are more flexible than the gravity version of Eaton and Kortum (2002). In particular, Adao et al. (2017) demonstrate that counterfactual predictions, such as those concerning the gains from trade, in nonparametric Ricardian models can be constructed from the knowledge of countries' import demand systems alone, and that such systems can be estimated using data on bilateral trade flows, tariffs, and exogeneity restrictions that generalize (19). With a single cross-section of data, as we have used here, import demand systems can only be estimated under the assumption that they are identical across countries. In the context of the competitive model in Section 2 this will obtain when F features exchangeability across countries, a condition that is implied by the i.i.d. condition. The similarity of the light and dark blue lines in Figure 7 therefore gives a sense of how little is restricted when researchers assume that countries' import demand systems are identical.

5.6 Summary

Stepping back, the combined lesson from the restrictions imposed throughout this section is that gains from trade in the Ricardian model can be many times larger or smaller than they are in the workhorse Eaton and Kortum (2002) version, even in the face of reasonable restrictions. Interestingly, our estimated bounds are wide when one departs from that model's productivity distribution—namely, that countries' productivities are drawn i.i.d. from $Frechet(1, \theta^*)$ marginals—in *either* of two high-level directions that are available. One such direction is to relax the independence assumption while maintaining that of common $Frechet(1, \theta^*)$ marginals, as we have explored in Figure 6. The other direction is to relax the assumption of common $Frechet(1, \theta^*)$ marginals while maintaining the assumptions of independence and common marginals, as we have seen in Figure 7. Either way, for the finite neighborhood within divergence $\delta = 1$ of the Eaton and Kortum (2002) distribution, we estimate bounds on gains from trade that are about 16% narrower than those for the completely unrestricted case but still contain Ricardian models with substantial (or tiny) gains from trade.

6 Concluding Remarks

Three modern revolutions have taken place in the study of international trade. A first is the desire to answer counterfactual questions using quantitative simulations. A second is the incorporation of micro-level production heterogeneity across goods and countries, either in Ricardian models a la Eaton and Kortum (2002) or monopolistically competitive models a la Melitz (2003). And a third is the desire to ensure that models exactly match all baseline trade flows, and their implied empirical patterns such as the fit of the gravity equation or the magnitude of correlation between tariffs and trade flows, before being used for counterfactual simulations away from the baseline scenario.

The goal of this paper has been to explore the robustness of this three-part revolution to variations in the assumed distribution of cross-country and cross-good heterogeneity in production technologies. Workhorse models such as Eaton and Kortum (2002) and Chaney (2008) assume that this heterogeneity is drawn from particular distributions that have been chosen largely for their tractability and alignment with gravity model traditions. But given the impossibility of direct evidence about these distributions—since what matters for key counterfactual questions such as the gains from trade is the distribution of a country's productivity among goods it does not produce—it seems important to probe the sensitivity of model predictions to variations in the underlying distribution of assumed production heterogeneity.

Our results convey a strong sense in which quantitative trade models are highly sensitive to the underlying distribution of production heterogeneity that is assumed. And this is true despite ensuring that every distribution of heterogeneity that we consider is capable of generating, in full general equilibrium, the same values for all bilateral trade flows as occurs at the empirical baseline, as well as satisfying the same orthogonality conditions that are typically used for parameter estimation. Indeed, this finding continues to hold even when one maintains components of the workhorse distributions—e.g. that countries' draws are independent but with arbitrary country-specific marginals, or independent with arbitrary but common marginals, or with benchmark common marginals but arbitrary amounts of (positive or negative) dependence.

A natural direction for future work is to extend the methods in this paper to a wider set of model features. In principle, our model can be augmented to

include elements such as domestic input-output linkages, international supply chains, multiple sectors, and multiple factors—all of which typically increase the gains from trade in quantitative models, implying that their omission here has been conservative. Likewise, phenomena such as factor mobility and consumer demand are typically modeled via distributional assumptions about the heterogeneity of households' preferences. The tools developed in this paper could be applied to such settings as well.

References

Adao, Rodrigo, Arnaud Costinot, and Dave Donaldson, "Nonparametric Counterfactual Predictions in Neoclassical Models of International Trade," *American Economic Review*, 2017, *107* (3).

_ , Costas Arkolakis, and Sharat Ganapati, "From Heterogeneous Firms to Heterogeneous Trade Elasticities: The Aggregate Implications of Firm Export Decisions," *Journal of Political Economy*, 2026, *forthcoming*.

Allen, Treb and Costas Arkolakis, "Quantitative regional economics," in Dave Donaldson and Stephen J. Redding, eds., *Handbook of Regional and Urban Economics*, Vol. 6, Elsevier, 2025, pp. 1–72.

Arkolakis, Costas, Arnaud Costinot, and Andres Rodríguez-Clare, "New Trade Models, Same Old Gains?," *American Economic Review*, 2012, *102* (1), 94–130.

Armington, Paul S., "A Theory of Demand for Products Distinguished by Place of Production," *International Monetary Fund Staff Papers*, 1969, *16*, 159–178.

Bas, Maria, Thierry Mayer, and Mathias Thoenig, "From Micro to Macro: Demand, Supply, and Heterogeneity in the Trade Elasticity," *Journal of International Economics*, 2017, *108*, 1–19.

Caliendo, Lorenzo and Fernando Parro, "Estimates of the Trade and Welfare Effects of NAFTA," *Review of Economic Studies*, 2015, *82* (1), 1–44.

**_ and _ **, "Trade policy," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. 5, Elsevier, 2022, pp. 219–295.

_ , Robert C Feenstra, John Romalis, and Alan M Taylor, "Tariff reductions, entry, and welfare: Theory and evidence for the last two decades," Technical Report, National Bureau of Economic Research 2015.

Chaney, Thomas, "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *The American Economic Review*, 2008, *98* (4), 1707–1721.

Christensen, Timothy and Benjamin Connault, "Counterfactual sensitivity and robustness," *Econometrica*, 2023, *91* (1), 263–298.

Costinot, Arnaud and Andrés Rodríguez-Clare, "Trade theory with numbers: Quantifying the consequences of globalization," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. 4, Elsevier, 2014, pp. 197–261.

_ and _, "The US gains from trade: Valuation using the demand for foreign factor services," *Journal of Economic Perspectives*, 2018, *32* (2), 3–24.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Global Rebalancing with Gravity: Measuring the Burden of Adjustment," *IMF Staff Papers*, 2008, (3).

Eaton, Jonathan and Samuel Kortum, "Technology, Geography and Trade," *Econometrica*, 2002, *70* (5), 1741–1779.

_ , _ , and Francis Kramarz, "An Anatomy of International Trade: Evidence from French Firms," *Econometrica*, 2011, *79* (5), 1453–1498.

Einav, Liran, Peter J. Klenow, Jonathan D. Levin, and Raviv Murciano-Goroff, "Customers and retail growth," 2022.

Erhardt, Katharina, Peter Egger, and Sergey Nigai, "Empirical Productivity Distributions and International Trade," 2021. ZBW-Leibniz Information Centre for Economics.

Feenstra, Robert C, "Restoring the product variety and pro-competitive gains from trade with heterogeneous firms and bounded productivity," *Journal of International Economics*, 2018, *110*, 16–27.

Fernandes, Ana M, Peter J Klenow, Sergii Meleshchuk, Martha Denisse Pierola, and Andrés Rodríguez-Clare, "The intensive margin in trade: How big and how important?," *American Economic Journal: Macroeconomics*, 2023, *15* (3), 320–354.

Hansen, Lars Peter and Thomas J Sargent, "Robust control and model uncertainty," *American Economic Review*, 2001, *91* (2), 60–66.

Head, Keith and Thierry Mayer, "Gravity Equations: Toolkit, Cookbook, Workhorse," in Gita Gopinath, Elhanan Helpman, and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. 4, New York: Elsevier, 2014.

_ , _ , and Mathias Thoenig, "Welfare and trade without Pareto," *American Economic Review*, 2014, *104* (5), 310–316.

Helpman, Elhanan, Marc J. Melitz, and Stephen R. Yeaple, "Export Versus FDI with Heterogeneous Firms," *The American Economic Review*, 2004, *94* (1), 300–316.

_ , Marc Melitz, and Yona Rubinstein, "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics*, 2008, *2* (5), 441–487.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

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

Lind, Nelson and Natalia Ramondo, "Trade with correlation," *American Economic Review*, 2023, *113* (2), 317–353.

Melitz, Marc and Stephen Redding, *Heterogeneous Firms and Trade*, Vol. 4 of *Handbook of International Economics*, Elsevier, 2013.

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

Melitz, Marc J and Stephen J Redding, "Heterogeneous firms and trade," *Handbook of international economics*, 2014, *4*, 1–54.

**_ and _ **, "New trade models, new welfare implications," *American Economic Review*, 2015, *105* (3), 1105–1146.

Nigai, Sergey, "A tale of two tails: Productivity distribution and the gains from trade," *Journal of International Economics*, 2017, *104*, 44–62.

Redding, Stephen J., "Quantitative urban economics," in Dave Donaldson and Stephen J. Redding, eds., *Handbook of Regional and Urban Economics*, Vol. 6, Elsevier, 2025, pp. 73–141.

Teti, Feodora A., "Missing Tariffs," Technical Report 11590, CESifo Working Papers 2024. Feodora Teti's Global Tariff Database (v_beta1-2024-12).

A Proofs

We first prove a lemma that establishes that, for the class of distributions relevant to our problem, it is always possible to find a (unique) vector of unobservable trade shifters, $\{A\}$, that can rationalize the trade shares in the Ricardian model.

Lemma 1 (Trade-share inversion)**.** *Let strictly positive wages, $\{w\}$ and tariffs $\{1+t_{od}\}$ be given. Let a distribution F be given that is absolutely continuous with respect to F^* and has strictly positive density on $\mathbb{R}^D_{++}$. Suppose that F satisfies $\mathbb{E}_F[z_o^{\sigma-1}] < \infty$ for all o.*

Then, for every vector of strictly positive empirical trade shares, $\{\tilde{\lambda}\}$, there exists a unique (up to a destination-specific multiplier) vector $\{A\} \in \mathbb{R}^{D^2}_{++}$ such that the Ricardian model generates $\lambda_{od} = \tilde{\lambda}_{od}$ for all od.

Proof. As wages are fixed and the Ricardian model features head-to-head competition only within a destination, it is sufficient to prove inversion destination-by-destination. Thus, for a given d, define $v_o \equiv (\sigma-1)\ln(A_{od}/(w_o(1+t_{od}))$. Then, for any vector of productivities z, the lowest delivered price satisfies $\max_o e^{v_o} z_o^{\sigma-1}$. Because $F \ll F^*$, ties occur with probability zero. Hence, the trade share is

$$\lambda_{od}(v) = \frac{\mathbb{E}_F\left[e^{v_o} z_o^{\sigma-1} 1\{o = \arg\max_m e^{v_m} z_m^{\sigma-1}\right]}{\mathbb{E}_F\left[\max_m e^{v_m} z_m^{\sigma-1}\right]}.$$

Define $Q(v) \equiv \ln \mathbb{E}_F\left[\max_k e^{v_k} z_k^{\sigma-1}\right] - \sum_o \tilde{\lambda}_{od} v_o$. As $F \ll F^*$ and $\mathbb{E}_F[z_o^{\sigma-1}] < \infty$, where $\partial Q/\partial v_o = \lambda_{od}(v) - \tilde{\lambda}_{od}$. Standard arguments then establish that there exists a unique minimizer, v^*, which means the FOC implies $\lambda_{od}(v^*) = \tilde{\lambda}_{od}$.[45] Then, there exists a unique (to scale) vector $\{A^*_{od}\}_o$, with $A^*_{od} \equiv \exp(v^*_o/(\sigma-1))w_o(1+t_{od})$, such that the trade shares in the model match the empirical trade shares for d. The same argument can then be repeated for each destination. □

[45] Since $Q(v)$ is invariant to increasing each v_o by a common constant, normalize $\max_o v_o = 0$. For existence, since $Q(v) \to \infty$ as any $v_o \to -\infty$, there exists some sufficiently large $M > 0$ such that if $\min_o v_o < -M$, then $Q(v) > Q(0)$. Thus, any minimizer of $Q(v)$ is in $[-M, 0]^D$. Then, since $Q(v)$ is continuous and $[-M, 0]^D$ is compact, $Q(v)$ attains its minimum in $[-M, 0]^D$, with minimizer denoted v^*. Moreover, the minimizer v^* to the normalized problem also minimizes $Q(v)$ globally. For uniqueness, Holder's Inequality implies that $Q(v)$ is convex, and full support ensures the inequality is strict, so v^* is unique (up to scale). In particular, define a function $f(v) \equiv \max_o e^{v_o} z_o^{\sigma-1}$. Since, $f(\alpha v + (1-\alpha)w) = \max_o (e^{v_o} z_o^{\sigma-1})^\alpha (e^{w_o} z_o^{\sigma-1})^{1-\alpha} \leqslant f(v)^\alpha f(w)^{1-\alpha}$, taking expectations gives $\mathbb{E}_F[f(\alpha v + (1-\alpha)w)] \leqslant \mathbb{E}_F[f(v)^\alpha f(w)^{1-\alpha}] \leqslant \mathbb{E}_F[f(v)]^\alpha \mathbb{E}_F[f(w)]^{1-\alpha}$, where the second inequality is from Holder's Inequality. Taking logs establishes that $Q(v)$ is convex. As F has full support, the inequality applied to the normalized problem is strict unless $v = w$.

A.1 Proofs of Propositions 1(a) and 1(b)

For simplicity of notation, throughout this and subsequent Appendix sections we continue to write the orthogonality restriction in the form $\sum_o \sum_d \Delta_o \Delta_d \ln(1 + t_{od}) \Delta_o \Delta_d \varepsilon_{od} = 0$; as discussed in footnote 33, we restrict the estimation sample to improve identification, so this notation should be understood as shorthand for residualizing against origin and destination fixed effects.

Proposition 1(a) (Bounds for competitive model). *Suppose that the baseline trade data satisfy, for country j of interest, the condition:*

$$\lambda_{jj}/\lambda_{jd} \geqslant \lambda_{oj}/\lambda_{od} \quad \forall o, d \neq j \tag{C1}$$

Consider all possible productivity distributions F. Then the smallest and largest possible values for j's gains from trade in the competitive model, as defined in program (23), are given by

$$\underline{GT_j} = 0 \qquad \textit{and} \qquad \overline{GT_j} = 1 - \lambda_{jj}^{\frac{1}{\sigma-1}}. \tag{41}$$

Proof. First, we prove that $0 \leqslant GT_j \leqslant 1 - \lambda_{jj}^{1/(\sigma-1)}$ for any distribution and dataset. We show that GT_j is fundamentally about the extent to which j can domestically replace goods that it imports at baseline. The agnostic upper bound is approached when j's latent productivity on such goods is arbitrarily low, and the agnostic lower bound is approached when j's latent productivity on such goods is arbitrarily close to the baseline foreign delivered price. Then, we prove that the bounds are sharp by constructing distributions that achieve GT_j arbitrarily close to the bounds. The core logic is that, if (C1) holds, then, under F^*, if j does not sell a good ω domestically at baseline, j also does not export ω. This means that we can adjust j's productivity on such goods without affecting the baseline equilibrium.

Step 1: Agnostic Bounds We show that $0 \leqslant GT_j \leqslant 1 - \lambda_{jj}^{1/(\sigma-1)}$ at any feasible point of program (23). Begin with the upper bound. The trade share is

$$\lambda_{od} = \frac{1}{\gamma_d} \mathbb{E}_F[p_{od}(\omega)^{1-\sigma} \mathbb{1}\{o = \arg\min_{o'} p_{od}(\omega)\}]. \tag{42}$$

Re-arranging yields an expression for γ_d (with an analogous expression for γ'_d). Taking the case where $o = d = j$, substitute into $GT_j = 1 - (\gamma'_j/\gamma_j)^{1/(\sigma-1)}$ to obtain

$$GT_j = 1 - \left[\lambda_{jj}\frac{\mathbb{E}_F[z_j(\omega)^{\sigma-1}]}{\mathbb{E}_F[z_j(\omega)^{\sigma-1}\mathbb{1}\{j \text{ wins } \omega\}]}\right]^{\frac{1}{\sigma-1}} = 1 - \lambda_{jj}^{\frac{1}{\sigma-1}}\left[1 + \frac{\mathbb{E}_F\left[z_j(\omega)^{\sigma-1}\mathbb{1}\{j \text{ loses } \omega\}\right]}{\mathbb{E}_F\left[z_j(\omega)^{\sigma-1}\mathbb{1}\{j \text{ wins } \omega\}\right]}\right]^{\frac{1}{\sigma-1}}.$$

The term in square brackets is weakly greater than 1, so we have $GT_j \leqslant 1 - \lambda_{jj}^{\frac{1}{\sigma-1}}$. For the lower end, we can simply compare the two price indices directly:

$$\gamma_j \;=\; \mathbb{E}_F\Big[\big(\min_o MC_{oj}\big)^{1-\sigma}\Big] \;\geqslant\; \mathbb{E}_F\big[MC_{jj}^{1-\sigma}\big] \;=\; \gamma_j'. \tag{43}$$

Hence $\gamma_j'/\gamma_j \leqslant 1$ and, thus, $GT_j \geqslant 0$.

Step 2: Define the upper bound distribution We focus on the upper bound first. Let $\{A^*\}$ and $\{w^*\}$ denote the bilateral productivity shifters and wages that rationalize the data under F^*, and let $\kappa_{od}^* \equiv w_o^*(1+t_{od})/A_{od}^*$ denote the implied fixed component of trade costs. Define two random variables: a vector $\tilde{z}(\omega) \sim F^*$, and a scalar $u(\omega) \sim U(0,1)$. Then, for $\eta > 1$, define the random variable

$$z_o(\omega) = \begin{cases} \tilde{z}_o(\omega) & \text{if } o \neq j \text{ or } (o = j \text{ and } j = \arg\min_{o'} \kappa_{o'j}/\tilde{z}_{o'}(\omega)\ ) \\ u(\omega)^\eta \tilde{z}_o(\omega) & \text{if } o = j \text{ and } j \neq \arg\min_{o'} \kappa_{o'j}/\tilde{z}_{o'}(\omega). \end{cases}$$

We denote the distribution of this random variable F_η^+. Intuitively, this is simply the original Frechet distribution except that j's draws have been multiplied by a smaller value for any goods that j imports under the original Frechet. We can make j arbitrarily unproductive on such goods by making η large, while ensuring that F_η^+ remains continuous with strictly positive density on the full support of F^*.

Step 3: Show that trade shares are identical under F_η^+ and F^* if (C1) holds Since the orthogonality restriction is satisfied at $\{A^*\}$, if we can show that the economy under F_η^+ generates the same trade shares at $\{\kappa^*\}$, it follows that F^+ satisfies all the constraints in program (23). The only change between F^* and F^+ is that j's productivity has been lowered on goods that it imports under F^*. It does not change the set of goods that j wins at home, which means it does not affect j's domestic sales. Moreover, it can only affect j's exports on goods that j exported under F^* but did not sell domestically. Formally, j loses on good ω to o at home iff $\kappa_{oj}^*/\tilde{z}_o(\omega) < \kappa_{jj}^*/\tilde{z}_j(\omega)$, or $\kappa_{oj}^*/\kappa_{jj}^* < \tilde{z}_o(\omega)/\tilde{z}_j(\omega)$, i.e., if relative idiosyncratic productivity between o and

j exceeds relative bilateral cost shifters. However, under F^*, the gravity structure implies that relative bilateral cost shifters are a constant elasticity function of relative trade shares, $\kappa^*_{oj}/\kappa^*_{jj} = \left(\lambda_{oj}/\lambda_{jj}\right)^{-1/\theta^*}$. Therefore, if j loses on good ω at home o, some algebra gives

$$\frac{\tilde{z}_o(\omega)}{\tilde{z}_j(\omega)} > \frac{\kappa^*_{oj}}{\kappa^*_{jj}} \geqslant \frac{\kappa^*_{od}}{\kappa^*_{jd}}, \tag{44}$$

which implies that o beats j in d, too. Intuitively, the gravity structure of F^* means that the data condition translates into j having an absolute advantage over o at selling to j, compared to selling to other destinations d. Therefore, if j's idiosyncratic draw for good ω is sufficiently poor relative to o's that j loses ω to o domestically, then *a fortiori* it will lose to o abroad, too. Thus, prices and trade shares are identical under F^* and F^+_η, which also implies that $\{w^*\}$ satisfies trade balance. Finally, the orthogonality restrictions continues to hold as we have held $\{A\}$ at $\{A^*\}$.

Step 4: Show that F^+_η approaches the agnostic upper bound Since $\mathbb{E}[u^x] = 1/(1+x)$ for $x > -1$ if $u \sim U(0,1)$, we can re-use the expression for GT_j from Step 1 and noting that $u(\omega)$ and $\tilde{z}(\omega)$ are independent, we can write,

$$GT_j = 1 - \lambda_{jj}^{\frac{1}{\sigma-1}} \left[1 + \frac{1}{1+\eta(\sigma-1)} \frac{\mathbb{E}_{F^*}\left[z_j(\omega)^{\sigma-1}\mathbb{1}\{j \text{ loses } \omega\}\right]}{\mathbb{E}_{F^*}\left[z_j(\omega)^{\sigma-1}\mathbb{1}\{j \text{ wins } \omega\}\right]}\right]^{\frac{1}{\sigma-1}}. \tag{45}$$

Since $GT_j \to 1 - \lambda_{jj}^{1/(\sigma-1)}$ as $\eta \to \infty$, for any $\epsilon > 0$ we can find a distribution F^+_η that achieves $GT_j > 1 - \lambda_{jj}^{1/(\sigma-1)} - \epsilon$ by setting η sufficiently large. It is also worth noting that the fact that the equilibrium has not changed means that the above steps would all remain valid even if all countries had their marginals adjusted in the same manner as j's, so all countries for which the data condition holds (which is every country in our dataset) can achieve the agnostic bound simultaneously.

Step 5: Analogous steps show that the lower bound is also sharp We follow the same steps, except we modify slightly the adjusted random variable. Define $\bar{z}_j(\omega) \equiv \kappa_{jj}/\min_{o\neq j}(\kappa_{oj}/\tilde{z}_o(\omega))$ as the value that j's productivity would have to take to exactly tie with the minimum imported price under F^*. Then, define

$$z_o(\omega) = \begin{cases} \tilde{z}_o(\omega) & \text{if } o \neq j \text{ or } (o = j \text{ and } j = \arg\min_{o'} \kappa_{o'j}/\tilde{z}_{o'}(\omega)) \\ u(\omega)^{\frac{1}{\eta}} \tilde{z}_j(\omega) & \text{if } o = j \text{ and } j \neq \arg\min_{o'} \kappa_{o'j}/\tilde{z}_{o'}(\omega), \end{cases}$$

and denote its distribution F_η^-, where, again, $\eta \in (1, \infty)$. This implies that $\min_o MC_{oj}(\omega) = MC_{jj}(\omega)$ for goods where j wins, and $\min_o MC_{oj}(\omega) = u(\omega)^{1/\eta} MC_{jj}(\omega)$ for goods where j loses. By the same argument as in Step 3, j continues to lose in every destination for the transformed goods, so prices, trade shares, and the orthogonality restriction remain unchanged. Thus, writing GT_j as $1 - (\gamma_j'/\gamma_j)^{1/(\sigma-1)}$ gives

$$GT_j = 1 - \left(\frac{\mathbb{E}_F[MC_{jj}(\omega)^{1-\sigma}]}{\mathbb{E}_F[(\min_o MC_{oj}(\omega))^{1-\sigma}]} \right)^{\frac{1}{\sigma-1}} = 1 - \left(\lambda_{jj} + (1 - \lambda_{jj}) \frac{1}{1 + (\sigma - 1)/\eta} \right)^{\frac{1}{\sigma-1}}.$$

Since $GT_j \to 0$ as $\eta \to \infty$, for any $\epsilon > 0$ we can find a distribution F_η^- that achieves $GT_j < \epsilon$ by setting η sufficiently large. □

A.1.1 Extending the bounds to multiple sectors with input-output linkages

We can extend the agnostic bounds in Proposition 1(a) to the multi-sector roundabout input-output linkage model of Caliendo and Parro (2015). Suppose there are $s = 1, 2, ..., S$ sectors. Final consumption is Cobb-Douglas over sectors and CES within sector,

$$U_d = \prod_s \left(\frac{Q_{d,s}}{\alpha_{d,s}} \right)^{\alpha_{d,s}}, \qquad Q_{d,s} = \left(\int_0^1 q_{d,s}^{\frac{\sigma_s - 1}{\sigma}}(\omega) d\omega \right)^{\frac{\sigma_s}{\sigma_s - 1}}, \tag{46}$$

with $\sum_s \alpha_{d,s} = 1$. Writing $P_{d,s}$ as the price index of sector s in country d, real income is given by $W_d = w_d L_d / \prod_s P_{d,s}^{\alpha_{d,s}}$. The price of supplying variety ω of sector s from origin o to destination d is $p_{od,s}(\omega) = c_{o,s}(1 + t_{od,s})/(A_{od,s} z_{o,s}(\omega))$, where $c_{o,s} = w_o^{\xi_{o,s}} \prod_{k=1}^S P_{o,k}^{\xi_{o,ks}}$ and $\xi_{o,ks}$ is the value-added share of sector k in producing sector s in country o and $\xi_{o,s} = 1 - \sum_k \xi_{o,ks}$ is the labor share.

The expenditure share of d on goods from o within sector s is given by

$$\lambda_{od,s} = \frac{1}{P_{d,s}^{1-\sigma_s}} \mathbb{E}[p_{od,s}(\omega)^{1-\sigma_s} 1\{o = \arg\min_{o'} p_{o'd,s}(\omega)\}]. \tag{47}$$

Write $M_{o,s} \equiv \mathbb{E}[z_{o,s}(\omega)^{\sigma_s-1} 1\{o = \arg\min_{o'} p_{o'd,s}(\omega)\}]/\lambda_{od,s}$ for simplicity. Then, as in Proposition 1(a), the special case where $o = d$ with the normalization $w_d = w'_d = 1$ implies that changes in the price index are given by

$$\log \hat{P}_{d,s} = \log \hat{c}_{d,s} - \frac{1}{\sigma_s - 1} \log \hat{M}_{d,s} = \sum_k \xi_{d,ks} \log \hat{P}_{d,k} - \frac{1}{\sigma_s - 1} \log \hat{M}_{d,s}. \tag{48}$$

Stacking sectors and defining Ξ_d as a matrix with elements $\xi_{d,ks}$, we can write

$$\log \hat{P}_d = \Xi_d \log \hat{P}_d + M_d = (I - \Xi_d)^{-1} M_d. \tag{49}$$

where M_d is an $S \times 1$ vector with typical element $-1/(\sigma_s - 1) \log \hat{M}_{d,s}$ and $(I - \Xi_d)^{-1}$ is the standard Leontief inverse. The log change in real income is given by

$$\log \hat{W}_d = \log \hat{w}_d - \alpha'_d \log \hat{P}_d = \log \hat{w}_d - \alpha'_d (I - \Xi_d)^{-1} \log \hat{M}_d = \log \hat{w}_d + \sum_s \beta_{d,s} \frac{1}{\sigma_s - 1} \log \hat{M}_{d,s},$$

where $\beta_{d,s}$ is element s of $\alpha'_d (I - \Xi_d)^{-1}$. Exponentiating gives $\hat{W}_d = \prod_s \hat{M}_{d,s}^{\frac{\beta_{d,s}}{\sigma_s - 1}}$. It then follows that gains from trade satisfies an equivalent to Equation (26),

$$GT_d = 1 - \prod_s \lambda_{dd,s}^{\frac{\beta_{d,s}}{\sigma_s - 1}} \left(1 + \frac{\mathbb{E}[z_{d,s}(\omega)^{\sigma_s - 1} \mathbb{1}\{d \neq \arg\min_o MC_{od,s}(\omega)\}]}{\mathbb{E}[z_{d,s}(\omega)^{\sigma_s - 1} \mathbb{1}\{d = \arg\min_o MC_{od,s}(\omega)\}]} \right)^{\frac{\beta_{d,s}}{\sigma_s - 1}}, \tag{50}$$

and the agnostic upper bound is given by $\overline{GT}_d = 1 - \prod_s \lambda_{dd,s}^{\frac{\beta_{d,s}}{\sigma_s - 1}}$. The lower bound is analogous. Sharpness follows under a sector-by-sector analogue of (C1).

Proposition 1(b) (Bounds for monopolistic competition model). *Suppose that the baseline tariff data satisfy the condition that,*

$$\exists i_1 \neq i_2, i_1 \neq j: \qquad \Delta_o \Delta_d \ln(1 + t_{i_1 i_2}) > 0 \tag{C2}$$

Consider all possible productivity distributions F. Then the smallest and largest possible values for j's gains from trade in the monopolistic competition model, as defined in program (28), are given by

$$\underline{GT}_j = 0 \qquad \textit{and} \qquad \overline{GT}_j = 1 - \lambda_{jj}^{\frac{1}{\sigma - 1}}. \tag{51}$$

Proof. We first establish the universal bounds. Then, we use a heterogeneous

Pareto construction and show that the bounds are approached by sending the focal country's shape parameters to $\sigma - 1$ and ∞. Finally, we show that it remains possible to satisfy orthogonality along this sequence.

Step 1: Distribution-agnostic bounds The first step is to show that changes in welfare are summarized entirely by changes in the domestic cutoff. By definition, the cutoff $\hat{z}_{jj}$ is such that $X_{jj}(\hat{z}_{jj}) = \sigma w_j f_{jj}$. Substituting in and re-arranging yields $(\hat{z}_{jj})^{\sigma-1} = \sigma w_j f_{jj} \kappa_{jj}^{\sigma-1} \gamma_j / E_j$. Importantly, everything on the RHS is unchanged between baseline and autarky except γ_j. Dividing the baseline and autarky versions thus gives $\hat{z}'_{jj}/\hat{z}_{jj} = (\gamma'_j/\gamma_j)^{1/(\sigma-1)}$, so $GT_j = 1 - \hat{z}'_{jj}/\hat{z}_{jj}$. Intuitively, the profit to domestic entry depends on residual demand and thus the aggregate price index. Since the costs of entry do not change between baseline and autarky, the change in the profitability of entry co-moves exactly with the change in the price index.

The lower bound is straightforward as we only need to show that the cutoff cannot rise when moving to autarky. This follows directly from free entry: exporting gives an additional source of expected profits at the same entry cost ($w_j f_{Ej}$), so the domestic cutoff must be higher in the trading equilibrium to ensure aggregate profits are the same as in autarky. Formally, profits from selling domestically are

$$\mathbb{E}\left[\left(\frac{1}{\sigma}X_{jj}(z) - w_j f_{jj}\right) 1\{z \geqslant \hat{z}_{jj}\}\right] = \mathbb{E}\left[\left(\frac{1}{\sigma}X_{jj}(\hat{z}_{jj})\left(\frac{z}{\hat{z}_{jj}}\right)^{\sigma-1} - w_j f_{jj}\right) 1\{z \geqslant \hat{z}_{jj}\}\right] = w_j f_{jj} J(\hat{z}_{jj}),$$

where $J(x) \equiv \mathbb{E}[((z/x)^{\sigma-1} - 1)1\{z > x\}]$. Free entry then implies $w'_j f_{jj} J(\hat{z}'_{jj}) = w_j f_{jj} J(\hat{z}_{jj}) + \sum_{d \neq j} w_j f_{jd} J(\hat{z}_{jd})$, i.e., Equation (8) in Melitz and Redding (2014). Since the last term is positive and $J(x)$ is strictly decreasing, it follows that $\hat{z}'_{jj} \leqslant \hat{z}_{jj}$. Thus, $GT_j \geqslant 0$.

The upper bound proceeds similarly. Free entry implies aggregate profits are

$$\mathbb{E}[\pi_j(z(\omega))] = \sum_d \frac{1}{\sigma N_j} X_{jd} - \sum_d w_j f_{jd} \Pr(z_j \geqslant \hat{z}_{jd}) = \frac{1}{\sigma}\frac{E_j}{N_j} - w_j \sum_d f_{jd} \Pr(z_j \geqslant \hat{z}_{jd}),$$

where the second equality uses trade balance to replace total sales with total expenditure. Recall that $E_j = w_j L_j$, and by choice of numeraire we can write $w_j = w'_j = 1$. Then, since free entry in autarky implies $\mathbb{E}_F[\pi'(\omega)] = w_j f_{jj} J(\hat{z}'_{jj})$ and entry costs are unchanged, set $\mathbb{E}_F[\pi_j(\omega)] = \mathbb{E}_F[\pi'_j(\omega)]$ and split up the summation term to obtain

$$\frac{L_j}{\sigma N_j} = f_{jj}\left(J(\hat{z}'_{jj}) + 1 - F_j(\hat{z}_{jj})\right) + \sum_{d \neq j} f_{jd}(1 - F_j(\hat{z}_{jd})). \tag{52}$$

We can use the cutoff equation, $X_{jj}(\hat{z}_{jj}) = \sigma w_j f_{jj}$, to replace f_{jj} with an expression involving the cutoff. Since $X_{jj}(\hat{z}_{jj})$ also depends on γ_j, we can substitute out γ_j by using j's domestic trade share equation, which gives

$$f_{jj} = \frac{L_j}{\sigma N_j}\lambda_{jj}\frac{\hat{z}_{jj}^{\sigma-1}}{\mathbb{E}[z_j^{\sigma-1}1\{z_j \geqslant \hat{z}_{jj}\}]}. \tag{53}$$

Finally, we still need the autarky cutoff, $\hat{z}'_{jj}$. Note that $J(\hat{z}'_{jj}) + 1 - F(\hat{z}_{jj}) = \mathbb{E}[(z_j/\hat{z}'_{jj})^{\sigma-1}1\{z_j \geqslant \hat{z}'_{jj}\}] - \Pr(\hat{z}'_{jj} \leqslant z \leqslant \hat{z}_{jj})$. Substituting (53) into (52) therefore gives

$$1 = \left(\frac{\hat{z}_{jj}}{\hat{z}'_{jj}}\right)^{\sigma-1}\lambda_{jj}\frac{\mathbb{E}[z_j^{\sigma-1}1\{z \geqslant \hat{z}'_{jj}\}] - (\hat{z}'_{jj})^{\sigma-1}\Pr(\hat{z}'_{jj} \leqslant z_j \leqslant \hat{z}_{jj})}{\mathbb{E}[z_j^{\sigma-1}1\{z_j \geqslant \hat{z}_{jj}\}]} + \frac{\sigma N_j}{L_j}\sum_{d \neq j} f_{jd}(1 - F_j(\hat{z}_{jd})).$$

Since $GT_j = 1 - \hat{z}'_{jj}/\hat{z}_{jj}$, re-arranging gives

$$GT_j = 1 - \lambda_{jj}^{\frac{1}{\sigma-1}}\left[\left(1 - \frac{N_j\sum_{d \neq j} f_{jd}\Pr(z_j \geqslant \hat{z}_{jd})}{L_j/\sigma}\right)^{-1}\left(1 + \frac{\mathbb{E}[(z_j^{\sigma-1} - (\hat{z}'_{jj})^{\sigma-1})1\{\hat{z}'_{jj} \leqslant z_j \leqslant \hat{z}_{jj}\}]}{\mathbb{E}[z_j^{\sigma-1}1\{z_j \geqslant \hat{z}_{jj}\}]}\right)\right]^{\frac{1}{\sigma-1}}.$$

The second term is clearly weakly greater than 1. For the first term, observe that the numerator is total export fixed costs paid, while the denominator is total variable profits (as trade balance implies that total expenditure, L_j, equals total income). As variable profits are used to finance all fixed costs, this ratio is in $(0, 1)$ and the overall term in parentheses is greater than 1. Thus, we have $GT_j \leqslant 1 - \lambda_{jj}^{1/(\sigma-1)}$. The first term captures the wedge due to having to pay fixed export costs, f_{jd}, to access the gains from trade, and becomes small by rationalizing aggregate trade flows as the result of a vanishingly small mass of highly productive firms. The second term captures the loss of domestic varieties caused by the cutoff rising when opening up to trade and has an intuition similar to that in the Ricardian model: the welfare loss from moving to autarky will be large if the average productivity of new domestic firms that start producing in response is low.

Step 2: Define the attaining distribution Let j's marginal be $Pareto(1, \theta_j)$ and all other countries' marginals be $Pareto(1, \theta^*)$. While this departs from the standard i.i.d. Pareto case, the usual steps that establish that $N_o = N'_o$ and that $GT_d = 1 - \lambda_{dd}^{1/\theta_d}$ continue to apply country by country. For the upper bound, take $\theta_j \to \sigma - 1$; for the lower bound, take $\theta_j \to \infty$. The only remaining question is whether it is possible to do this without violating the orthogonality restrictions.

Step 3: Verify orthogonality Pareto productivity means that A_{od} and f_{od} are log-linear functions of the cutoff, $\hat{z}_{od}$. This means that the orthogonality conditions can be rewritten as restrictions on the cutoffs. It also means that any cutoff can be rationalized with appropriate choices of A_{od} and f_{od}, which implies that orthogonality has little bite. Formally, simply plugging in the closed-form Pareto conditional expectation to (1) and taking log double differences gives

$$\Delta_o \Delta_d \log A_{od} = \Delta_o \Delta_d \log(1 + t_{od}) + \frac{1}{\sigma - 1} \Delta_o \Delta_d (\theta_o - (\sigma - 1)) \ln \hat{z}_{od} + \frac{1}{\sigma - 1} \Delta_o \Delta_d \ln \lambda_{od}.$$

The cutoff definition gives a similar expression for $\Delta_o \Delta_d \log f_{od}$:

$$\Delta_o \Delta_d \ln f_{od} = \Delta_o \Delta_d \ln \lambda_{od} + \Delta_o \Delta_d \theta_o \ln \hat{z}_{od}. \tag{54}$$

Substituting these into the orthogonality restrictions and using the empirical gravity coefficient of $-\theta^*$ to replace the terms involving $\Delta_o \Delta_d \ln \lambda_{od}$, we obtain

$$\sum_{o,d} \Delta_o \Delta_d \ln \hat{z}_{od} \Delta_o \Delta_d \ln(1 + t_{od}) = \sum_{o,d} (\Delta_o \Delta_d \ln(1 + t_{od}))^2, \text{ and} \qquad \text{(Orth } A\text{)}$$

$$\sum_{o,d} \theta_o \Delta_o \Delta_d \ln \hat{z}_{od} \Delta_o \Delta_d \ln(1 + t_{od}) = \theta^* \sum_{o,d} (\Delta_o \Delta_d \ln(1 + t_{od}))^2 \qquad \text{(Orth } f\text{)}$$

Since we can select cutoffs freely by appropriate choices of $\{A\}$ and $\{f\}$, we can simply set $\hat{z}_{jd} = \hat{z}_{jj}$ for all d, which means j's export row drops out, which implies that θ_j can be freely chosen. Then, since $\theta_o = \theta^*$ in our construction, the (Orth f) becomes equivalent to (Orth A). Simplifying even further, we can also set $\hat{z}_{od} = \hat{z}_{oo}$ for all o, with the exception of one cell $o, d = o_1, d_1$ that satisfies $\Delta_o \Delta_d \ln(1 + t_{o_1 d_1}) > 0$ and $o_1 \neq j$ and $o_1 \neq d_1$ (but can otherwise be chosen arbitrarily) where we set $\ln \hat{z}_{o_1 d_1} = \ln \hat{z}_{o_1 o_1} + \sum_{o,d} (\Delta_o \Delta_d \ln(1 + t_{od}))^2 / \Delta_o \Delta_d \ln(1 + t_{o_1 d_1})$, which verifies the condition. Intuitively, since all foreign origins have $\theta_o = \theta^*$ in our construction,

the relationship between $\{A\}$ and $\{f\}$ among foreign origins is exactly the same as under F^*, which means we can rationalize the trade data while still having variation in the cutoffs that satisfies orthogonality. □

A.2 Proof of Proposition 2

Proposition 2 (Bounds for competitive model with independence). *Define ω_j as the share of residualized tariff variation in rows where $o = j$, and β_j as the negative of the slope of a regression of $\ln(\lambda_{jd}/(1-\lambda_{jd}))$ on residualized tariffs in rows where $o = j$. Suppose the following data condition holds:*

$$\omega_j \beta_j < \sigma - 1. \tag{55}$$

Then, consider all possible productivity distributions F under which countries draw their productivities independently. Then the largest possible value for country j's gains from trade in the competitive model, as defined in program (23), is given by

$$\overline{GT}_j^{\text{indep.}} = \overline{GT}_j = 1 - \lambda_{jj}^{\frac{1}{\sigma-1}}. \tag{56}$$

Further, this value is approached arbitrarily closely by a sequence of distributions with Frechet marginals.

Proof. The proof works by constructing a distribution where the focal country j's marginal is $Frechet(1, \theta_j)$, while all countries have marginal $Frechet(1, \theta_F)$. As the agnostic upper bound is approached when $\theta_j \to \sigma - 1$, the construction raises all other θ_F to compensate and thus satisfy orthogonality. As identifying variation for θ_j comes primarily from j's exports, the data condition requires that j's exports are sufficiently inelastic and/or occupy a sufficiently small share of the overall tariff variation in order to it to be possible to use $o \neq j$ to compensate.

Step 1: Derive and invert the trade shares Consider a distribution where the focal country j's marginal is $Frechet(1, \theta_j)$, while all countries have (common) marginal $Frechet(1, \theta_F)$, and draws are independent. Since the foreign block is i.i.d. $Frechet(1, \theta_F)$, trade shares within this block follow standard gravity expressions,

$$\frac{\lambda_{od}}{1-\lambda_{jd}} = \frac{\kappa_{od}^{-\theta_F}}{\Phi_d}, \qquad \kappa_{od} \equiv \frac{w_o(1+t_{od})}{A_{od}}, \qquad \Phi_d \equiv \sum_{o'} \kappa_{o'd}^{-\theta_F}. \tag{57}$$

The structural residuals required to verify the orthogonality restriction are

$$\ln A_{od} = \ln w_o + \ln(1+t_{od}) + \frac{1}{\theta_F} \ln \frac{\lambda_{od}}{1-\lambda_{jd}} + \frac{1}{\theta_F} \ln \Phi_d. \tag{58}$$

As the expression for λ_{jd} is more involved, we express all other A_{od} relative to it, noting that Lemma 1 ensures that a suitable A_{jd} exists.

Step 2: Verify the orthogonality restriction To difference out $\ln A_{jd}$, define a function $x_d(\theta_j, \theta_F) \equiv \kappa_{jd}^{-\theta_j} \Phi_d^{-\theta_j/\theta_F}$. Substituting into (58) gives

$$\ln \frac{A_{od}}{A_{jd}} = \ln \frac{w_o}{w_j} + \ln\left(\frac{1+t_{od}}{1+t_{jd}}\right) + \frac{1}{\theta_F} \ln \frac{\lambda_{od}}{1-\lambda_{jd}} - \frac{1}{\theta_j} \ln x_d(\theta_j, \theta_F). \tag{59}$$

Define the orthogonality function as $K(\theta_j, \theta_F) \equiv \sum_{o,d} \Delta_o \Delta_d \ln A_{od} \Delta_o \Delta_d \ln(1+t_{od})$. As $\sum_{o,d} \Delta_o \Delta_d \ln A_{od} \Delta_o \Delta_d \ln(1+t_{od}) = \sum_{o,d} \Delta_o \ln A_{od} \Delta_o \Delta_d \ln(1+t_{od})$, we have[46]

$$\begin{aligned} K(\theta_j, \theta_F) = &\sum_o \sum_d \Delta_o \Delta_d \ln(1+t_{od}) \Delta_o \Delta_d \ln(1+t_{od}) + \frac{1}{\theta_F} \sum_{o \neq j} \sum_d \ln\left(\frac{\lambda_{od}}{1-\lambda_{jd}}\right) \Delta_o \Delta_d \ln(1+t_{od}) \\ &+ \frac{1}{\theta_j} \sum_d \ln x_d(\theta_j, \theta_F) \Delta_o \Delta_d \ln(1+t_{jd}). \end{aligned}$$

As the agnostic upper bound is approached when $\theta_j \to \sigma - 1$, we need to check whether there exist values of θ_F such that $K(\sigma-1, \theta_F^+) > 0$ and $K(\sigma-1, \theta_F^-) < 0$. Continuity arguments then establish that, for θ_j sufficiently close to $\sigma - 1$, we can find some $\hat{\theta}_F(\theta_j)$ such that $K(\theta_j, \hat{\theta}_F(\theta_j)) = 0$. Finding θ_F^- is trivial: simply set $\theta_F = \theta_j$, in which case the distribution is i.i.d. $Frechet(1, \theta_j)$, in which case standard gravity applies and we have $\Delta_o \Delta_d \ln A_{od} = \Delta_o \Delta_d \ln(1+t_{od}) + (1/\theta_j) \Delta_o \Delta_d \ln \lambda_{od}$ and $\theta_j < \theta^*$

[46] Write $T_{od} \equiv \ln(1+t_{od})$, $S_{od} \equiv \Delta_o \Delta_d T_{od}$. The wage term cancels as $\sum_{o \neq j} \ln(w_o/w_j) \sum_d S_{od} = 0$. For the x_d term, we have $\sum_d x_d \sum_{o \neq j} S_{od} = -\sum_d x_d S_{jd}$ as $\sum_o S_{od} = 0$. Finally, for the tariff term,

$$\sum_{o \neq j, d} (T_{od} - T_{jd}) S_{od} = \sum_{o \neq j, d} T_{od} S_{od} - \sum_d T_{jd} \sum_{o \neq j} S_{od} = \sum_{o \neq j, d} T_{od} S_{od} + \sum_d T_{jd} S_{jd} = \sum_{o,d} T_{od} S_{od} = \sum_{o,d} S_{od}^2.$$

yields $K(\theta_j, \theta_j) < 0$.

The case for θ_F^+ is more involved, but in what follows we show that the limit of $K(\theta_j, \theta_F)$ as $\theta_F \to \infty$ and $\theta_j \to \sigma - 1$ provides an intuitive and easily checkable data condition. We can see that the second term in $K(\theta_j, \theta_F)$ goes to zero when $\theta_F \to \infty$. We need to evaluate what happens to $\ln x_d$. For a given d, standard Frechet algebra shows that the distribution of the best foreign price is $\Pr(\min_{o \neq j} p_{od}(\omega) < p) = 1 - \exp((p/\Phi_d^{-1/\theta_F})^{\theta_F})$. Then, defining $a_d \equiv \lim_{\theta_F \to \infty} \Phi_d^{-1/\theta_F}$, the best foreign price converges to a degenerate distribution at a_d. It is therefore as if each d is choosing between the best of a $Frechet(1, \theta_j)$ draw and a fixed outside option of a_d. The expenditure on j relative to Foreign is then

$$R_d \equiv \frac{\lambda_{jd}}{1 - \lambda_{jd}} = \left(\frac{\kappa_{jd}}{a_d}\right)^{1-\sigma} \frac{\int_{\bar{z}_d}^{\infty} z^{\sigma-1} dF_j(z)}{F_j(\bar{z}_d)} = \left(\frac{\kappa_{jd}}{a_d}\right)^{1-\sigma} e^{\bar{z}_d^{-\theta_j}} \int_{\bar{z}_d}^{\infty} z^{\sigma-1} dF_j(z), \quad (60)$$

where $\bar{z}_d \equiv \kappa_{jd}/a_d = \lim_{\theta_F \to \infty} x_d(\theta_j, \theta_F)^{-1/\theta_j}$. Define $\epsilon \equiv \theta_j - (\sigma - 1)$. Since $dF_j(z) = \theta_j z^{-\theta_j - 1} e^{-z^{-\theta_j}}$, the integral term is bounded between

$$e^{-\bar{z}_d^{-\theta_j}} \int_{\bar{z}_d}^{\infty} z^{-1-\epsilon} dz \leqslant \int_{\bar{z}_d}^{\infty} z^{-1-\epsilon} e^{-z^{-\theta_j}} dz \leqslant \int_{\bar{z}_d}^{\infty} z^{-1-\epsilon} dz. \quad (61)$$

Since $\int_{\bar{z}_d}^{\infty} z^{-1-\epsilon} dz = \bar{z}_d^{-\epsilon}/\epsilon$, this is equivalent to $(\theta_j/\epsilon) x_d \leqslant R_d \leqslant (\theta_j/\epsilon) x_d e^{x_d}$. Re-arranging gives $R_d e^{-x_d} \leqslant (\theta_j/\epsilon) x_d \leqslant R_d$. Then, taking log differences across d gives $\Delta_d R_d - x_d \leqslant \Delta_d \ln x_d \leqslant \Delta_d \ln R_d + x_{d_0}$. Since $x_d \to 0$ for all d as $\theta_j \to \sigma - 1$, we have $\Delta_d \log R_d \to \Delta_d \log x_d$. Since $\sum_d \ln x_d \Delta_o \Delta_d \ln(1 + t_{od}) = \sum_d \Delta_d \ln x_d \Delta_o \Delta_d \ln(1 + t_{od})$, substituting this back into the expression for K gives

$$K(\sigma - 1, \infty) = \sum_o \sum_d \Delta_o \Delta_d \ln(1 + t_{od}) \Delta_o \Delta_d \ln(1 + t_{od}) + \frac{1}{\sigma - 1} \sum_d \ln\left(\frac{\lambda_{jd}}{1 - \lambda_{jd}}\right) \Delta_o \Delta_d \ln(1 + t_{jd}),$$

where we have again used the fact that $\sum_d \Delta_d x_d \Delta_o \Delta_d y_{od} = \sum_d x_d \Delta_o \Delta_d y_{od}$. The implication of this limiting expression is that if $K(\sigma - 1, \infty) > 0$, then, so long as θ_j is sufficiently close to $\sigma - 1$, it will always be possible to find some $\hat{\theta}_F$ such that $K(\theta_j, \hat{\theta}_F) = 0$. In fact, $K(\sigma - 1, \infty)$ can be re-arranged into a more intuitive form. The condition $K(\sigma - 1, \infty) > 0$ is equivalent to $\omega_j \beta_j < \sigma - 1$, where

$$\omega_j \equiv \frac{\text{Var}_{o=j}\Delta_o\Delta_d \ln(1+t_{od})}{\text{Var}\Delta_o\Delta_d \ln(1+t_{od})}, \qquad \beta_j \equiv -\frac{\text{Cov}_{o=j}\left(\ln\left(\frac{\lambda_{jd}}{1-\lambda_{jd}}\right), \Delta_o\Delta_d \ln(1+t_{od})\right)}{\text{Var}_{o=j}\Delta_o\Delta_d \ln(1+t_{od})}.$$

In particular, ω_j is the share of total residualized tariff variation in the focal country's export row, and β_j is the slope of a regression of j's exports to every d, relative to d's non-j expenditure, on tariffs. Intuitively, the proof works by taking j's shape parameter to $\sigma - 1$ and raising other countries' common shape parameter to compensate. Since changing j's shape parameter primarily affects its behavior as an origin, it is most directly identified by exporting behavior. Thus, the condition essentially requires that French exports cannot be very elastic and occupy a large share of the tariff variation, as otherwise it not be possible for the model to rationalize the empirical gravity coefficient by using other cells to compensate.

Step 3: Show that the bound is approached as $\theta_j \to \sigma - 1$ Following the steps in 1(a), GT_j can be expressed as $GT_j = 1 - (\lambda_{dd}/(1-L))^{1/(\sigma-1)}$, where $L \equiv \mathbb{E}_F[z_j(\omega)^{\sigma-1}1\{j \text{ loses } \omega \text{ in } j\}]/\mathbb{E}_F[z_j(\omega)^{\sigma-1}]$. Let $\bar{z}_j(\omega)$ denote the value of z that j would need for its price to equal the lowest foreign price for ω. Then, we have

$$L = \frac{\mathbb{E}_F[z_j(\omega)^{\sigma-1}1\{z_j(\omega) < \bar{z}_j(\omega)\}]}{\mathbb{E}_F[z_j(\omega)^{\sigma-1}]} \leqslant \frac{\mathbb{E}_F[\sqrt{z_j(\omega)^{\sigma-1}\bar{z}_j(\omega)^{\sigma-1}}]}{\mathbb{E}_F[z_j(\omega)^{\sigma-1}]} \leqslant \frac{\mathbb{E}_F[\sqrt{z_j(\omega)^{\sigma-1}}]\mathbb{E}_F[\sqrt{\bar{z}_j(\omega)^{\sigma-1}}]}{\mathbb{E}_F[z_j(\omega)^{\sigma-1}]}.$$

If we can bound $\mathbb{E}_F[\sqrt{\bar{z}_j(\omega)^{\sigma-1}}]$ then we are done, as the denominator goes to ∞ while both numerator terms are bounded. The trade share equation implies $\lambda_{jj} = \mathbb{E}_F[z_j(\omega)^{\sigma-1}1\{z_j(\omega) > \bar{z}_j(\omega)\}]/\mathbb{E}_F[\max\{z_j(\omega)^{\sigma-1}, \bar{z}_j(\omega)^{\sigma-1}\}]$, which in turn implies $\sqrt{\mathbb{E}_F[\bar{z}_j(\omega)^{\sigma-1}]} \leqslant \sqrt{\mathbb{E}_F[z_j(\omega)^{\sigma-1}]}/\sqrt{\lambda_{jj}}$. Using Jensen's Inequality then gives $L < (\mathbb{E}_F[\sqrt{z_j(\omega)^{\sigma-1}}]/\sqrt{\mathbb{E}_F[z_j(\omega)^{\sigma-1}]})/\sqrt{\lambda_{jj}}$. Since the denominator goes to ∞ while the numerator converges to a finite number, we have that L can be made arbitrarily close to 0 and thus GT_j arbitrarily close to the agnostic upper bound. □

A.3 Proof of Proposition 3

Proposition 3 (Bounds for competitive model with common marginal distributions). *Let j index a country of interest. Consider a partition of all foreign countries, $o \neq j$, into K mutually exclusive blocks, $S_1, S_2, ..., S_K$, with trade flows assigned to blocks based*

on origin. Then, define $g_k \equiv \nu_k \beta_k$, as block k's contribution to the full-sample gravity regression, where ν_k is the share of residualized tariff variation in block k and β_k is the negative of the slope of a regression of trade shares against residualized tariffs in block k. Suppose that the following condition holds:

$$\sum_{k:g_k>0} g_k > \theta^* - (\sigma - 1). \tag{62}$$

Then, consider all possible productivity distributions F under which countries draw their productivities from a common marginal distribution but with potential dependence across countries. The largest possible value for country j's gains from trade in the competitive model, as defined in program (23), is given by

$$\overline{GT}_j^{\text{CM}} = \overline{GT}_j = 1 - \lambda_{jj}^{\frac{1}{\sigma-1}}. \tag{63}$$

Further, this bound is approached arbitrarily closely by a distribution with Frechet$(1, \theta)$ marginals.

Proof. The proof works by constructing an economy where all marginals are *Frechet*$(1, \theta)$ and then letting θ approach $\sigma - 1$, using the dependence structure to preserve the orthogonality condition. It constructs blocks of foreign countries with positive within-block correlation. Positive correlation increases the model-implied trade elasticity because it decreases the importance of idiosyncratic productivity differences as compared to fixed trade shifters, such as tariffs. By maximizing this correlation, the model interprets a negative empirical covariance between trade flows and tariffs within foreign blocks as evidence of a high trade elasticity. The data condition states it is possible to construct blocks such that the share of the observed gravity relationship that comes from intra-block variation is large enough to rationalize the full gap between the estimated elasticity, θ^*, and $\sigma - 1$.

Step 1: Derive Trade Shares under Nested Gravity Partition all countries other than the focal country, j, into K mutually exclusive blocks, $S_1, S_2, ..., S_K$. Let $\{z_o\}$ be distributed according to a two-level nested-logistic (GEV) max-stable law

$$\Pr(\{z\} < \{\bar{z}\}) = \exp\Big(-\bar{z}_j^{-\theta} - \sum_{k=1}^{K} \Big(\sum_{o \in S_k} \bar{z}_o^{-\theta/(1-\rho_k)} \Big)^{1-\rho_k} \Big). \tag{64}$$

This has the property that j and each block S_k are independent (the upper nest) and each block S_k has intra-block dependence $\rho_k \in [0,1)$ (the lower nest). Each marginal is $Frechet(1,\theta)$. Note that if $\rho_k = 0$ for all k, this becomes i.i.d. $Frechet(1,\theta)$.

Denote the fixed component of bilateral trade costs as $\kappa_{od} \equiv w_o(1+t_{od})/A_{od}$. Then, following standard Frechet derivations, the distribution of the minimum price is given by $\Pr(\min_o p_{od}(\omega) \leqslant p) = 1 - \exp(-\Phi_d p^\theta)$ with

$$\Phi_d \equiv \kappa_{jd}^{-\theta} + \sum_k \bar{\kappa}_{kd}^{-\theta}, \qquad \bar{\kappa}_{kd} \equiv \left(\sum_{o \in S_k} \kappa_{od}^{-\varepsilon_k} \right)^{-\frac{1}{\varepsilon_k}}, \qquad \varepsilon_k \equiv \frac{\theta}{1-\rho_k}.$$

Define $\Lambda_{kd} \equiv \sum_{o \in S_k} \lambda_{od}$ as the share of d's total expenditure allocated to block S_k. Then, the system satisfies gravity in the upper nest (across blocks) with elasticity θ and gravity in the lower nest (within blocks) with elasticity $\varepsilon_k \in [\theta, \infty)$,

$$\lambda_{jd} = \frac{\kappa_{jd}^{-\theta}}{\Phi_d}, \qquad \Lambda_{kd} = \frac{\bar{\kappa}_{od}^{-\theta}}{\Phi_d}, \qquad \frac{\lambda_{od}}{\Lambda_{kd}} = \frac{\kappa_{od}^{-\varepsilon_k}}{\bar{\kappa}_{kd}^{-\varepsilon_k}} = \frac{\kappa_{od}^{-\varepsilon_k}}{\sum_{o' \in S_k} \kappa_{o'd}^{-\varepsilon_k}}. \tag{65}$$

Note that the within-block elasticity, ε_k, goes to ∞ as ρ_k goes to 1. Intuitively, increasing intra-block correlation reduces the importance of idiosyncratic productivity draws in generating trade relative to fixed trade shifters, such as tariffs.

Step 2: Observe that θ is the Welfare Elasticity We can see from (65) that, for $d = j$, a uniform shift $q > 1$ to foreign prices decreases Λ_{kd} with elasticity θ, i.e., the relevant welfare elasticity is θ. This implies that the usual formula from Arkolakis et al. (2012) holds, so gains from trade is given by $GT_j = 1 - \lambda_{jj}^{\frac{1}{\theta}}$. The nested gravity model therefore attains the agnostic upper bound if and only if $\theta \to \sigma - 1$.

Step 3: Invert Trade Shares to Solve for A_{od} Having derived the trade shares, we need to invert them to obtain $\{A\}$ to verify the orthogonality restriction. For $o = j$, standard gravity holds with elasticity θ, $\ln \kappa_{jd} = -\frac{1}{\theta} \ln \lambda_{jd} - \ln \frac{1}{\theta} \Phi_d$. For $o \neq j$, define $\ell_{od}^k \equiv 1\{o \in S_k\} \ln \frac{\lambda_{od}}{\Lambda_{kd}}$ as the log share of d's within-k expenditure on o (and 0 if $o \notin S_k$). Then, some algebra yields $\ln \kappa_{od} = -\frac{1}{\theta} \ln \Lambda_{kd} - \frac{1}{\varepsilon_k} \ell_{od}^k - \frac{1}{\theta} \ln \Phi_d$. Defining $\ell_{od}^{\text{upper}} = 1\{o = j\} \ln \lambda_{od} + \sum_k 1\{o \in S_k\} \ln \Lambda_{kd}$, we have

$$\Delta_o \Delta_d \ln A_{od} = \Delta_o \Delta_d \ln(1 + t_{od}) + \frac{1}{\theta} \Delta_o \Delta_d \ell_{od}^{\text{upper}} + \sum_k \frac{1}{\varepsilon_k} \Delta_o \Delta_d \ell_{od}^k. \tag{66}$$

Step 4: Impose Orthogonality Our objective is to construct a distribution for which orthogonality holds. Define a function $K(\theta, \{\rho_k\}) \equiv \sum_o \sum_d \Delta_o \Delta_d \ln A_{od} \Delta_o \Delta_d \ln(1 + t_{od})$, where orthogonality holds when $K = 0$. Substituting in gives

$$\begin{aligned} K(\theta, \{\rho_k\}) = &\sum_{o,d} (\Delta_o \Delta_d \ln(1 + t_{od}))^2 + \frac{1}{\theta} \sum_{o,d} \Delta_o \Delta_d \ell_{od}^{\text{upper}} \Delta_o \Delta_d \ln(1 + t_{od}) \\ &+ \sum_k \frac{1}{\varepsilon_k} \sum_{o,d} \Delta_o \Delta_d \ell_{od}^k \Delta_o \Delta_d \ln(1 + t_{od}). \end{aligned}$$

To simplify further, note that $\ell_{od}^{upper} = \ln \lambda_{od} - \sum_k \ell_{od}^k$ and that the empirical gravity regression implies $\sum_{o,d} \Delta_o \Delta_d \ln \lambda_{od} \Delta_o \Delta_d \ln(1 + t_{od}) = -\theta^* \sum_{o,d} (\Delta_o \Delta_d \ln(1 + t_{od}))^2$. Some algebra then yields $K(\theta, \{\rho_k\}) = \frac{V_\tau}{\theta} \left[(\theta - \theta^*) + \sum_k g_k \left(1 - \frac{\theta}{\varepsilon_k}\right) \right]$, where $V_\tau \equiv \sum_{o,d} (\Delta_o \Delta_d \ln(1 + t_{od}))^2$ and

$$g_k \equiv \underbrace{\frac{\sum_{o \in S_k, d} (\Delta_o \Delta_d \ln(1 + t_{od}))^2}{\sum_{o,d} (\Delta_o \Delta_d \ln(1 + t_{od}))^2}}_{\nu_k} \underbrace{\left(- \frac{\sum_{o \in S_k, d} \ell_{od}^k \Delta_o \Delta_d \ln(1 + t_{od})}{\sum_{o \in S_k, d} (\Delta_o \Delta_d \ln(1 + t_{od}))^2} \right)}_{\beta_k} \tag{67}$$

is the block k's contribution to the slope of the full sample relationship between trade shares and residualized tariffs. This can be decomposed into the negative of the slope of a within-k gravity regression of residualized trade shares on residualized tariffs weighted by k's share of total residualized tariff variance.

Setting $K = 0$ yields a convenient expression: orthogonality holds if and only if

$$\theta^* = \theta + \sum_k g_k \rho_k. \tag{68}$$

This condition says that if the orthogonality condition holds, then the estimated gravity coefficient, θ^*, can be decomposed into the elasticity of the outer nest, $\theta < \theta^*$, and an additional term capturing the within-nest dependence weighted by the within-nest gravity slope. Intuitively, if the model has an outer elasticity θ that is lower than the overall empirical elasticity, then it must compensate by generating a higher elasticity elsewhere. This can be done by inducing positive dependence

within the foreign nests. Intuitively, positive within-nest dependence increases the within-nest trade elasticity as it decreases the importance of idiosyncratic productivity draws relative to fixed cost shifters. Indeed, recall that the within-nest trade elasticity is $\theta/(1-\rho_k)$, which goes to ∞ as ρ_k goes to 1.

Recall that the nested gravity model achieves the agnostic upper bound if and only if $\theta \to \sigma - 1$. This amounts to asking whether we can construct blocks $S_1, S_2, ..., S_K$ with intra-block dependence parameters $\rho_1, \rho_2, ..., \rho_K$ such that (68) holds with $\theta = \sigma - 1$. To facilitate a simple continuity argument, define $\hat{K}(\{\rho_k\}) \equiv (\sigma-1) - \theta^* + \sum_k g_k \rho_k$. We can set $\rho_k = 0$ for all k to obtain some $\hat{K}^- < 0$. The largest value that $K(\{\rho_k\})$ can take is when $\rho_k \to 1$ for all k where $g_k > 0$ and $\rho_k = 0$ for all k where $g_k < 0$. In other words, generate maximum intra-block dependence in blocks where the gravity slope is positive. Thus, a sufficient condition for this construction to achieve the agnostic upper bound is $\sum_{k:g_k>0} g_k > \theta^* - (\sigma - 1)$. □

A.4 Proof of Proposition 4

Condition 1 (Three-part distribution data condition)**.** *Let j index a country of interest. The following is true:*

$$\lambda_{jj} > \max_{o \neq j} \lambda_{oj}. \tag{69}$$

Sort j's expenditure shares for all $o \neq j$ in descending order, $\lambda_{(1)j}, \lambda_{(2)j}, ..., \lambda_{(D-1)j}$, and let $\underline{\nu} \in (0,1)$ denote the lowest value of ν for which $\nu(1-\nu^k)\lambda_{jj} > (1-\nu)\sum_{r=2}^{k+1} \lambda_{(r)j}$ holds for all $k = 1, 2, ..., D-1$. There exists some $\nu \in (\underline{\nu}, 1)$ such that

$$K(\nu) < 0, \tag{70}$$

with $K(\nu)$ defined as in (78).

Proposition 4 (Bounds for the competitive model with i.i.d. productivities)**.** *Suppose the baseline tariff and trade data satisfy Condition 1. Then, there exists a distribution under which countries draw productivities independently from a common marginal distribution that attains the agnostic upper bound in (24), i.e.,*

$$\overline{GT}_j^{\text{i.i.d.}} \;=\; 1 - \lambda_{jj}^{1/(\sigma-1)}.$$

Proof. The proof proceeds by constructing a discrete distribution with three possible

values: 0 , 1, and $H > 1$. We construct A_{od} to be such that j always wins among countries with the same draw, so it loses only when its own draw is 0 or it draws 1 and another country draws H. Then, by letting the probability of drawing H go to zero, we approach an economy where j's average productivity of goods it imports at baseline is 0 and thus j achieves the agnostic upper bound. As the discrete distribution implies a rigid relationship between bilateral cost shifters and trade shares, foreign countries are combined into blocks with identical bilateral trade shifters to create flexibility by allowing ties to be broken judiciously. The first part of Condition 1 ensures that this procedure is consistent with j always winning when it is active. To satisfy orthogonality, the value of H grows at the same rate as its probability goes to zero, which means it maintains a first-order impact on expenditure. The value of this impact affects the orthogonality restriction, so adjusting it provides a lever to adjust orthogonality, and the second part of Condition 1 states that it can adjust it enough for orthogonality to hold exactly.

Step 1: Construct the distribution and verify that it can match trade shares Let the common i.i.d. marginal be given by

$$z_o^{\sigma-1} = \begin{cases} 0 & \text{with probability } \nu \\ 1 & \text{with probability } 1-\nu-\eta \\ H & \text{with probability } \eta, \end{cases} \tag{71}$$

where $H = 1 + c/\eta$ and $c > 0$. We now need to calculate the trade shares. For a given good, d buys from the o with the lowest price, which combines the draw from (71) and the fixed component $\kappa_{od} \equiv w_o(1+t_{od})/A_{od}$. Since we will be interested in the limiting case where $\eta \to 0$, let η be sufficiently small that a country receiving the H draw is always cheaper than a country receiving the 1 draw. This means that we only need to compare κ_{od} between countries with the same draw.

Consider a rule that assigns each origin to a priority rank k such that d always buy from the origin with lowest k in the event that multiple origins have the same draw (and no origin has a better draw). If κ_{od} are not equal across origins, the ranking is simply the ranking of κ_{od}. However, we will be interested in the case where κ_{od} are equal across some origins, in which case the priority ranking describes a tie-breaking rule. There are two ways that rank k can win. If rank k draws H, then

it wins if and only if no higher-ranked origin also draws H. This has contribution to expenditure $H\eta(1-\eta)^{k-1} = (\eta + c)(1-\eta)^{k-1}$. If rank k draws 1, then it wins if and only if no other higher-ranked origin also ranks 1, and if no other origin draws H. This has expenditure contribution $(1-\nu-\eta)\nu^{k-1}(1-\eta)^{D-k}$. Therefore, the expected contribution to d's expenditure of an origin of rank k is

$$r_k(\nu, c, \eta) \equiv (\eta + c)(1-\eta)^{k-1} + (1-\nu-\eta)\nu^{k-1}(1-\eta)^{D-k}. \tag{72}$$

As $\eta \to 0$, $r_k(\nu, c, \eta) \to (1-\nu)\nu^{k-1} + c$, which shows that the the H draw can have vanishingly small probability but maintain a first-order impact on expenditure. We will write $r_k(\nu, c) \equiv \lim_{\eta \to 0} r_k(\nu, c, \eta)$ to denote the limiting case.

Trade shares are given by $\lambda_{(k)d} = \kappa_{(k)d}^{1-\sigma} r_k / \sum_l \kappa_{(l)d}^{1-\sigma} r_l$, which implies $\kappa_{(k)d}^{1-\sigma} \propto \lambda_{(k)d}/r_k$. The condition required for a given priority rule and set of probabilities (ν, η) to be consistent with the sourcing structure of the Ricardian model is $\kappa_{(1)d} \leqslant \kappa_{(2)d} \leqslant ... \leqslant \kappa_{(D)d}$, which is equivalent to $\lambda_{(1)d}/r_1 \geqslant \lambda_{(2)d}/r_2 \geqslant ... \geqslant \lambda_{(D)d}/r_D$. The construction has two levels of rankings. An outer level assigns countries to blocks, where the ranking is over blocks: d always buys from a country in a block of rank k over a country in a block of rank $k' > k$ when both have the same draw. We require that all countries in the same block share the same $\kappa_{(k)d}$, which we can set to $\kappa_{(k)d}^{1-\sigma} = \Lambda_{(k)d}/R_k$, where $\Lambda_{(k)d} \equiv \sum_{o \in (k)} \lambda_{od}$ and $R_k \equiv \sum_{o \in (k)} r_k$. Thus, a priority ranking over $K \leqslant D$ blocks is consistent if and only if

$$\Lambda_{(1)d}/R_1 \geqslant \Lambda_{(2)d}/R_2 \geqslant ... \geqslant \Lambda_{(K)d}/R_K. \tag{73}$$

To determine the outer ranking, consider a simple algorithm. First, let each block be a singleton, and propose an initial ranking based on sorting $\{\lambda_{od}\}_o$ in descending order. If (73) holds, select this ranking. If (73) does not hold, merge the lowest-ranked adjacent pair for which it does not hold and again check (73). Repeat this process until (73) holds. This process generates a mapping of origins to blocks and a priority ranking across blocks that allocates expenditure $\Lambda_{(k)d}$ to the k-ranked block in a way that is consistent with the sourcing structure of the Ricardian model.

Within each block, there is a ranking across origins. As all origins within a block have the same κ_{od}, allocating expenditure across origins must be done entirely through breaking ties. By construction, the total model-implied expenditure in each block is exactly sufficient to match all of the empirical trade shares within that block.

A tie-breaking rule can be implemented as a lottery over priority rankings, so the core question is whether there exists a convex combination of priority rankings (within a given block) that matches the (within-block) empirical trade shares exactly. To answer this, we will use majorization. Consider a block B containing origins with ranks $\{r, r+1, ..., s\}$. Let k denote a member origin's original ranking in $\{\lambda_{od}\}_o$. Then, define $q_{k,d} \equiv \lambda_{(k)d}/\Lambda_{Bd}$ as (k)'s within-block empirical expenditure share and $a_k \equiv r_k/R_B$ as k's within-block model-implied expenditure contribution. Stack these into two vectors, $q_d = (q_{r,d}, q_{r+1,d}, ..., q_{s,d})$ and $a \equiv (a_r, a_{r+1}, ..., a_s)$. Since both vectors sum to 1, q_d is in the convex hull of a if and only if a majorizes q_d, i.e.,

$$\sum_{l=r}^{m} q_{l,d} \leqslant \sum_{l=r}^{m} a_l \qquad \text{for } m = r, r+1, ..., s-1. \tag{74}$$

We now show that the algorithm used to construct blocks guarantees that (74) holds. Define $v_{kd} \equiv \Lambda_{kd}/R_k$ for a block of rank k. Consider a merge of blocks in ranks i and j with $i < j$. Since they were merged, it must be that $v_{id} < v_{jd}$. Then, the merged block has $v_{\text{new},d} \in (v_{id}, v_{jd})$. This implies

$$\frac{\Lambda_{id}}{R_i} < \frac{\Lambda_{id} + \Lambda_{jd}}{R_i + R_j} \iff \frac{\Lambda_{id}}{\Lambda_{id} + \Lambda_{jd}} < \frac{R_i}{R_i + R_j}, \tag{75}$$

which is the majorization property required. Since this applies at all merge events, including initial merges with $\lambda_{(k)d}/r_k$, it follows that (74) holds.

Step 2: Ensure j achieves the agnostic upper bound We have shown that it is possible to match trade shares for all o, d. We know from Proposition 1(a), j achieves the agnostic upper bound is achieved if $\mathbb{E}_F[z_j(\nu)^{\sigma-1}1\{j \text{ imports } \omega\}]/\mathbb{E}_F[z_j(\omega)^{\sigma-1}] \to 0$. In this construction, this requires j to have priority rank 1 for $d = j$. To see this, note that if j is priority rank 1 then the only event in which j imports ω and has non-zero latent productivity on ω is if j draws 1 but some $o \neq j$ draws H. This implies $\mathbb{E}_F[z_j(\omega)^{\sigma-1}1\{j \text{ imports } \omega\}] = (1-\nu-\eta)(1-(1-\eta)^{D-1})$. As this goes to zero as η goes to zero, while $\mathbb{E}_F[z_j(\omega)^{\sigma-1}]$ converges to a positive number, j achieves the agnostic upper bound in the limit as $\eta \to 0$ if j has priority rank 1.

As we will require c to adjust to satisfy orthogonality, a sufficient condition is that j is priority rank 1 for every $c > 0$. As the original ordering was in descending order

of λ_{od}, priority rank 2 could in principle consist of country (2) only, or countries $\{(2),(3)\}$, or countries $\{(2),(3),(4)\}$, and so on, i.e., $\{(2),(3),...,(s)\}$ for $s \leqslant D$. Thus, a condition that ensures that j is never merged with position 2 in the limiting economy as $\eta \to 0$ is:

$$\frac{\lambda_{jj}}{r_1(\nu,c)} > \frac{\sum_{k=2}^{K+1}\lambda_{(k)j}}{\sum_{k=2}^{K+1} r_k(\nu,c)} \qquad K = 1,2,...,D-1. \tag{76}$$

Strictness ensures that if (76) holds, then it also holds in the $\eta > 0$ economy if η is sufficiently small. Re-arranging and noting that $r_k(\nu,c) = (1-\nu)\nu^{k-1} + c$ gives

$$\lambda_{jj}\sum_{k=2}^{K+1}(1-\nu)\nu^{k-1} - (1-\nu)\sum_{k=2}^{K+1}\lambda_{(k)j} + c\left(K\lambda_{jj} - \sum_{k=2}^{K+1}\lambda_{(k)j}\right) > 0. \tag{77}$$

We want (77) to hold for all $c > 0$. The first part of the data condition, (69), ensures that the term in parentheses is strictly positive. Thus, since the sum of the first two terms is exactly 0 if $\nu = 1$, we can always ensure holds by setting ν sufficiently close to 1. Overall, then, Steps 1 and 2 show that if we take η sufficiently small we can construct a Ricardian economy in which all trade shares are matched exactly and j's gains from trade becomes arbitrarily close to the agnostic upper bound. All that remains is to check whether this economy can be made to satisfy orthogonality.

Step 3: Verify orthogonality The algorithm used to construct the economy in Steps 1 and 2 generates a specific vector of κ_{od}, denoted κ_{od}^B. Since $\Delta_o\Delta_d \ln A_{od} = \Delta_o\Delta_d \ln(1+t_{od}) - \Delta_o\Delta_d \ln \kappa_{od}$, define a function

$$K(\nu,c) = \text{Cov}(\Delta_o\Delta_d \ln \kappa_{od}^B, \Delta_o\Delta_d \ln(1+t_{od})) - \text{Var}(\Delta_o\Delta_d \ln(1+t_{od})), \tag{78}$$

where orthogonality holds if and only if $K(\nu,c) = 0$. Since K is continuous in c, if we can find values c_1, c_2 such that $K(\nu,c_1) < 0 < K(\nu,c_2)$, then there exists some $c^* \in (c_1,c_2)$ for which orthogonality holds.

For the positive side, consider the limit as $c \to \infty$. Since the expenditure contribution of country of rank k in the limiting economy as $\eta \to 0$ is $(1-\nu)\nu^{k-1} + c$, when c becomes sufficiently large all expenditure contributions are approximately equal and we have $\kappa_{od}^{1-\sigma} \propto \lambda_{od}$ in the limit, which implies $(1-\sigma)\Delta_o\Delta_d \ln \kappa_{od} \to \Delta_o\Delta_d \ln \lambda_{od}$.

This implies that $\lim_{c\to\infty} K(\nu, c) = -1/(\sigma-1)\text{Cov}(\Delta_o\Delta_d \ln \lambda_{od}, \Delta_o\Delta_d \ln(1+t_{od})) - \text{Var}(\Delta_o\Delta_d \ln(1+t_{od}))$. Since the empirical gravity coefficient is $-\theta^*$, substituting in implies that $\lim_{c\to\infty} K(\nu, c) > 0$ if and only if $\theta^* > \sigma - 1$. Since this is strictly positive (as we require the economy to have a finite price index under F^*), there exists some c_2 such that $K(\nu, c_2) > 0$ for any $\nu \in (0, 1)$.

The negative side is more challenging as the algorithmic construction does not yield a simple expression for κ^B_{od}. The data condition therefore amounts to checking whether $K(\nu, c) < 0$ for some ν, c. For simplicity, the data condition checks only $c = 0$ as then the limiting economy is Bernoulli and thus simpler to evaluate, but this can be relaxed to only require $K(\nu, c) < 0$ for some c. As the chosen ν must also satisfy (77), which becomes tighter as ν decreases, in practice we need only consider ν above the minimum value that satisfies (77), denoted $\underline{\nu}$.

Thus, as Steps 1 and 2 are valid for any $c > 0$, if data condition (70) holds, then orthogonality is verified. Again, we require strictness to ensure that in the $\eta > 0$ economy we can still find c_1, c_2 such that $K(\nu, c_1) < 0 < K(\nu, c_2)$ and thus satisfy orthogonality exactly if we take η to be sufficiently small.

Step 4: Approximate the Three-Part Distribution with a Continuous Distribution We have shown that, if the data condition holds, then we can construct a discrete distribution that generates a Ricardian economy that matches all trade shares exactly, yields GT_j arbitrarily close to the agnostic upper bound, and verifies the orthogonality restriction exactly. To obtain a distribution that is absolutely continuous with respect to F^*, we can approximate the discrete distribution as follows. Consider a distribution of $x_o \equiv z_o^{\sigma-1}$, $G_{\nu,\epsilon,\rho,\eta}(x)$, with marginal defined by the following mixture

$$G_{\nu,\epsilon,\rho,\eta} = \nu U[0,\epsilon] + (1-\nu-\eta-\rho)U[1,1+\epsilon] + \eta U\left(1+\frac{c}{\eta}, 1+\frac{c}{\eta}+\epsilon\right) + \rho G^*,$$

where $\epsilon, \rho, \eta > 0$ and G^* is $Frechet(1, \theta^*/(\sigma-1))$. This is continuous with positive density on $(0, \infty)$, and ties occur with probability zero. By Lemma 1, this economy yields a unique vector $\{\kappa^G\}$ that rationalizes the trade shares. Moreover, by letting ϵ, ρ, η be sufficiently small, standard arguments establish that $\{\kappa^G\}$ become arbitrarily close to the $\{\kappa\}$ that obtains under the discrete distribution. As the data condition ensuring j wins when active holds strictly, j never has ties in the discrete economy,

so by taking ϵ, ρ, η to 0 the convergence in $\{\kappa\}$ ensures this property holds in the continuous economy. For orthogonality, since the data condition holds strictly and K is continuous in $\{\kappa\}$, for sufficiently small ϵ, ρ, η we also have $K(\nu, c_1) < 0 < K(\nu, c_2)$ in the continuous economy and thus orthogonality holds exactly. □

A.5 Additional Theoretical Results

The following Lemma shows that, in the monopolistic competition model, if there are no restrictions on the absolute level of the bilateral fixed costs, $\{f\}$, then any achievable GT_d can be achieved at arbitrarily low divergence.

Lemma 2 (Unrestricted Pareto tail dilutes divergence)**.** *Define a feasible monopolistic competition economy as a combination of a distribution F and a parameter vector ψ that satisfy all of the constraints in Equation (28). Suppose there exists a feasible monopolistic competition economy that achieves gains from trade g with divergence $D(F \mid\mid F^*) < \infty$. Then, for any $\epsilon > 0$, there exists another feasible monopolistic competition economy, $(F_\epsilon, \psi_\epsilon)$, that achieves gains from trade g with divergence $D(F_\epsilon \mid\mid F^*) < \epsilon$.*

Proof. We will show that for any scalar $c > 1$, there exists a distribution F_c and vector of parameters ψ^c that satisfy all of the constraints in (28) and achieve gains from trade g with divergence $D(F_c \mid\mid F^*) \to 0$ as $c \to \infty$. We will construct F_c to apply the same change of measure to F^* as F, but shifted into the tail of F^*.

Define the shifted distribution: Let $M(\{z\}) \equiv \frac{dF(\{z\})}{dF^*(\{z\})}$ be the change of measure that defines the joint distribution F, and let $m(z) \equiv \frac{dF(z)}{dF^*(z)}$ be the change of measure for a particular origin's marginal, where the origin will be clear from context. We will apply the construction separately for each origin, which simplifies notation and is without loss as the absence of head-to-head competition implies that aggregate quantities depend only on one origin's marginal at a time.

Consider a function $m_c(z) = 1$ for $1 \leqslant z < c$ and $m_c(z) = m(z/c)$ for $z \geqslant c$. If $m_c(z)$ induces a valid distribution, F_c, this distribution is identical to F^* in $[1, c)$ and then applies the change of measure $m(z)$ from c onwards. To verify that $m_c(z)$ induces a valid distribution, we need $m_c(z) \geqslant 0$ for all z in the support of F^* and $\mathbb{E}_{F^*}[m_c(z)] = 1$ (see Section B.1 for details). The former is straightforward as $m(z/c) \geqslant 0$. The latter holds as F^* is $Pareto(1, \theta^*)$, which has the special property

that if $z \sim F^*$, then $z/c \mid z \geqslant c \sim F^*$, which implies that $\mathbb{E}_{F^*}[g(z/c)|z \geqslant c] = \mathbb{E}_{F^*}[g(z)]$ for a function $g(z)$. In particular,

$$\int_1^\infty m_c(z)dF^*(z) = \int_1^c dF^*(z) + \int_c^\infty m\left(\frac{z}{c}\right)dF^*(z) = F^*(c) + (1-F^*(c))\mathbb{E}_{F^*}\left[m\left(z\right)\right] = 1.$$

Show that a feasible model exists under F_c: We now show that the constraints hold under F_c. Define shifted parameters $f^c_{od} = c^{\theta^*} f_{od}$ and $A^c_{od} = c^{\frac{\theta^*}{\sigma-1}-1}A_{od}$, and keep all other parameters the same. We will use the following relationship between the conditional Pareto tails of F_c and F:

$$\mathbb{E}_{F_c}[z^{\sigma-1}1\{z > c\hat{z}\}] = \int_{c\hat{z}}^\infty z^{\sigma-1}m\left(\frac{z}{c}\right)dF^*(z) = c^{\sigma-1-\theta^*}\int_{\hat{z}}^\infty u^{\sigma-1}m(u)dF^*(u) = c^{\sigma-1-\theta^*}\mathbb{E}_F[z^{\sigma-1}1\{z > z\}].$$

Conjecture that the cutoff satisfies $\hat{z}^c_{od} = c\hat{z}_{od}$ and consider the price index:

$$\gamma^c_d = \sum_o \mathbb{E}_{F_c}[p^c_{od}(z)^{1-\sigma}1\{z > c\hat{z}_{od}\}] = \sum_o \mathbb{E}_F[p_{od}(z)^{1-\sigma}1\{z > \hat{z}_{od}\}] = \gamma_d. \tag{79}$$

Doing the same for the autarky price index yields $(\gamma^c_d)' = \gamma'_d$. Using this, it is then straightforward to verify the conjecture by substituting into the cutoff equation. Analogous steps establish that the trade flow constraints and aggregate variable profits are unchanged. Finally, the orthogonality restrictions continue to hold as c is a proportional shift. Gains from trade remains the same as γ_d, γ'_d are unchanged.

The shifted distribution has lower divergence: Let $S \subseteq \{1, 2, ..., D\}$ denote the set of origins whose draws z_o exceed c. Then, we write the event that this set S equals a particular set s as $A_{c,s} \equiv \{S = s\}$. The change of measure is then given by $M_c(\{z\}) = \prod_o m_c(z_o) = \prod_{o\in s} m\left(\frac{z_o}{c}\right)$, as mass below c has $m_c(z_o) = 1$. Divergence conditional on event $A_{c,s}$, i.e., $z_o \geqslant c$ for all $o \in S$, is

$$\mathbb{E}_{F^*}[\phi(M_c(\{z\})|A_{c,s}] = \mathbb{E}_{F^*}\left[\phi\left(\prod_{o\in s} m\left(z_o\right)\right)\right]. \tag{80}$$

Unconditional divergence is then the probability weighted sum of (80):

$$D(F_c||F^*) = \sum_{S\subseteq\{1,2,...,D\}} (c^{-\theta^*})^{|S|}(1-c^{-\theta^*})^{D-|S|}\mathbb{E}_{F^*}\left[\phi\left(\prod_{o\in S} m\left(z_o\right)\right)\right].$$

As $c \to \infty$, the RHS becomes arbitrarily small, which yields $D(F_c \mid\mid F^*) < \varepsilon$. Intuitively, under F^*, it is always possible to shift the economic activity further into the tail, lowering the divergence cost in proportion to the mass below the tail. □

A.6 Dimensionality reduction of the program (28)

Define the reduced program that retains only the domestic country's block:

$$\max_{F,\{A_{jd}\},\{f_{jd}\},f_{Ej},w_j,w'_j,\gamma'_j,N_j,N'_j} \quad 1 - \left(\frac{\gamma_j}{\gamma'_j}\right)^{\frac{1}{1-\sigma}} \tag{81}$$

$$\text{subject to:} \begin{cases} \gamma'_j = N'_j \mathbb{E}_F[p'_{jj}(\omega)^{1-\sigma}], \\ \sum_d \mathbb{E}_F\left[\max\{0, \frac{1}{\sigma} X_{jd}(\omega) - w_j f_{jd}\}\right] - w_j f_{Ej} = 0, \\ \mathbb{E}_F\left[\max\{0, \frac{1}{\sigma} X'_{jj}(\omega) - w'_j f_{jj}\}\right] - w'_j f_{Ej} = 0, \\ N_j \mathbb{E}_F[X_{jd}(\omega)] = \widetilde{X}_{jd} \text{ for all } d, \\ \hat{z}_{jd} \geqslant \hat{z}_{jj} \text{ for all } d \neq j, \\ (\hat{z}_{jj})' = 1, \\ \text{and } F \in \mathcal{N}_\delta, \end{cases}$$

Proposition 5 (Dimensionality reduction). *Suppose the baseline tariff data contain at least two origin countries, $i_1, i_2 \neq j$, where residualized tariffs, $\Delta_o \Delta_d \ln(1+t)$ contain off-diagonal cells of both signs. Assume that $\overline{GT}_j(\delta)$ and $\underline{GT}_j(\delta)$ are continuous in δ. Then, for every $\delta \in (0, \infty)$, the upper and lower bound values of the full program (31) and of the reduced program (81) coincide.*

Proof. We proceed in three steps. First, we show that, absent the orthogonality restriction, every feasible solution to (81) can be completed to a solution of the full program (28) with identical divergence. Second, we show that the two orthogonality conditions can be satisfied by choosing cutoffs appropriately. Third, we will show that the step in 3 can be done at arbitrarily low divergence cost.

Step 1: The reduced program implies a completed full program Let a feasible point of the reduced program (81) be given, with marginal F_j. Then, in absence of the orthogonality restriction (discussed below), note that we can always select

appropriate $\{A, f, f_E, N, N'\}_{-j}$ to satisfy all of the restrictions of the full program, (28), without affecting any restrictions of the reduced program (while γ_d for $d \neq j$ does appear in the free entry condition of the reduced program, it is without loss to normalize it to 1). Moreover,GT_j is the same, as it is fixed by the parameters γ'_j and γ_j. We call this a completed full program. We can construct this as follows. First, set the omitted foreign marginals to the Pareto benchmark. Second, choose admissible bilateral cutoffs. Third, for each od pair, invert the trade share restriction to obtain A_{od}. Fourth, define f_{od} to be the value that rationalizes the chosen cutoff and implied A_{od}. Fifth, set f_{Eo} equal to aggregate variable profits divided by w_o. Sixth, note that the price index restrictions automatically hold from summing bilateral expenditure, and market clearing automatically holds once trade shares are matched. The question is whether it is possible for the completed full program to satisfy the orthogonality restrictions, and what the divergence implications are.

Step 2: The orthogonality restriction can always be satisfied In Proposition 1(b), we showed that with Pareto marginals we can write both A_{od} and f_{od} as log-linear functions of the cutoffs, $\hat{z}_{od}$. We do the same here, except the A_{od} the function will not be log-linear. In particular, take the definition of the trade share in (1) and use the fact that the empirical gravity coefficient is θ^* to eliminate the $\Delta_o\Delta_d \ln \lambda_{od}\Delta_o\Delta_d \ln(1+t_{od})$ term to obtain

$$\sum_{o,d} \ln \mathbb{E}_{F_o}[z_o^{\sigma-1}1\{z > \hat{z}_{od}\}]\Delta_o\Delta_d \ln(1+t_{od}) = (\sigma - 1 - \theta^*)\sum_{o,d}(\Delta_o\Delta_d \ln(1+t_{od}))^2.$$

Doing the same for the cutoff definition, using (1) to eliminate the A_{od}, and using the previous expression to eliminate the expectation term, we get

$$\sum_{o,d} \Delta_o\Delta_d \ln \hat{z}_{od}\Delta_o\Delta_d \ln(1+t_{od}) = \sum_{o,d} \left(\Delta_o\Delta_d \ln(1+t_{od})\right)^2. \tag{82}$$

Combining these gives

$$\sum_{o,d} \left(\ln \mathbb{E}_{F_o}[z_o^{\sigma-1}1\{z > \hat{z}_{od}\}] - (\sigma - 1 - \theta^*)\Delta_o\Delta_d \ln \hat{z}_{od}\right) \Delta_o\Delta_d \ln(1+t_{od}) = 0, \tag{83}$$

which is 0 for any F_o that is $Pareto(1, \theta^*)$. As cutoffs can be freely chosen with an appropriate choice of $\{A\}$ and $\{f\}$, we will choose one country, i_1, to satisfy

(83) and another country, i_2, to satisfy (82). In particular, let i_1's distribution be a mixture: with probability $1-\varepsilon$ it draws from F^*, and with probability ε it draws from $Pareto(1,\theta)$ for $\theta \in (\sigma-1,\theta^*)$. The term in parentheses for $o = i_1$ evaluates to

$$\ln\left((1-\varepsilon)\frac{\theta^*}{\theta^*-(\sigma-1)}+\varepsilon\frac{\theta}{\theta-(\sigma-1)}\hat{z}_{od}^{\theta^*-\theta}\right). \tag{84}$$

This is strictly increasing in $\hat{z}_{od}$ and goes to ∞ as $\hat{z}_{od}$ goes to ∞. Thus, for any vector of parameters, we can always ensure (83) holds by making $\hat{z}_{od}$ sufficiently large for a d where $\Delta_o\Delta_d \ln(1+t_{od}) > 0$ (if the condition was previously violated in a negative direction) or for a d where $\Delta_o\Delta_d \ln(1+t_{od}) < 0$ (if the condition was previously violated in a positive direction). The data condition simply ensures there exists a country that has residualized tariff variation in the necessary direction in an off-diagonal column (to maintain selection into exporting).

For the other restriction, (82), consider another country i_2 with residualized tariff variation in off-diagonal rows, and let its marginal be $Pareto(1,\theta^*)$. Then, we can do the same exercise of letting $\hat{z}_{od}$ become arbitrarily large for a d that has a residualized tariff in the appropriate direction. Tthis does not affect (83) as the term in parentheses is zero if $F_o \sim Pareto(1,\theta^*)$. Thus, we can satisfy both restrictions.

Step 3: Divergence cost of satisfying the restrictions is vanishingly small Thus-far, we have the joint distribution F characterized by independent draws with marginals F_j given from the reduced program, $F_o = F^*$ for all $o \neq j, i_1$, and $F_o = (1-\varepsilon)F^* + \varepsilon Pareto(1,\theta)$ for $o = i_1$. A property of ϕ-divergence is that adding dimensions independently at the benchmark marginal does not affect divergence. Therefore, the only non-j element that could increase divergence is the ε component of i_1. Consider a given divergence target for the full program, δ. Then, consider the solution to the reduced program at divergence constraint $\delta-\eta$ for small $\eta > 0$, denoted $\overline{GT}_j^{\text{reduced}}(\delta-\eta)$. By continuity, by choosing ε sufficiently small we can produce a full program solution with the same gains from trade and with divergence $D(F||F^*) \in (\delta-\eta,\delta)$. Therefore, $\overline{GT}_j^{\text{reduced}}(\delta-\eta) \leqslant \overline{GT}_j^{\text{full}}(\delta) \leqslant \overline{GT}_j^{\text{reduced}}(\delta)$, where the second inequality holds because the reduced program is a relaxation of the full program. Taking the limit as $\eta \to 0$ gives $\overline{GT}_j^{\text{full}}(\delta) = \overline{GT}_j^{\text{reduced}}(\delta)$. □

B Computational Details

In this Appendix, we give further detail on the theoretical and numerical procedure used to implement the divergence bounds. Throughout the Appendix, we index the focal country (France) with j. Section B.1 provides more detail on inner loop that converts the infinite-dimensional optimization over distributions into a finite-dimensional convex program. Section B.2 provides further detail on the numerical implementation of the Ricardian model. Section B.3 provides further detail on the implementation of the monopolistically competitive model.

B.1 Further Detail on the Inner Program

This section provides an intuitive overview of how the inner loop converts the infinite-dimensional optimization over distributions into a finite-dimensional convex program. Christensen and Connault (2023) provide a formal treatment.

For a fixed vector of parameters ψ, let $g(z;\psi)$ denote the vector of moment functions, so the restrictions in (23) can be written as $\mathbb{E}_F[g(z;\psi)] = 0$. Since the objective in (23) is purely a function of ψ and thus has no direct dependence on F, it is convenient to consider an auxiliary problem of choosing F to minimize divergence between F^* and F subject to satisfying all of the restrictions evaluated at ψ (with the convention that divergence is ∞ if no such F exists). Denote the minimized divergence $\Delta^*(\psi)$. Then, the original problem can be written as $\max_\psi 1 - (\gamma_j'/\gamma_j)^{1/(\sigma-1)}$ subject to $\Delta^*(\psi) \leqslant \delta$. We now describe in more detail how to characterize $\Delta^*(\psi)$.

Recall that the definition of divergence in (30) is

$$D(F||F^*) \equiv \int \phi\left(\frac{dF(z)}{dF^*(z)}\right) dF^*(z), \tag{85}$$

where $\phi(x)$ is the chosen divergence penalty function described in Section 3.3. The primal of the auxiliary problem is then

$$\min_F \mathbb{E}_{F^*}\left[\phi\left(\frac{dF(z)}{dF^*(z)}\right)\right] \text{ subject to } \mathbb{E}_F[g(z;\psi)] = 0. \tag{86}$$

Since $dF(z)/dF^*(z)$ is a function mapping z onto $[0,\infty)$, it is convenient to simply write $m(z) \equiv dF(z)/dF^*(z)$. Then, notice that $\mathbb{E}_F[g(z;\psi)] = \mathbb{E}_{F^*}[m(z)g(z;\psi)]$, i.e.,

$m(z)$ can (loosely speaking) be interpreted as a weighting function that reweights the mass that F^* places on z such that any expectation under F can be expressed as a weighted expectation under F^*.[47] Searching over distributions F is then equivalent to directly searching over weights $m(z)$, where we require $m(z) \geqslant 0$ for all z and $\mathbb{E}_{F^*}[m(z)] = 1$ to ensure that the distributions implied by $m(z)$ are valid. We can thus express (86) entirely in terms of expectations under F^* as

$$\min_{m(z)} \mathbb{E}_{F^*}[\phi(m(z))] \tag{87}$$
$$\text{subject to: } \begin{cases} \mathbb{E}_{F^*}[m(z)g(z;\psi)] = 0 \\ \mathbb{E}_{F^*}[m(z)] = 1 \\ \text{and } m(z) \geqslant 0 \text{ for all } z. \end{cases}$$

Of course, $m(z)$ is still an infinite-dimensional object. However, as we show below, for suitable differentiable and strictly convex functions ϕ it is straightforward to solve for $m(z)$ pointwise as a closed-form function of the multipliers of the Lagrangean. In particular, the Lagrangean is

$$L(m(z), \lambda, \zeta) = \mathbb{E}_{F^*}[\phi(m(z)) + m(z)(\lambda' g(z;\psi) + \zeta)] - \zeta, \tag{88}$$

where λ are the multipliers on the constraints $g(z;\psi)$ and ζ is the multiplier on the adding-up constraint (the non-negativity constraint will always be slack for our choice of ϕ). Conditional on (λ, ζ), the FOC yields a solution for $m(z)$:

$$\hat{m}(z) = (\phi')^{-1}(-\lambda' g(z;\psi) - \zeta). \tag{89}$$

For the specific divergence function that we use, $\phi(x) = x \ln x - x + 1$ for $x \leqslant e$ and $\phi(x) = (x-e)^2/(2e) + (x-e) + 1$ for $x > e$, and writing $y \equiv -\lambda' g(z;\psi) - \zeta$ for convenience, this is $\hat{m}(z) = e^y$ for $y \leqslant 1$ and $\hat{m}(z) = ey$ for $y > 1$. Substituting this into (88) and doing some algebra yields the dual problem,

[47] To see this, note that, under appropriate regularity conditions, we can write

$$\mathbb{E}_F[g(z;\psi)] = \int g(z;\psi) dF(z) = \int g(z;\psi) \frac{dF(z)}{dF^*(z)} dF^*(z) = \int g(z;\psi) m(z) dF^*(z) = \mathbb{E}_{F^*}[m(z)g(z;\psi)].$$

$$\sup_{\lambda,\zeta} -\mathbb{E}_{F^*}[\phi^*(-\lambda' g(z;\psi) - \zeta)] - \zeta, \tag{90}$$

where $\phi^*(x) = e^x - 1$ for $x \leqslant 1$ and $\phi^*(x) = \frac{e}{2}(x^2+1) - 1$ for $x > 1$ is the convex conjugate of $\phi(x)$.[48] The program (90) is a convex optimization problem of dimension equal to the number of constraints plus one. Christensen and Connault (2023) establish the conditions under which the solution to the finite-dimensional dual in (90) is equal to the solution to the original infinite-dimensional primal in (86).

We implement (90) using quasi-Monte Carlo integration. That is, we simulate draws z_ω for $\omega = 1, 2, ..., W$ from F^* using scrambled Sobol sequences and solve

$$\sup_{\lambda,\zeta} -\frac{1}{W}\sum_{\omega=1}^{W} \phi^*(-\lambda' g(z_\omega;\psi) - \zeta) - \zeta. \tag{91}$$

B.2 Implementation Details for the Ricardian Model

This section provides more detail on the numerical program to calculate the divergence bounds for the Ricardian model. Our aim is to solve the problem in (23) with the additional constraint $F \in \mathcal{N}_\delta$. Define $\psi \equiv \{\{A\}, \{w\}, \{w'\}, \{\gamma\}, \{\gamma'\}\}$ as the parameter vector. Then, consider the following minimum divergence problem,

$$\Delta^*(\psi) \equiv \min_F D(F \mid\mid F^*) \quad \text{subject to all constraints in (23) evaluated at } \psi. \tag{92}$$

Christensen and Connault (2023) show that the dual to (92) is convex and of dimension equal to the number of constraints, and establish conditions under which the solution to the dual equals the solution to the primal. The full program in (23) that searches over both ψ and F is equivalent to $\max_\psi GT_j(\psi)$ subject to $\Delta^*(\psi) \leqslant \delta$.

Several observations allow us to reduce the dimensionality of ψ. First, the market clearing condition in the counterfactual economy can be omitted, as we do not need to know w_o' for any $o \neq j$. We can also normalize $w_j' = 1$. Second, substituting the trade flow moments into the baseline market clearing gives $w_o L_o = \sum_d \tilde{\lambda}_{od} w_d L_d$. This does not depend on ω, so it can be solved directly for $\{w_o\}$ outside of the

[48] The convex conjugate of a function $f(x)$ is defined as $f^*(x) = -(\inf_t f(t) - tx)$. Letting $t = m(z)$ and $x = -\lambda' g(z;\psi) - \zeta$, and writing the associated dual problem to (88) as $\sup_{\lambda,\zeta} \inf_{m(z)} \mathbb{E}_{F^*}[\phi(m(z)) - m(z)(-\lambda' g(z;\psi) - \zeta)] - \zeta$, note that the inf takes this form pointwise.

program. Third, we do not need γ'_d for any $d \neq j$ at any point, so we can omit these. Fourth, we can normalize $\gamma_d = 1$ for all d by choosing appropriate units of $\{A\}$. Fifth, as we do not include rows where $d = ROW$ in the orthogonality restriction, we do not directly need the corresponding A_{od} in this program. As, for any F, it is possible to find a vector $\{A\}$ that rationalizes the trade shares (see Lemma 1), we can omit the trade flow moments for $d = ROW$. We are therefore left with $\hat{\psi} \equiv \{\hat{\gamma}_j, \{A_{od}\}_{d \neq ROW}\}$, a total of 361 parameters, and the trade flow moments, j's autarky price index moment, and the orthogonality restriction.

We implement this numerically using various gradient-based solvers in Artelys Knitro, with expectations computed using 500,000 quasi-Monte Carlo draws constructed with scrambled Sobol sequences. The two key requirements for this to be computationally feasible are that the Δ^* inner optimization can be computed at low cost and that high quality gradients can be supplied to the outer solver. We address these in turn.

Following the notation of Section B.1, the inner Δ^* program is given by

$$\Delta^*(\psi) = \sup_{\lambda,\zeta} -\zeta - \frac{1}{W}\sum_{\omega=1}^{W} \phi^*(-\zeta - \lambda' g(z_\omega; \psi)), \tag{93}$$

where $\phi^*(x) = e^x - 1$ for $x \leqslant 1$ and $\phi^*(x) = \frac{e}{2}(x^2 + 1) - 1$ for $x > 1$ is the convex conjugate of $\phi(x)$. This program is smooth and convex, allowing for specialized algorithms to converge in relatively few iterations if supplied with the exact gradient and Hessian. Straightforward derivations show that these can be calculated in closed form. Moreover, they can be constructed efficiently by exploiting the sparse structure of $g(z_\omega, \psi)$ to substantially reduce the number of matrix operations required: perfect competition ensures that the trade share moment, $X_{od}(\omega)$, is positive for exactly one o in each d, and zero otherwise. We further accelerate computation of (93) by implementing several low-cost pre-screens to detect infeasible ψ (for which $\Delta^* = \infty$) without requiring numerical maximization, and maintaining a bank of cached prior solutions to use for warm starts.

As to the outer solver gradient, we calculate exact derivatives for the smooth CES expenditure changes on the intensive margin of incumbent origins and use finite differences to account for changes in origin induced by changes in A_{od}. Again, the sparsity of the trade share moments permits efficient computation, and the

separability across destinations allows for large gains from parallelization and high quality aggregation from individual coordinates to predicted step value. The Envelope Theorem ensures that we do not need to re-compute Δ^* to evaluate the outer gradient. We initialize the outer solver at randomized starting points and iteratively feed prior solutions until convergence.

B.3 Implementation Details for the Monopolistically Competitive Model

The program to be solved is (28), with the additional constraint $(\hat{z}_{jj})' = 1$ (discussed in footnote 28) and $F \in \mathcal{N}_\delta$. As discussed in Proposition 5, it is sufficient to consider a reduced program that only explicitly imposes trade shares and parameters corresponding to the focal country's exports. Note that we can omit the constraint $\hat{z}_{jj} \geqslant 1$ as the cutoff can never rise when moving to autarky, so it is already implied by the restriction $(\hat{z}_{jj})' = 1$.

First, as in the Ricardian model, we can normalize $w_j = w_j' = 1$ and $\gamma_d = 1$ for all d by measuring A_{od} in appropriate destination-specific units. Second, we can normalize $N_j = 1$ by measuring A_{jd}, f_{jd}, and f_{Ej} in appropriate origin-specific units.

Next, we can eliminate f_{Ej} by combining the baseline and autarky free-entry conditions. In particular, moving f_{Ej} to the RHS of the baseline and autarky conditions and then equating them gives

$$\mathbb{E}_F \left[\sum_d \max\{0, \frac{1}{\sigma} X_{jd}(\omega) - f_{jd}\} - \max\{0, \frac{1}{\sigma} X_{jj}'(\omega) - f_{jj}\} \right] = 0. \tag{94}$$

This imposes the same economic restriction but does not require us to explicitly specify f_{Eo}. By definition, the sales of a firm with productivity equal to the zero profit root satisfy $\hat{X}_{jj} = \sigma w_j f_{jj}$. The restriction $(\hat{z}_{jj})' = 1$ combined with the observation that $\hat{X}_{jj}'/\hat{X}_{jj} = 1$ implies $\gamma_j' = \hat{z}_{jj}^{1-\sigma}$, which allows us to express γ_j' as a function of $\hat{z}_{jj}$ and thus eliminate it from the search.

Given A_{od}, the definition of the cutoff provides a one-to-one mapping between f_{od} and $\hat{z}_{od}$. It is therefore equivalent, and generally preferable numerically, to search over $\hat{z}_{od}$ instead of f_{od}. In what follows, we write $f_{od}(\hat{z}_{od})$ to emphasize that f_{od} is a (log-linear) function of $\hat{z}_{od}$.

Finally, we can also omit the autarky price index moment as we can always

select N_j' ex post to satisfy it. Thus, define $\psi \equiv (\{A_{jd}\}, \{\hat{z}_{jd}\})$. Define the minimum divergence program as

$$\Delta^*(\psi) \equiv \min_F D(F \mid\mid F^*) \tag{95}$$

$$\text{subject to: } \begin{cases} \mathbb{E}_F[p_{jd}(\omega)^{1-\sigma} 1\{z_j(\omega) > \hat{z}_{jd}\} - \tilde{\lambda}_{jd}] = 0 \text{ for all } d, \\ \mathbb{E}_F\left[\sum_d \max\{0, \frac{1}{\sigma} X_{jd}(\omega) - f_{jd}(\hat{z}_{jd})\} - \max\{0, \frac{1}{\sigma} X_{jj}'(\omega) - f_{jj}(\hat{z}_{jj})\}\right] = 0. \end{cases}$$

with overall optimization program $\max_\psi \; 1 - 1/\hat{z}_{jj}$ subject to $\Delta^*(\psi) \leqslant \delta$ and $\hat{z}_{jd} \geqslant \hat{z}_{jj}$ for all d. This is a maximization program over $2D$ parameters. The computational details are similar to the Ricardian model, described in Section B.2, though in much lower dimension. Again, we implement this numerically using a combination of gradient-based solvers in Artelys Knitro. Two aspects of the problem are particularly convenient. First, the problem is smooth in A_{jd} and thus exact gradients can be provided. Second, the trade share matrix is both highly sparse and highly structured as, in a given od pair, $X_{od}(\omega) = 0$ for all ω below the cutoff. This implies that we can sort the trade shares moment indices and compute running sums ex ante, meaning that we need only locate the index of the cutoff crossing.

C Supplementary Analysis

Table C1: Upper Bound for Alternative Values of σ

	$\delta = 0$	$\delta = 1$		$\delta = \infty$
σ		Ricardian	Monopolistic Competition	
2	2.4%	9.5%	11.2%	16.5%
3	2.4%	7.9%	8.1%	8.6%
4	2.4%	5.6%	5.8%	5.9%

Notes: This table reports the value of the upper bound on gains from trade in the (unrestricted) Ricardian model and the Monopolistic Competition model for different values of the elasticity of substitution, σ. The values in the $\delta = 0$ column correspond to gains from trade under the benchmark distribution, F^*, which follow the formula in Arkolakis et al. (2012), $1 - \lambda_{dd}^{1/\theta^*}$. The values in the $\delta = \infty$ column correspond to the distribution-agnostic bounds derived in Proposition 1(a) and Proposition 1(b), $1 - \lambda_{dd}^{1/(\sigma-1)}$. The values in the $\delta = 1$ columns correspond to the bounds when F is restricted to lie within divergence neighborhood $\mathcal{N}_\delta$ where $\delta = 1$, where the $\sigma = 3$ values can be seen directly in Figure 3.