Quantifying the Sensitivity of Quantitative Trade Models*

Habib Ansari
Birkbeck

Dave Donaldson
MIT

Edward Wiles
Harvard

Abstract

The core ingredient of modern quantitative trade models is cross-country heterogeneity in productive capabilities across goods. In workhorse models this heterogeneity is assumed to follow specific functional forms—such as i.i.d. Frechet or Pareto—but little is known about the sensitivity of these models' counterfactual predictions to alternative distributional assumptions. In this paper we quantify this sensitivity by estimating sharp bounds on counterfactuals across all distributions within given neighborhoods of the workhorse model. Importantly, our method only considers distributions that can generate baseline equilibria that exactly match all trade flow data (and hence empirical facts such as the fit of the gravity equation or an estimated trade elasticity). Our main finding is that there exist distributions under which the gains from trade are four times larger than in the workhorse model and that even small departures from the workhorse distributional assumption can come close to these extremes. Imposing high-level restrictions on the distributions considered—for example, maintaining Frechet/Pareto marginals but relaxing independence, or vice versa—does shrink these bounds, but only by about 20%.

*This version: March 2026. Author contacts: mhansari@mit.edu, ddonald@mit.edu, edwardwiles@fas.harvard.edu. We are grateful to Arnaud Costinot and Federico Huneeus for helpful comments and to Noah Siderhurst for outstanding research assistance.

1 Introduction

A major advance in the field of international economics has 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] The core ingredient of these models 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 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 imports and doesn't produce?

To make progress, researchers have posited specific functional forms to describe production heterogeneity around the world. Two choices known as "gravity models"—that followed by Eaton and Kortum (2002) for competitive models and by Chaney (2008) for monopolistically competitive models—have become enormously influential for workhorse quantitative modeling. But how would trade economists' counterfactual predictions change if alternative forms of production heterogeneity were considered instead?

In this paper we conduct a comprehensive quantification of the sensitivity of quantitative trade models to departures from the gravity form. Our approach embodies two key constraints. First, in order to focus on the role of production heterogeneity, we hold constant all other model components, such as preferences, factor market operation, and market conduct. Second, we ensure that every bilateral trade flow—and hence also all derived properties of such flows, such as the fit of the gravity equation or the elasticity of trade with respect to tariffs—is matched perfectly by every variant we consider, just as they are by gravity models. Then, subject to these constraints, we estimate sharp bounds on counterfactual predictions that hold across *every* distribution of production heterogeneity within a fixed neighborhood of that used in gravity models. Letting this neighborhood grow provides a systematic exploration of model sensitivity.

We focus primarily (in this version) on one key counterfactual question: how much do countries gain from trading? That is, what would be the welfare cost of a move from the baseline trading equilibrium to a counterfactual scenario of complete autarky.

Our main finding is that workhorse quantitative models are highly sensitive to alternative distributions of production heterogeneity that are "close" to the gravity bench-

[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 surveyed in two recent chapters of the *Handbook of Regional and Urban Economics*, Allen and Arkolakis (2025) and Redding (2025).

mark. To interpret this result, begin first with the case in which one is completely agnostic about the distribution of production heterogeneity across countries and goods—that is, the neighborhood of distributions considered is infinitely large. We derive new analytical bounds on the gains from trade for this case and show that the bounds are sharp (i.e. there exists a distribution that can achieve them). They are also identical across the competitive and monopolistically competitive models. The lower bound is zero and the upper bound is given by a formula involving one data ingredient (the openness of the country in question) and one parameter (the elasticity of substitution across goods in consumer preferences).[2]

Evaluating this agnostic bounds formula for the case of France, a country that has a representative level of openness among those in our sample, we find that the upper bound on French gains from trade is about five times larger than the typical gravity model benchmark. Complete agnosticism about production heterogeneity would clearly lead to wide bounds on predictions about this important object.

We therefore turn away from infinitely large neighborhoods of distributions and consider only those that are "close" to the gravity model benchmark (the Frechet distribution in Eaton and Kortum (2002) and the Pareto distribution in Chaney (2008)). To evaluate what is "close", we use the concept of divergence, a standard way of measuring the proximity between two probability distributions. We then consider a level of divergence that is smaller than what is typically considered in conventional robustness checks that are routinely performed within the use of gravity models (such as evaluating a limited range of values for the trade elasticity). When evaluating all possible distributions that are within a neighborhood constrained by this amount of divergence we find that the minimal gains from trade for France are again zero, and the maximal gains are about four times the gravity benchmark. That is, even when considering only "small" neighborhoods of alternative distributions, the gains from trade get very close to their agnostic limits of somewhere between zero and five times the gravity benchmark.

Evaluating the infinitely large set of candidate distributions within a neighborhood is of course impossible, and matters are harder still when solving the global equilibrium at each candidate and retaining only those that are consistent with trade data. To circumvent this challenge we draw on advances in the field of distribution-robust econometrics. In particular, Christensen and Connault (2023) consider an optimization problem over all distributions of a random variable, where the objectives and constraints are expectations

[2]This upper bound formula bears some resemblance to that used by Arkolakis et al. (2012) to answer a different question, which is what is the exact (not upper bound on the) value of the gains from trade in the case of one particular distribution of heterogeneity, the gravity model. But the parameter involved relates to the shape of this distribution, not preferences.

of (known functions of) the random variable. They show that this infinite-dimensional program has a dual that is finite-dimensional, with dimension given by the number of constraints in the original problem. Following this logic, we write our problem—maximize gains from trade subject to the constraints of model equilibrium and matching all trade data—as one with an inner loop (given all unknown model parameters and equilibrium variables in the factual and counterfactual equilibria) and an outer loop (that searches over these parameters and variables). The inner loop is infinite-dimensional but benefits from the finite-dimensional dual formulation in Christensen and Connault (2023), and the outer loop is finite-dimensional. Put 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 seek to restrict the admissible distributions further. Our final set of results do this by imposing high-level restrictions. To guide this, consider three core features of the benchmark distribution: (i) countries take draws independently from one another; (ii) countries all have the same marginal distribution; and (iii) the common marginal distribution takes a particular form, namely 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—to incorporate 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. For example, a simple subset of (i) is zero correlation, which can be imposed by adding constraints to ensure that the expectation of the product of any pair of countries' draws is equal to the product of the countries' expectations.

Across a wide range of restrictions imposed, the basic message is that gains from trade bounds still remain wide. In fact, it turns out that for the case of the monopolistically competitive model the imposition of restrictions (i) and (ii) has no effect on our bounds. In this model, every firm in the world produces a unique product. When France goes autarkic its consumers will lose access to a given mass of products, but the value of that mass is pinned down by the trade flow data points we match at baseline when considering any distribution, restricted or not. What does then uniquely matter to French consumers in autarky is their country's own marginal distribution. If we impose (iii) then this corresponds to the gravity model and the gains from trade are known. If we leave this marginal to be flexible (within a given neighborhood of Pareto) the gains from trade can be many times larger or smaller, as discussed above.

The competitive model is qualitatively different. Here, competition is head-to-head.

In autarky, French consumers will therefore replace the draws they received for any imported goods with draws from France's marginal distribution. Since the distribution of imported draws depends on the joint properties of the foreign and French distribution, the restrictions (i) and (ii) do have bite—at least for the upper bound. Perhaps surprisingly, we find that the reduction in the upper bound (at fixed neighborhood size) is always about one fifth: imposing independence is 18%, zero correlation is 18%, common (but flexible marginals) is 20%, common (and Frechet) marginals is 20%, and the joint imposition of independence and common (flexible) marginals is 23%. That is, even if countries draw productivities independently from one another like in Eaton and Kortum (2002), and even if they all draw from the same distribution like in Eaton and Kortum (2002), there exist versions of this distribution that, when fed into the standard Ricardian trade model, continue to match all available data and yet have gains from trade that can be as low as zero or as high as 3.5 times the value in Eaton and Kortum (2002).

We believe these results are surprising for two reasons. First, one might conjecture that once key facts about trade flows are matched—facts such as the extent of openness, the impressive fit of the gravity equation, and the elasticity of trade with respect to tariffs—then there may be little room to maneuver in search of alternative models with distributions of production heterogeneity and yet considerably different gains from trade. To the contrary, our findings demonstrate that the empirical fit of the gravity equation has little connection to the extent to which any given country's trade flows will respond to tariff restrictions, and hence about the shape of that country's import demand curve (the area under which quantifies the welfare losses from an inability to trade).

Second, one might conjecture that the predictions of the competitive model may inherit robustness from the extremal types theorem. This result is analogous to the central limit theorem but for extremal statistics such as the price of the cheapest global supplier of any good, which is what matters for welfare in the competitive model. Our findings imply that the relevant rate of convergence in the extremal types theorem in our context must be "slow" in the economic sense that it does not appear to provide much robustness for the problem that interests us here (i.e. quantifying the gains from trade in the case of 20 countries). We are unaware of attempts to quantify this element of potential robustness in settings with a larger number of spatial units.

2 Quantitative trade model environments

In this section we describe the core economic environment of 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) and then a monopolistically competitive version.

2.1 A competitive model of trade

We begin with a standard Ricardian model of international 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 which is immobile across countries.

Demand. Utility in each country d takes the CES form across the continuum of goods:

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

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. $\sigma > 1$).[3]

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 the 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 $\gamma_d \equiv \sum_o \int_0^1 p_{od}(\omega)^{1-\sigma} d\omega$ 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 efficiency $z_o(\omega)$ for producing

[3] This restriction is needed only 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.

any good ω. The marginal cost of a producer of ω in o when selling in country d is given by

$$MC_{od}(\omega) \equiv \frac{w_o(1+t_{od})}{A_{od}z_o(\omega)}, \tag{2}$$

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.[4] The efficiency level $z_o(\omega)$ augments production by shifting downwards the marginal cost of producing ω in o.

Because the good ω is homogeneous and the Ricardian model features perfect competition, 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

$$p_{od}(\omega) = \begin{cases} MC_{od}(\omega) & \text{if } o = \arg\min_{o' \in D} MC_{o'd}(\omega), \\ \infty & \text{otherwise,} \end{cases} \tag{3}$$

with the convention that country d only buys from the minimum-price supplier of any good (since $p_{od}(\omega) = \infty$ 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 $w_o L_o$. We assume that tariff revenue is spent on projects that do not directly contribute to utility.[5] To account for trade imbalances, we define country d's trade deficit as the difference between expenditure and income $B_d \equiv E_d - w_d L_d$. A competitive global equilibrium in this Ricardian model is then one in which the factor market clears,

$$w_o L_o = \sum_d \int_0^1 \frac{1}{1+t_{od}} X_{od}(\omega) d\omega \tag{4}$$

in every country o.

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

[5] This assumption simplifies the formulas, which makes our core focus—the role of unobserved heterogeneity in productivity—clearer. In practice, the share of tariff revenue in expenditure is negligible for our reference country (France).

2.2 Stochastic representation

The key ingredient of this Ricardian model is the joint distribution of countries' efficiencies, $z_o(\omega)$ for all $o \in D$. For example, if efficiencies were to be distributed in such a way that they are perfectly correlated across countries (i.e. $z_o(\omega) = \alpha z_{o'}(\omega)$ for some $\alpha > 0$ for any pair of countries o and o') 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 convention in Eaton and Kortum (2002) and treat the efficiency $z_o(\omega)$ as a random variable that is drawn from a multi-country stochastic process with *arbitrary* cumulative distribution function F. 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 better draws than country o') and that countries' draws display forms of dependence (e.g. that the draws of country o are correlated with those of country o').

Using this stochastic representation we follow the law-of-large-numbers convention of using the expectation operator $\mathbb{E}_F[\cdot]$ to replace integrals across the continuum of goods. That is, a model equilibrium is defined by

$$\gamma_d = \sum_o \mathbb{E}_F\left[p_{od}(\omega)^{1-\sigma}\right], \tag{5}$$

$$w_o L_o = \sum_d \frac{1}{1+t_{od}} \mathbb{E}_F\left[X_{od}(\omega)\right], \tag{6}$$

alongside the definition of the random variables $p_{od}(\omega)$ from (3) and $X_{od}(\omega)$ from (1). For future reference, we let the total value of bilateral trade be denoted $X_{od} \equiv \mathbb{E}_F[X_{od}(\omega)]$.

2.3 Counterfactuals

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 all other exogenous model features—endowments $\{L\}$, bilateral productivity shifters $\{A\}$, preferences σ, the deficit to GDP ratio $\{B/Y\}$, and the multi-country efficiency distribution F—constant.

The counterfactual equilibrium is one in which all endogenous variables have changed

correspondingly (e.g. the factor price in country o has changed from w_o to w'_o). This means that counterfactual marginal costs become $MC'_{od}(\omega) = w'_o(1+t'_{od})A_{od}^{-1}[z_o(\omega)]^{-1}$, counterfactual prices $p'_{od}(\omega)$ follow from (3) but evaluated at $MC'_{od}(\omega)$, and equilibrium holds at new goods and factor prices:

$$\begin{aligned} X'_{od}(\omega) &= (w'_d L_d + B'_d)(\gamma'_d)^{-1} p'_{od}(\omega)^{1-\sigma}, & (7) \\ \gamma'_d &= \sum_o \mathbb{E}_F \left[p'_{od}(\omega)^{1-\sigma} \right], & (8) \\ w'_o L_o &= \sum_d \frac{1}{1+t'_{od}} \mathbb{E}_F \left[X'_{od}(\omega) \right]. & (9) \end{aligned}$$

Gains from trade. A key counterfactual exercise of interest concerns the gains from trade for country d. This is calculated by changing tariffs from $\{t_d\}$ to the tariffs that would place d in autarky (i.e. $t_{od} = \infty$ for all $o \neq d$) and then solving for the change in real income (relative to baseline). That is, we define

$$GT_d \equiv 1 - \left(\frac{w'_d L_d + B'_d}{w_d L_d + B_d} \right) \left(\frac{\gamma_d}{\gamma'_d} \right)^{\frac{1}{1-\sigma}}. \tag{10}$$

Given the numeraire choice of $w_d = w'_d = 1$, as well as the specification that the deficit is fixed as a share of GDP, this further simplifies to

$$GT_d = 1 - \left(\frac{\gamma_d}{\gamma'_d} \right)^{\frac{1}{1-\sigma}}. \tag{11}$$

Additional sets of counterfactual calculations are just as feasible.[6]

Analytical bounds on the gains from trade. Our quantitative application will explore how sensitive gains from trade are to the choice of distribution, F. It is possible, however, to derive analytically the maximum and minimum gains from trade that can obtain.

Proposition 1. *Across all distributions F, gains from trade satisfy $0 \leq GT_d \leq 1 - \lambda_{dd}^{\frac{1}{\sigma-1}}$. Moreover, there exist F^- and F^+ that attain the lower and upper bound, respectively.*

Proof. Consider first the upper bound. By definition, the own-trade share at baseline satisfies $\lambda_{dd} = \gamma_d^{-1} \mathbb{E}_F[MC_{dd}(\omega)^{1-\sigma} 1\{d = \arg\min_o MC_{od}(\omega)\}]$. The equivalent expression in the counterfactual is $1 = (\lambda'_d)^{-1} \mathbb{E}_F[MC'_{dd}(\omega)^{1-\sigma}]$, as $t_{od} = \infty$ for $o \neq d$ implies

[6] In particular, for any known function $k(\cdot)$, the methods described below can be applied to any counterfactual calculation that takes the form $\kappa = \mathbb{E}_F[k(\{z\}\{w, w'\}, \{\gamma, \gamma'\}, \{t, t'\}, \{L, L'\}, \{A\}, \sigma)]$.

that all goods are sourced domestically. Substituting these expressions into (11), using $w_d = w'_d = 1$, and simplifying gives

$$GT_d = 1 - \lambda_{dd}^{\frac{1}{\sigma-1}} \left(\frac{\mathbb{E}_F[z_d(\omega)^{\sigma-1}]}{\mathbb{E}_F[z_d(\omega)^{\sigma-1} 1\{d = \arg\min_o MC_{od}(\omega)\}]} \right)^{\frac{1}{\sigma-1}}. \tag{12}$$

The expression in parentheses is minimized at 1, which occurs when $z_d(\omega)^{\sigma-1} = 0$ for all goods that d imports at baseline, which means that GT_d is maximized at $1 - \lambda_{dd}^{\frac{1}{\sigma-1}}$. We can construct a distribution that achieves this upper bound as follows. Suppose that all countries $o \neq d$ have independent draws from identical Frechet distributions, as in Eaton and Kortum (2002). Then, suppose that country d's draws are given by the following process. First, some $\tilde{z}_d(\omega)$ is drawn independently from the same Frechet. Then, if $\widetilde{MC}_{dd}(\omega) < \min_{o \neq d} MC_{od}(\omega)$, set $z_d(\omega) = \tilde{z}_d(\omega)$; otherwise, set $z_d(\omega) = 0$.

For the lower bound, the definition in (5) and perfect competition together gives $\gamma_d = \mathbb{E}_F[(\min_o MC_{od}(\omega))^{1-\sigma}]$. Again, noting that $MC'_{dd}(\omega) = MC_{dd}(\omega)$, we can substitute into equation (11) to obtain

$$GT_d = 1 - \left(\frac{\mathbb{E}_F[MC_{dd}(\omega)^{\sigma-1}]}{\mathbb{E}_F[(\min_o MC_{od}(\omega))^{\sigma-1}]} \right)^{\frac{1}{\sigma-1}}. \tag{13}$$

The expression in parentheses is maximized at 1, which occurs when d's marginal cost at baseline is exactly equal to the marginal cost of imported goods (gross of tariffs), which means that GT_d is minimized at 0. We can construct a distribution that obtains this lower bound as follows. Let $z_o(\omega)$ for $o \neq d$ be drawn from any F that satisfies baseline trade shares between all non-d pairs, and then $z_d(\omega) = \alpha_{o(\omega)} z_{o(\omega)}(\omega)$, where $o(\omega) \equiv \arg\min_{o' \neq d} MC_{od}(\omega)$ and $\alpha_{o(\omega)} = \frac{A_{o(\omega)d} w_d (1+t_{dd})}{A_{dd} w_{o(\omega)} (1+t_{o(\omega)d})}$. □

Intuitively, gains from trade in the competitive Ricardian model are ultimately given by the difference between the price index when some goods are imported and the price index if those goods were instead produced domestically (using more costly techniques). It is maximized if d is incapable of producing these goods domestically ($MC_{dd}(\omega) = \infty$), in which case the price index aggregator shrinks to a share λ_{dd}, which translates into a welfare loss according to the elasticity $1/(1-\sigma)$. Indeed, this is the same formula that obtains under an Armington model (though with a different structural interpretation of σ) which occurs because in the Armington model country d, by definition, cannot produce domestically the goods that other countries produce. In the opposite direction, gains from trade are minimized if d can produce the currently imported goods at the same tariff-

inclusive price, in which case trade at baseline simply reflects consumer indifference.

2.4 The Eaton and Kortum (2002) Ricardian model

Our presentation of the Ricardian model so far has held for any given distribution of global efficiencies F. However, one particularly tractable version of F, proposed in the seminal model of Eaton and Kortum (2002), is that where countries' efficiency levels are drawn independently from a Frechet distribution with scale parameter equal to one and shape parameter equal to θ^*, i.e. $z_o(\omega) \sim Frechet(1, \theta^*)$ for all o. We refer to this distribution as F^* in what follows and treat it as a key reference point with which to compare the Ricardian model based on other distributions F.

It is worth emphasizing 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 will be 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 (2), on the basis of efficiency draws $z_o(\omega)$ from an *arbitrary* global distribution that we continue to denote by F.

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 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. 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.[7]

[7] As is standard, we assume that, in equilibrium, there is only one producer per good ω in any destination market d.

Collecting these observations and applying them to the case of CES demand implies that

$$p_{od}(\omega) = \begin{cases} \frac{\sigma}{\sigma-1} MC_{od}(\omega) & \text{if } X_{od}(\omega) > \sigma f_{od} w_o, \\ \infty & \text{otherwise.} \end{cases} \tag{14}$$

The inequality $X_{od}(\omega) > \sigma f_{od} w_o$ is equivalent to a condition that productivity exceeds a cutoff, i.e., $z_o(\omega) > z^*_{od}$. We impose two standard restrictions. First, we assume that parameters are such that $z^*_{od} \geq z^*_{dd}$ for all $o \neq d$, which ensures selection into exporting. Second, we assume that all cutoffs are interior—i.e., for all od, there exists some mass of goods for which $z_o(\omega) > z^*_{od}$ and some mass of goods for which $z_o(\omega) < z^*_{od}$.[8]

The second major difference is that the equilibrium mass N_o of potential producers in o who choose to take an efficiency 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 efficiency will be equal to the cost $w_o f_{Eo}$. That is

$$\sum_d \mathbb{E}_F \left[\max\{0, \frac{1}{\sigma} X_{od}(\omega) - w_o f_{od}\} \right] = w_o f_{Eo}. \tag{15}$$

We restrict attention to economies in which $N_o = N'_o$, i.e., where the counterfactual of interest does not change the mass of entrants.[9] Since profits are zero, household income in o continues to be given by $Y_o = w_o L_o$.

The market clearing and price index expressions of (5) and (6) are similar to the competitive model, except that they now include N_o,

$$\gamma_d = \sum_o N_o \mathbb{E}_F \left[p_{od}(\omega)^{1-\sigma} \right], \tag{16}$$

$$w_o L_o = \sum_d \left(\frac{\sigma - 1}{\sigma} \frac{1}{1 + t_{od}} + \frac{1}{\sigma} \right) N_o \mathbb{E}_F \left[X_{od}(\omega) \right]. \tag{17}$$

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 $\{w\}$ and masses $\{N\}$ of potential producers such that equations (1), (2) and (14)-(17) hold for all countries.

[8] A sufficient condition for at least some firms to be active in each od pair is that the upper bound of the support of F is ∞, which holds in most applications.

[9] In the Pareto case, which is focal below, this occurs because $N_o = \frac{\sigma-1}{\sigma\theta} \frac{L_o}{f_{E_o}}$ always.

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 by 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 efficiencies F—and solve for the new equilibrium at counterfactual $\{t'\}$ and where all endogenous variables 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.

Analytical bounds on the gains from trade. As with the case of the competitive model, it is possible to derive analytical bounds on the gains from trade for any country d that would be possible across every distribution F. Indeed, as shown in Proposition 2 in Appendix A, these bounds turn out to be identical to those of the competitive model; that is, gains from trade must satisfy $0 \leq GT_d \leq 1 - \lambda_{dd}^{\frac{1}{\sigma-1}}$.

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 efficiency 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 efficiency $z_o(\omega)$ independently from a Pareto distribution with scale parameter equal to one and a shape parameter equal to θ^*. 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 reference distribution F^* imposes three restrictions types of restrictions: full independence, common marginals, and particular (i.e. Pareto) 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 technology shifters $\{A\}$), plus the additional parameters 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 ideas 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 Ricardian 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. 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 the same for both the Eaton and Kortum (2002) and Chaney (2008) models. We therefore discuss the general approach that applies to both of these models simultaneously here.

The first step is to estimate the single unknown parameter inside F^*, referred to in both models as θ^*. This involves estimating the slope—the so-called trade elasticity—in the gravity equation. Gravity models predict that total 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}. A standard approach to estimating the parameter θ^* therefore involves using data on observed bilateral trade flows, which we denote by $\widetilde{X}_{od}$, and ad-valorem bilateral tariffs in order to estimate equation (18) by OLS, typically with the inclusion of origin- and destination-country fixed effects as controls. This procedure effectively invokes the assumption that $\ln(1 + t_{od})$ is mean independent of the unobserved error ε_{od} and then uses this assumption, plus the model-derived gravity form, to estimate the coefficient on $\ln(1 + t_{od})$, which corresponds to $-\theta^*$.

The second step is to note that, once an unbiased estimate of θ^* has been obtained, each bilateral residual in equation (18)—i.e. $\ln \widetilde{X}_{od} + \theta^* \ln(1 + t_{od})$—provides an unbiased estimate of ε_{od} up to two irrelevant o-specific and d-specific scale factors. These estimated values of ε_{od} then provide unbiased estimates of all of the remaining exogenous variables

in the model that are required for counterfactual simulations (such as computing GT_d).[10] Notably, this step involves choosing values of model parameters such that the estimated model perfectly matches all of the D^2 bilateral values of observed trade, $\widetilde{X}_{od}$. That is, this step ensures

$$\mathbb{E}_{F^*}[X_{od}(\omega)] = \widetilde{X}_{od} \text{ for all } o \text{ and } d. \tag{19}$$

Finally, the third step is to solve for GT_d, as defined in equation (11), according to the estimated model.[11] In the case of the Eaton and Kortum (2002) version of the gravity model, this amounts to solving the system of equations (5) and (6) for the factual equilibrium and (8) and (9) for the counterfactual one, each time with the expectations $\mathbb{E}_F[\cdot]$ evaluated at $F = F^*$ (and F^* being independent draws from a $Frechet(1, \theta^*)$ distribution at the estimated value of θ^*). Similarly, in the Chaney (2008) case, step 3 requires solving equations (16) and (17) 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.[12]

3.2 Beyond gravity models

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 efficiency 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?

[10] For example, in the Eaton and Kortum (2002) case, the fixed effects are $\alpha_o = -\theta^* \ln w_o$, $\alpha_d = \ln E_d - \ln\left(\sum_l A_{ld}(w_l(1+t_{ld}))^{-\theta^*}\right)$, and the residuals are $\varepsilon_{od} = \theta^* \ln(A_{od})$. And in the Chaney (2008) case, $\alpha_o = \ln N_o + \left(1 - \frac{\sigma}{\sigma-1}\theta^*\right) \ln w_o$, $\alpha_d = \ln E_d - \ln\left(\sum_l N_l w_l^{1-\frac{\sigma\theta^*}{\sigma-1}} A_{ld}^{\theta^*} f_{ld}^{1-\frac{\theta^*}{\sigma-1}} (1+t_{ld})^{-\theta^*}\right)$, and the residuals are $\varepsilon_{od} = \theta^* \ln A_{od} + \left(1 - \frac{\theta^*}{\sigma-1}\right) \ln f_{od}$. The bilateral technology shifters, A_{od}, and fixed costs, f_{od}, are not separately identified from aggregate trade flows in the Chaney (2008) case. However, such a separation is not needed to compute GT_d.

[11] The second and third steps of this procedure are often combined into an algebraically equivalent single step via the "exact-hat" algebra procedure pioneered by Dekle et al. (2007).

[12] Strikingly, the parameter σ is irrelevant for the calculation of GT_d in either the Eaton and Kortum (2002) or the Chaney (2008) model. But this is not true for the more general case of $F \neq F^*$.

One direction would be to repeat the analogs of the above steps for particular alternative distributions. For example, one could imagine working with the Ricardian case and positing that the global distribution F is independent draws from a log-normal marginal distribution, with location parameter zero and scale parameter s. 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 efficiency distribution. While we are unaware of an implementation of this example in prior work, examples of GT_d being estimated for distributions beyond the Eaton and Kortum (2002) and Chaney (2008) special cases do exist.[13]

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. To circumvent this concern we remain agnostic about the choice of F. Instead, we develop a procedure for estimating sharp bounds on GT_d that hold for *every* distribution F that lies within a given "neighborhood"of the reference distribution F^*. This permits a systematic investigation of the question posed above: how would GT_d change at alternative values of $F \neq F^*$?

Neighborhoods of the F^* distribution To begin, we use a standard method for comparing two distributions from the statistics literature. In particular, the "divergence" $D(F||F^*) \geq 0$ of F from F^* is given by

$$D(F||F^*) \equiv \int_{x \in \mathbb{R}_D^+} \phi\left(\frac{f(x)}{f^*(x)}\right) f^*(x) dx, \tag{20}$$

for a suitable choice of $\phi(\cdot)$, and where $f(\cdot)$ denotes the PDF of $F(\cdot)$, etc.[14] For our purposes, an attractive feature of divergence is that there is zero divergence between F^* and any F that can be achieved by a change in units of the random variable $z_o(\omega)$. This means that our choice to model F^* with a common location parameter, equal to one, in all countries does not affect any of our conclusions.

Armed with this definition of divergence, we can define a neighborhood $\mathcal{N}_\delta$ of F^* as all distributions $F \in \mathcal{N}_\delta$ for which $D(F||F^*) \leq \delta$. An additional attractive feature of

[13] For example, Lind and Ramondo (2023) investigate versions of the Ricardian model with distributions that feature cross-country dependence.

[14] For example, a common choice of $\phi(\cdot)$ is the Kullback-Leibler form, $\phi(x) \equiv x \ln x - x + 1$.

divergence, for our purposes, is that as $\delta \to \infty$, all distributions F that are absolutely continuous with respect to F^* (i.e., that can be obtained by reweighting the baseline distribution within its support) will be contained inside neighborhood $\mathcal{N}_\delta$. Since the support of F^* is the positive reals for Eaton and Kortum (2002) and the positive reals greater than an arbitrary non-negative value (chosen to be one) for Chaney (2008), this means that essentially all continuous distributions are considered in $\mathcal{N}_\delta$ as $\delta \to \infty$.

Bounds on GT_d in the Ricardian case. We now seek to estimate sharp bounds on the size of country d's gains from trade GT_d that are consistent with the standard three-step estimation, but done for all efficiency distributions F within the neighborhood $\mathcal{N}_\delta$ of the F^* assumed by Eaton and Kortum (2002).

To do so, we proceed as follows. Consider the upper bound on GT_d within $\mathcal{N}_\delta$, which we denote by $\overline{GT}_d(\delta)$. Mathematically, this is given by

$$\overline{GT}_d(\delta) \equiv \max_{F \in \mathcal{N}_\delta, \{A\}, \{w\}, \{w'\}} GT_d \tag{21}$$

$$\text{subject to: } \begin{cases} \text{economy at } \{t\} \text{ is in equilibrium (i.e. equations 1-6 hold),} \\ \text{economy at } \{t'\} \text{ is in equilibrium (i.e. equations 8-9 hold),} \\ \text{and } \mathbb{E}_F[X_{od}(\omega)] = \widetilde{X}_{od} \text{ for all } o \text{ and } d. \end{cases}$$

That is, estimating $\overline{GT}_d(\delta)$ amounts to searching over all admissible distributions $F \in \mathcal{N}_\delta$ and bilateral productivity shifters $\{A\}$, while solving the resulting model at the initial $\{w\}$ and counterfactual $\{w'\}$ equilibria each time, in order to find the version of the Ricardian model that maximizes the welfare drop for country d when it moves from the initial equilibrium to autarky, always while ensuring that all trade flows in the initial equilibrium match those in the data. Clearly the value of $\overline{GT}_d(\delta)$ is sharp—that is, can always be achieved by some feasible choice—by construction. The lower-bound, which we denote by $\underline{GT}_d(\delta)$, is analogous but uses the min operator in (21).

While the program in (21) is easy 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. Fortunately, however, Christensen and Connault (2023) have shown that a particular sub-problem of (21) has a dual that is feasible. 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) estab-

lish 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^*.

The program in (21) 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 (21) directly involves model parameters such as the productivity shifters, $\{A\}$. Similarly, the factual and counterfactual wages, $\{w\}$ and $\{w'\}$, enter the objective function (11) and constraints in nonlinear ways that disrupt the pure expectations form. Indeed, the same is true for the price indices, $\{\gamma\}$ and $\{\gamma'\}$. 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 (21) 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.

We execute this nested two-step procedure below. But the large dimensionality of our problem—with a dimensionality of $9D + D^2$ in the inner loop and $D^2 + 4D$ in the outer loop—requires us to utilize a number of observations about dimensionality-reduction that apply to our problem. This is discussed in Section 3.3.

Bounds on GT_d in the Melitz (2003) case. We have seen how to compute bounds on the gains from trade GT_d in the case of Ricardian models. These bounds refer to all Ricardian models with efficiency distributions F that are within a given neighborhood $\mathcal{N}_\delta$ of the benchmark gravity model distribution F^* used by Eaton and Kortum (2002)—namely that countries draw their efficiencies $z_o(\omega)$ independently from the $Frechet(1, \theta^*)$ distribution.

We turn now to the case of the monopolistic competition model with heterogeneous producers, a la Melitz (2003).[15] In this case we seek to bound GT_d across all versions of the Melitz (2003) model, with efficiency distributions F within $\mathcal{N}_\delta$ of the $Pareto(1, \theta^*)$ distribution F^* used in the benchmark gravity model case of Chaney (2008). Yet, as above we do so while replicating the standard three-step estimation procedure, which ensures that

[15] We impose one further restriction, namely, that all firms produce in autarky—i.e., $(z^*_{dd})' = 1$. This is because the aggregate moments can be matched with either high fixed costs and a small mass of very productive firms, or low fixed costs and a large mass of unproductive firms, the choice of which affects the interpretation of the divergence metric, δ. In particular, by making fixed costs very high, the program can shrink the relevant support of F and thus decrease the cost of changing the distribution. We can impose a consistent interpretation by fixing the share of firms that produce, and we thus set $(z^*_{dd})' = 1$ so that the interpretation of δ is similar to the Ricardian model, in which all domestic firms produce in autarky.

all bilateral trade flow data points are matched by the model at any of the distributions F under consideration. Mathematically, this program is similar to that in (21). It can be written as

$$\overline{GT}_d(\delta) \equiv \max_{F \in \mathcal{N}_\delta, \{A\}, \{f\}, \{f_E\}, \{w\}, \{w'\}, \{N\}} GT_d \tag{22}$$

$$\text{subject to: } \begin{cases} \text{economy at } \{t\} \text{ is in equilibrium (i.e. equations 1, 2 and 14-17 hold),} \\ \text{economy at } \{t'\} \text{ is in equilibrium,} \\ \text{and } N_o \mathbb{E}_F[X_{od}(\omega)] = \widetilde{X}_{od} \text{ for all } o \text{ and } d. \end{cases}$$

Relative to the Ricardian model, the new considerations of the monopolistically competitive model naturally require an augmented set of model equations. However, it also requires a larger set of model parameters (i.e. the addition of bilateral fixed costs of production $\{f\}$ and entry costs $\{f_E\}$). In addition, the equilibrium number of potentially producing firms $\{N\}$ must be solved for at every candidate distribution F under consideration alongside the market-clearing wages $\{w\}$ and $\{w'\}$.

We solve the program (22), and its lower-bound equivalent, using the same inner- and outer-loop method described for the Ricardian case. However, the outer loop over parameters ψ is now substantially higher-dimensional due to the addition of D^2 additional bilateral fixed costs $\{f\}$, D additional entry costs $\{f_E\}$, and D additional firm masses $\{N\}$ to the vector ψ. The dimensionality-reduction techniques we discuss next are especially important in this case.

3.3 Implementation Details

We implement this two-step optimization program using Knitro, where an inner loop solves the search over distributions for a given vector of parameters, and an outer loop searches over the vector of parameters. We compute expectations using Monte Carlo integration with 800,000 draws.

In both versions of the model, there are a large number of parameters in the outer loop (in the order of at least D^2), which would be very challenging for any solver. However, we can substantially reduce the dimensionality by exploiting the fact that many of these parameters are irrelevant for computing gains from trade for d. For the Ricardian case, the parameters $A_{od'}$ for $d' \neq d$ do not appear in any equations relevant for determining GT_d, which means they can be omitted.[16] The only exception is that they do appear in

[16] We also then omit the corresponding trade share constraints, $\widetilde{X}_{od} = \mathbb{E}_F[X_{od'}(\omega)]$ for $d' \neq d$. This does not affect the model's ability to replicate the baseline equilibrium, as for any combination of F and the

the moment restriction corresponding to the gravity equation, discussed in Section 3.2, so we instead implement this moment using the subset of observations where $d' = d$.[17] The same logic allows us to omit $\{\gamma_o, \gamma'_o, w'_o\}_{o \neq d}$. Since baseline trade flows imply a unique (to scale) vector of wages, $\{w_o\}$, we can solve for these directly. We can also normalize $w_d = w'_d = 1$ as well as one element of $\{A_{od}\}_o$. Finally, we can omit equation (5) as it is directly implied by the full set of trade flow moments. This means we are left with a parameter vector that consists of $\{\gamma_d, \gamma'_d, \{A_{od}\}_{o \neq 1}\}$, which is of dimension $D + 1$.

The Melitz case is similar. In addition to omitting $A_{od'}$ for $d' \neq d$, we can also omit $\{A_{od}\}_{o \neq d}$. This is feasible because the X_{dd} trade flow moment—which can never be omitted as it is necessary to identify A_{dd} and thus γ'_d—has no direct dependence on other countries, except through d-level aggregate objects such as γ_d.[18] Moreover, combining the zero-profit condition (15) and the market clearing condition (17) allows us to eliminate $A_{od'}$ for which d is the exporter as aggregate profits are zero and these variables are not directly needed to determine d's consumption. These observations imply it is sufficient to aggregate all other countries $d' \neq d$ and reduce the overall set of parameters to $\{\gamma_d, \gamma'_d, A_{dd}, f_{dd}, f_{Ed}, A_{dx}, f_{dx}\}$, which has dimension 7.

3.4 Data

Our analysis focuses on data from 2016. We obtain aggregate bilateral trade data $\widetilde{X}_{od}$ from the OECD ICIO database and tariff data t_{od} from TRAINS, constructing a measure of the average bilateral tariffs by weighting 6-digit HS-level tariffs by their corresponding flows from the CEPII BACI 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 $D = 20$.

4 Quantifying the 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 productive efficiency? For now, our focus lies on estimates of GT_d for the case

other parameters, it is possible to find a vector $\{A_{od'}\}_{o,d' \neq d}$ that can rationalize the omitted trade shares. The dimensionality reduction comes from the observation that we do not need to know the values of these omitted A_{od} to compute our counterfactual of interest.

[17] As this removes the ability to include origin fixed effects, the moment is equivalent to an instrumental variables estimation of log trade flows on $\log w_o(1 + t_{od})$, using $\log 1 + t_{od}$ as an instrument.

[18] This is not possible in the Ricardian case because the X_{dd} trade flow moment contains all $A_{od'}$ for $d' = d$ in the min function.

where d is France, a country with import penetration that is representative among the countries in our sample.

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 efficiencies F are drawn from the particular distribution F^* (e.g. 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 point estimate of GT_d obtained when assuming $F = F^*$ is the point at which our upper and lower bounds from Section 3.2 coincide at the value $\overline{GT}_d(\delta = 0) = \underline{GT}_d(\delta = 0)$.

In addition, we know from the result of Arkolakis et al. (2012) that both the Eaton and Kortum (2002) and Chaney (2008) models will agree on the value of GT_d as long as they are estimated on the same underlying data in a manner that follows the same estimation procedure. That is, both models conclude that the gains from trade for country d follow the Arkolakis et al. (2012) formula

$$\overline{GT}_d(\delta = 0) = \underline{GT}_d(\delta = 0) = 1 - \lambda_{dd}^{\frac{1}{\theta^*}}, \tag{23}$$

where λ_{dd} is the share of d's expenditure at baseline that it sources from itself. That is, these two models' predictions will agree if they are calibrated to the same baseline value of λ_{dd}, and if their estimated values of θ^* agree. And the models will agree on the estimated value of θ^* if their value of θ^* is estimated as (minus) the trade trade elasticity in the gravity equation (18) based on the same dataset.

We replicate standard estimation procedures in order to obtain an estimate of $\overline{GT}_d(\delta = 0) = \underline{GT}_d(\delta = 0)$ for France in 2016. First, the OLS estimate of θ^* in equation (18) is 6.8. Second, in the 2016 data of our application, France has a value of $\lambda_{dd} = 0.835$ (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 $1 - (0.835)^{\frac{1}{6.8}} = 0.026$. In other words, according to the workhorse gravity model, aggregate welfare in France would drop by 2.6% 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 production efficiencies.

4.2 Analytical bounds on the gains from trade

As discussed above, a natural place to start this exploration of alternative efficiency distributions $F \neq F^*$ is with complete agnosticism about such alternatives. The results in Propositions 1 and 2 enable this, by characterizing the maximum and minimum values of GT_d that are possible if *any* distribution F is allowed for. That is, they establish a formula for sharp bounds on GT_d that would be obtained if one could search across all distributions within the neighborhoods $\mathcal{N}_\delta$ that lie within an infinite divergence ($\delta = \infty$) of F^* in either model. Hence, following the conclusions of Propositions 1 and 2, we have that—according to either the competitive or monopolistically competitive trade model—the analytical lower bound on GT_d is given by $\underline{GT}_d(\delta = \infty) = 0$. And the analogous upper bound is given by

$$\overline{GT}_d(\delta = \infty) = 1 - \lambda_{dd}^{\frac{1}{\sigma-1}}. \tag{24}$$

Like the Arkolakis et al. (2012) formula for GT_d under $\delta = 0$, this expression for the maximum possible value of GT_d depends on the baseline value of λ_{dd}. But rather than depending on the trade elasticity θ^*, $\overline{GT}_d(\delta = \infty)$ depends on σ, the elasticity of substitution across products. Perhaps surprisingly, the value of σ does not enter the GT_d formula under F^* 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 $F \neq F^*$. We therefore work with the value of $\sigma = 2.5$ and explore robustness to this value in ongoing work.

Using this value of σ implies that the maximal value that France's gains from trade can take, across all possible production heterogeneity distributions F, is 11.3% (i.e. $1 - 0.83^{\frac{1}{2.5-1}}$). This value is 4.3 times larger than that predicted by the workhorse model under F^*. Similarly, the lower bound of zero is meaningfully different from that implied at F^*. Evidently, if one were completely agnostic about the model's underlying production heterogeneity then a very wide range of predictions about GT_d are possible.

4.3 Interpreting measures of divergence δ

One reaction to the wide bounds on GT_d across *all* distributions (i.e. for $\delta = \infty$) discussed above is that complete agnosticism is too nihilistic. It is therefore natural to ask what our bounds would look like when we consider a smaller set of productivity distributions by narrowing the neighborhood size $\mathcal{N}_\delta$ such that δ is finite (but still greater than zero). In

order to explore this, we first pause to anticipate the question of how a divergence of any given value, such as δ, can be interpreted.

To provide one benchmark, consider the sensitivity analysis conducted inside gravity models when the value of the trade elasticity is adjusted. This is common practice among users of quantitative trade models. Begin with the Eaton and Kortum (2002) case, where the global productivity distribution F^* implies that $z_o(\omega)$ is distributed, across o and ω, such that each country o draws its own $z_o(\omega)$ independently from the distribution $Frechet(1, \theta^*)$. 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 trade elasticity, θ. That is, $F^{\text{alt.}}(\theta)$ is analogous to F^* but based on $Frechet(1, \theta)$ at some $\theta \neq \theta^*$. Because this distribution $F^{\text{alt.}}(\theta)$ differs from F^* we can calculate the resulting $D(F^{\text{alt.}}(\theta)||F^*)$. The particular version of divergence $D(F||F^*)$ that we use here and below is the "hybrid" divergence, recommended by Christensen and Connault (2023) for cases like ours where F^* has fat tails.[19]

The red line in Figure 1 plots the value of $D(F^{\text{alt.}}(\theta)||F^*)$ against θ for this Frechet case—that is, F^* is based on a $Frechet(1, \theta^*)$ distribution with $D = 20$ countries and using the value of $\theta^* = 6.8$ estimated in Section 4.1. Similarly, the blue line reports on an analogous exercise for the case of Chaney (2008), where F^* is based on a $Pareto(1, \theta^*)$ distribution. Naturally, in both cases the divergence is zero at $\theta = \theta^* = 6.8$, and the divergence is strictly greater than zero for all other values of $\theta \neq \theta^*$.

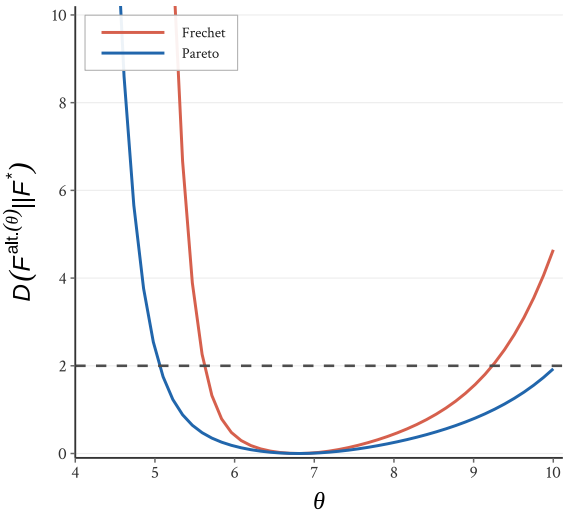


Figure 1: Examples of divergence between distributions

[19] In this case, the function $\phi(\cdot)$ from equation (20) is $\phi(x) \equiv x \ln x - x + 1$ for all $x \leq e$ and $\phi(x) \equiv (x-e)^2/(2e) + (x-e) + 1$ for all $x > e$. This means that $D(F||F^*)$ corresponds to the Kullback-Leibler divergence at low values and the χ^2-divergence at higher values.

This figure provides a quantitative comparison between conventional windows of sensitivity analysis and their corresponding values of divergence. For example, changing the trade elasticity from, say, $\theta^* = 6.8$ to $\theta \in [5.3, 9]$ corresponds to a divergence of approximately 1 in the Pareto case. Alternatively, a window that satisfies $D(F^{\text{alt.}}(\theta)||F^*) \leq 2$ corresponds to trade elasticities θ (centered on $\theta^* = 6.8$) that span the range $\theta = [5.5, 9]$ in the Frechet case and $\theta = [5, 10]$ 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. We therefore view a neighborhood such as $\mathcal{N}_\delta$ at $\delta = 2$—that is, the set of *all* distributions F that lie within $D(F||F^*) \leq 2$ of F^*—to be of substantial interest when conducting sensitivity analysis.

4.4 Sensitivity of the gains from trade to alternative efficiency distributions

We now turn to the main goal of our paper, which is to estimate bounds counterfactual predictions about the gains from trade that hold for a varying set of neighborhoods $\mathcal{N}_\delta$ centered around the workhorse gravity model.

As described in Section 3.2, 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. We focus on case where $d =$ France for now.

Figure 2 plots our baseline results. For purposes of comparison, the dashed orange line displays the value of $\overline{GT}_d(\delta = 0)$, the extent of gains from trade in the workhorse model (i.e. $\delta = 0$); as discussed in Section 4.1, we estimate this to be $GT_d(\delta = 0) = 0.026$. In addition, the dashed black lines report the analytical upper and lower bounds, $\overline{GT}_d(\delta = \infty)$ and $\underline{GT}_d(\delta = \infty)$ respectively, that are possible across all distributions; as discussed in Section 4.2, we estimate these bounds to be $\overline{GT}_d(\delta = \infty) = 0.113$ and $\underline{GT}_d(\delta = \infty) = 0$. On top of these lines we add 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 bounds converge to the standard gravity model estimate of $GT_d = 0.026$ as δ falls to zero. And these bounds asymptote towards the analytical bounds $\overline{GT}_d(\delta = \infty)$ and $\underline{GT}_d(\delta = \infty)$ as δ grows toward infinity.

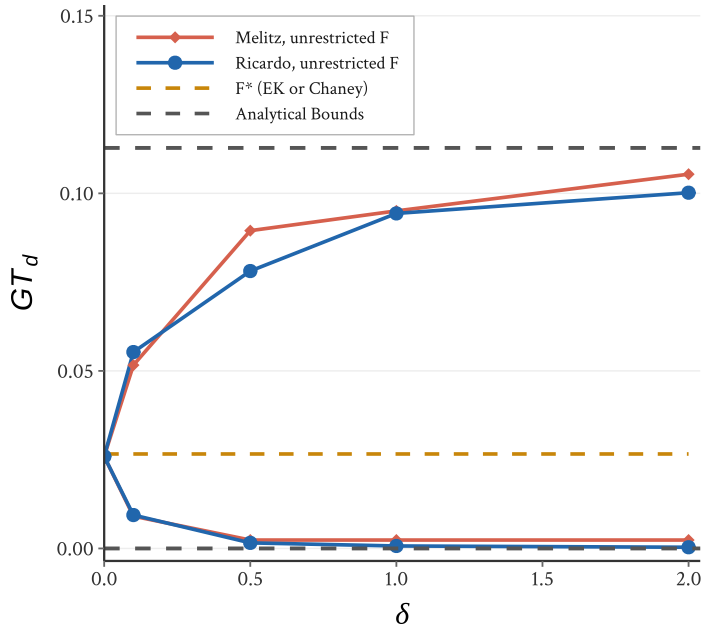


Figure 2: Bounds on Gains from Trade (for France) Across Alternative Productivity Distributions

Of greater 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 we see a striking amount of sensitivity. For example, at $\delta = 1$ the trade data support Ricardian models with gains from trade that are four times larger than in the gravity model, as well as Melitz (2003) models that are also four times larger. And at $\delta = 2$—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 upper bounds in Figure 2 are very close (88% of the way for Ricardo and 92% for Melitz (2003)) to the analytical maximal upper bound that is about five (4.78) times larger than in the gravity model. Evidently, a neighborhood of divergence within $\delta = 2$ of the gravity model F^* approximately spans the space of relevant distributions—it includes distributions that lead to viable general equilibrium trade models of either perfect or monopolistic competition, that are completely capable of matching all data on trade flows, and that can approximately achieve the maximal and minimal gains from trade that are possible across all distributions in these trade models.

These results have several implications. First, one might hope that empirically constraining trade models to match actual trade flows, and hence their regularities such as the fit of the gravity equation or the slope of the trade elasticity, would do a great deal to ensure the models' robustness to alternative ways of modeling the unobserved produc-

tivity distribution F. Unfortunately, Figure 2 demonstrates that this is not the case since all points inside the bound curves $\overline{GT}_d(\delta)$ and $\underline{GT}_d(\delta)$ are equally good at matching all trade data yet many of these points correspond to widely different values for GT_d.

Second, one might hope that the distribution of productivities F is relatively unimportant to the gains from trade, amidst so many other model features and parameters. Again, this view is inconsistent Figure 2, since all of the points within our bounds hold all model ingredients other than the productivity distribution F fixed.

Third, one might hope that the conventional way of exploring sensitivity in quantitative trade modeling, via considering alternative values for the trade elasticity, would span the space of likely possibilities. Yet, unless one views some distributions as inherently more likely than others, this view is also rejected by Figure 2 since, at any given δ, the bounds we compute across all distributions are are considerably wider than those across distributions that retain the gravity structure at alternative trade elasticities.

Finally, as discussed in the Introduction, one might hope that the Ricardian model would inherit robust properties from the logic of selection, due to perfect competition. The logic here is that the extremal types theorem predicts a tendency for selected draws—the ones that matter for welfare and hence GT_d—to converge to a fixed family of distributions regardless of the underlying stochastic process. To the extent that this logic could limit the sensitivity of the Ricardian model to F, it is not apparent in Figure 2.

5 Restricting Productivity Distributions

We have seen in Figure 2 how the estimates that quantitative trade models provide about gains from trade are highly sensitive to the model's assumption about the underlying productivity distribution F. That is, even for the neighborhood of distributions that are "close" to the gravity model benchmark of F^*, the gains from trade can be many times smaller or larger than they are in the model that assumes F^*. These findings beg the question: how much narrower would the bounds on GT_d be if one were to consider only a restricted class of distributions F?

In this section we provide a set of answers to this question. Of course, there are many ways to restrict the types of distributions under consideration.[20] We focus on a set of high-level characteristics that are assumed in F^* and see how the bounds on GT_d change as we impose some of these characteristics, one by one.

[20] Recall that the divergence between distributions is invariant to a change in the units of measurement in either distribution. So one type of restriction that is irrelevant to our conclusions is the pure location (or units of account) of the distributions we consider.

The global productivity distributions F^* used by Eaton and Kortum (2002) in the Ricardian model and Chaney (2008) in the Melitz model embody three high-level characteristics. First, each country takes independent draws—that is, there is no cross-country dependence or correlation across their draws. Second, every country draws from the same marginal distribution. And third, the common marginal distribution that all countries draw from takes the specific form of $Frechet(1, \theta*)$ in the Eaton and Kortum (2002) case and $Pareto(1, \theta*)$ in the Chaney (2008) case.

So far, our investigations have searched over the neighborhood $\mathcal{N}_\delta$ for distributions that, while always generating trade models that match the data, are otherwise unrestricted. This means that the distributions $F \in \mathcal{N}_\delta$ considered are relaxing some or all of the three high-level characteristics held by F^* as they move throughout the less restricted space of distributions. In this section we impose constraints on this flexibility—subsets of the three high-level conditions embodied in F^*—in order to learn about which conditions are more restrictive in the economic sense of constraining the gains from trade.

5.1 Restrictions in Ricardian vs Melitz models

One starting observation concerns a broad distinction between the Ricardian and Melitz models. The Ricardian model involves head-to-head competition, in every destination market d, over D worldwide producers of each good ω. 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 won this head-to-head competition) to a domestic supplier. The gains from trade therefore depend on the size of the foreign supplier's victory. The statistical properties of this victory depend on country d's marginal efficiency 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 correlation. In short, the gains in the Ricardian model clearly depend on on properties of domestic marginals, foreign marginals, and the dependence among them. We explore this below.

On the other hand, consider the Melitz model. In this case, competition is not head-to-head for any good; to the contrary, every firm pays an entry cost in order to be the unique 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 what matters for the gains is purely a function of the domestic versions. That is, conditional on knowing the marginal distribution of efficiency in country d the additional properties of F—such as the $D-1$ foreign countries' marginal distributions or the cross-country

dependence of efficiency draws—has no bearing at all on GT_d.

This observation means that the red line in Figure 2 describes the gains from trade in the Melitz model not only in the restricted case, but also in an additionally restricted case where the distributions F that are considered must satisfy some of the conditions on F^* such as independence and common marginals. Even the case in which all $D-1$ foreign marginals are restricted to be $Pareto(1,\theta^*)$ is covered too. This means that the bounds in the Melitz case of Figure 2 derive purely from uncertainty about the distribution of country d's own marginal distribution. When that marginal is $Pareto(1,\theta^*)$ then GT_d is given by the orange line corresponding to F^* but other marginal distributions for d generate gains from trade that lie within the red lines.

Given these points, there is nothing 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. We introduce the three types of high-level restrictions—independence, common marginals, and Frechet marginals—in turn.

5.2 Imposing independence and zero correlation

As discussed, the F^* used by the Eaton and Kortum (2002) version of the Ricardian model assumes that countries o take their efficiency draws $z_o(\omega)$ for any good ω independently from one another. We now seek to restrict the set of distributions F considered when constructing bounds such as $\overline{GT}_d(\delta)$ via program (21) such that any admissible F satisfies cross-country independence just like F^* does. In this case the only distinction between F^* and F that remains is particular marginal distributions that each countries draw from.

Recall that two jointly distributed random variables are independent if and only if it is the case that their joint CDF is equal to the product of their two marginal CDFs. We approximate this restriction about CDFs by using a finite number of expectations at L specific quantiles of the relevant CDFs. That is, for any pair of countries $o \neq o'$ and any quantile l, approximate independence requires

$$\mathbb{E}_F[\mathbb{1}\{z_o(\omega) < \bar{z}_l\}\mathbb{1}\{z_{o'}(\omega) < \bar{z}_l\}] = \mathbb{E}_F[\mathbb{1}\{z_o(\omega) < \bar{z}_l\}]\mathbb{E}_F[\mathbb{1}\{z_{o'}(\omega) < \bar{z}_l\}], \tag{25}$$

where, for example, $\bar{z}_l$ denotes the value of the lth quantile of a distribution of interest.[21]

That is, equation (25) imposes the necessary and sufficient condition for independence among all pairs of countries, at a finite number of points in the distribution. As L grows

[21] In practice, we use F^*, i.e. $Frechet(1,\theta^*)$, for this purpose. This has the advantage of using quantiles that are spaced out in the same proportion to their distribution under F^*.

this approximates the case of true independence.[22] We impose approximate independence on F by adding the $D(D-1)L$ constraints embodied in (25) to the program (21) when computing the upper bound $\overline{GT}_d(\delta)$ (and analogously for the lower bound $\underline{GT}_d(\delta)$).

Figure 3 displays the results from implementing this independence restriction: blue lines show the unrestricted bounds from Figure 2 and the dark green lines show those that impose independence. In the case of the lower bound $\underline{GT}_d(\delta)$, these lines are essentially coincident, implying that the bounds are unperturbed by the requirement that F features cross-country independence. On the other hand, the upper bound $\underline{GT}_d(\delta)$ with independence is somewhat lower than that facing no restrictions. For example, at the $\delta = 2$ level, $\underline{GT}_d(\delta)$ is 0.086 when F must exhibit cross-country independence and 0.1 when it F is unrestricted. This implies that French gains from trade would be about 18% (i.e. $1 - (0.086 - 0.026)/(0.1 - 0.026)$) less sensitive to the underlying efficiency distribution F if it were believed that countries draw their technologies independently.

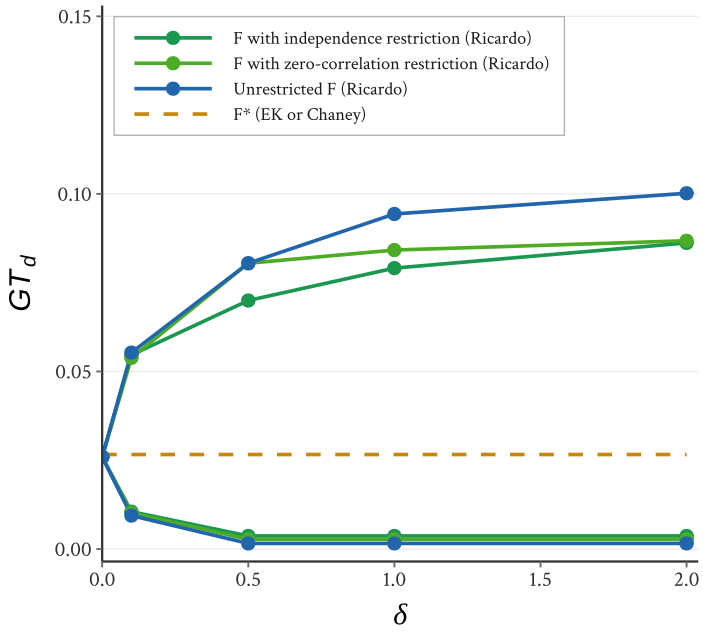


Figure 3: Gains from Trade and Cross-country Dependence

Stopping short of full independence, one may also wish to impose weaker restrictions on the dependence of efficiency draws across countries. For example, the condition of zero correlation can be obtained by imposing the far simpler version of (25) based on first moments alone:

$$\mathbb{E}_F[z_o(\omega)z_{o'}(\omega)] = \mathbb{E}_F[z_o(\omega)]\mathbb{E}_F[z_{o'}(\omega)] \quad \text{for all } o \neq o'. \tag{26}$$

[22] In practice, we set $L = 5$ due to the large number of pairwise combinations in (25).

This zero-correlation case is also shown in Figure 3. Naturally, since this restriction is weaker than that of full independence, the zero-correlation restriction bounds (shown with light green lines) are weakly wider than the independence restriction bounds (dark green lines). But the lines are close, especially for δ above 0.5. Perhaps surprisingly, therefore, there is not much of a distinction to be had from more complicated forms of dependence than linear correlation.

These results are also informative for the prior literature that has relaxed the independence assumption in Eaton and Kortum (2002). For example, Lind and Ramondo (2023) maintain the case of Frechet marginals but allow for non-zero amounts of correlation. Figure 3 illustrate an analogous idea but in the opposite direction: start with all distributions (blue lines) and then impose restrictions to as to keep only those with zero correlation (light green lines). We learn that this can reduce the maximal gains from trade by about 18%, but has essentially no effect on the minimal gains from trade.

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 efficiencies $z_o(\omega)$ 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 it is interesting to consider how the feature of commonality—which may derive from common primitive assumptions on where technologies derive form, or simply from a desire for parsimony—alone affects our results.

We again approximate this restriction about (marginal) CDFs by using a finite number of expectations at L quantiles l. All countries' marginal distributions are approximately common if the following condition holds

$$\mathbb{E}_F[\mathbb{1}\{z_o(\omega) < \bar{z}_l\}] = \mathbb{E}_F[\mathbb{1}\{z_{o'}(\omega) < \bar{z}_l\}] \text{ for all } l \leq L \text{ and } o' \neq o. \tag{27}$$

That is, if the expected mass of country o's distribution below any given point $\bar{z}_l$ is the same as that for any other country o'.[23] 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 $(D-1)L$ constraints embodied in (27) to the program (21).

While this method works arbitrarily well for large L, one challenge in practice occurs

[23] We again choose the points $\bar{z}_l$ as the L quantiles of the $Frechet(1, \theta^*)$ distribution. For any given L, he common marginals restriction in (27) requires far fewer restrictions then the pairwise restrictions required for independence in (25), so it is feasible to use a larger value of L. We set $L = 50$ here.

in regards to the likelihood of the highest-efficiency draws, which are particularly important for generating gains from trade. In particular, since gains from trade depend on the price index, which is a transformation of the $(1-\sigma)$th moment of the price draws that consumers face, we mitigate against the risks of inadequate upper-tail approximations by further imposing the additional $(D-1)L$ constraints

$$\mathbb{E}_F[z_o(\omega)^{(1-\sigma)}\mathbb{1}\{z_o(\omega) < \overline{z}_l\}] = \mathbb{E}_F[z_{o'}(\omega)^{(1-\sigma)}\mathbb{1}\{z_{o'}(\omega) < \overline{z}_l\}] \text{ for all } l \leq L \text{ and } o' \neq o. \tag{28}$$

This ensures that the draws within each quantile have the same $(1-\sigma)$th moment.

The orange lines in Figure 4 report the results from this exercise. (Again, the blue lines display the unrestricted baseline case for comparison.) The lesson from this figure is that imposing the common marginals constraint does not do much to reduce the sensitivity of quantitative Ricardian models. The lower bound is again unaffected by this constraint. And the upper bound is reduced, but only by one fifth (i.e. $1-(0.092-0.026)/(0.1-0.026)$) at a value like $\delta = 2$.

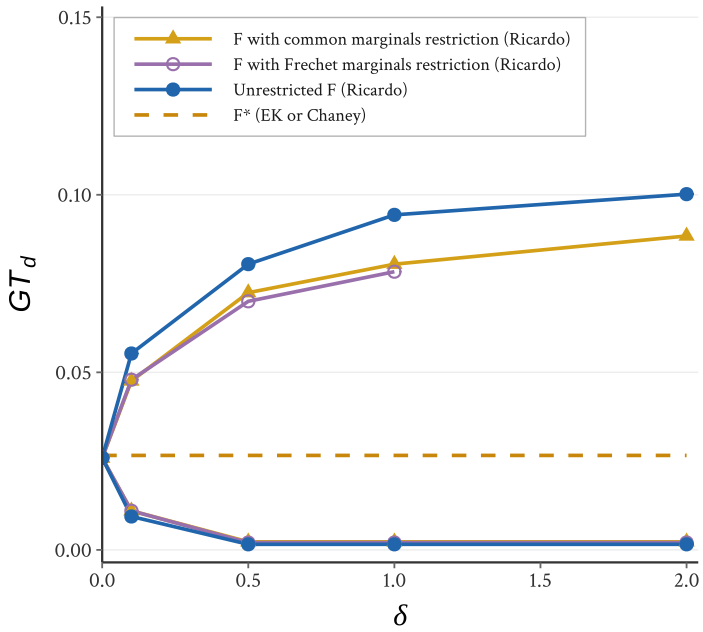


Figure 4: Gains from Trade and Countries' Marginal Distributions

5.4 Imposing common Frechet marginals

The previous restriction on F requires that all countries share the same marginal distribution, but 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 and Kortum (2002). 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. Lind and Ramondo (2023) explored this case by studying Ricardian models in which countries' marginal distributions were Frechet distributed, but the cross-country correlations in efficiency draws were parameterized in various respects. The restriction we study here is in the same spirit except that we allow for countries to be taking draws from *any* joint process (formally, any copula) with Frechet marginals (that leads to a Ricardian model that matches all data), not any particular correlation structure.

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 efficiency draws $z_o(\omega)$ in any country o has 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. Formally, this amounts to imposing the constraint

$$\mathbb{E}_F[\mathbb{1}\{z_o(\omega) < \overline{H}_h\}] = H(\overline{H}_h) \quad \text{for all } o \text{ and all } h \leq \overline{H}. \tag{29}$$

In addition, for the same reasons as above, we ensure that the $(1-\sigma)$th moment of each quantile of every country's distribution is the same as that under $Exp(1)$:

$$\mathbb{E}_F[z_o(\omega)^{1-\sigma}\mathbb{1}\{z_o(\omega) < \overline{H}_h\}] = \mathbb{E}_{F^*}[z_o(\omega)^{1-\sigma}\mathbb{1}\{z_o(\omega) < \overline{H}_h\}] \quad \text{for all } o \text{ and all } h \leq \overline{H}. \tag{30}$$

Adding the $2D\overline{H}$ constraints in (29) and (30) to the program (21) ensures that bounds on GT_d at any δ only allow distributions F that deviate from F^* in terms of the cross-country dependence of their technology draws and not at all in terms of the Frechet marginals that all countries share.

We report on these effects—with common-and-$Frechet(1, \theta^*)$ marginals—in Figure 4. The red lines report on the new restriction.[24] Again, counterfactual sensitivity remains substantial, despite restricting all countries to have the same marginal distributions as they do in Eaton and Kortum (2002). Perhaps surprisingly, the red lines (where all marginals are the same and Frechet) are very close to the orange lines (where all marginals are the same but can be *any* marginal distribution). So, at least when dependence is left

[24] In practice we set $\overline{H} = 50$.

flexible, the sensitivity of the Ricardian gravity model's predictions has little to do with the particular marginal distributions employed.

This result generalizes the case studied by Lind and Ramondo (2023)—that with common $Frechet(1, \theta^*)$ marginals for all countries, but with various amounts of non-zero correlation. We learn from the red lines in Figure 4 that there are widely differing amounts of gains from trade in Ricardian models, even those that all have $Frechet(1, \theta^*)$ marginals, if the cross-country dependence in efficiency draws is unknown and left flexible.

5.5 Imposing both independence and common marginal distributions

The four restrictions we have introduced so far have all been imposed one at a time. But they can also be imposed simultaneously. One combination of interest is the case where F displays cross-country independence and all countries share the same marginal distributions, in which case the only remaining distinction relative to F^* is that the marginals are flexible rather than $Frechet(1, \theta^*)$. We do this by adding to the program (21) the stacked set of constraints in (25), (27), and (28).

The results of this exercise are shown in Figure 5. In particular, the light blue lines report on the case where all permitted distributions F satisfy both independence and common marginals. For comparison, the dark blue lines again display the unrestricted case and the orange lines display that with common marginals (but not independence). Once again, the bounds conveyed by the light blue lines are still quite wide. Going from the unrestricted case to the doubly-restricted case of both independence and common marginals achieves, at $\delta = 2$ for example, a reduction in sensitivity of the upper bound by about 23.5% (i.e. $1 - (0.091 - 0.026)/(0.1 - 0.026)$). What is perhaps most surprising is that the combination of independence and common marginals (light blue in this figure) achieves little more robustness in $\overline{GT}_d(\delta)$ than that of either independence (Figure 3) or common marginals (orange line in this figure) on their own.

5.6 Summary

Stepping back, the combined impressions from the restrictions imposed in Figures 3 through 5 leaves the following lesson. Gains from trade in the Ricardian model can be many times larger or smaller than they are in the workhorse Eaton and Kortum (2002) case when one departs from that model's efficiency distribution F^* in *either* of the two directions available. One direction is to relax independence, as we have seen in the case with Frechet marginals but unrestricted dependence in Figure 3. The other direction is to relax the Frechet marginals, as we have seen in the case with independence and common

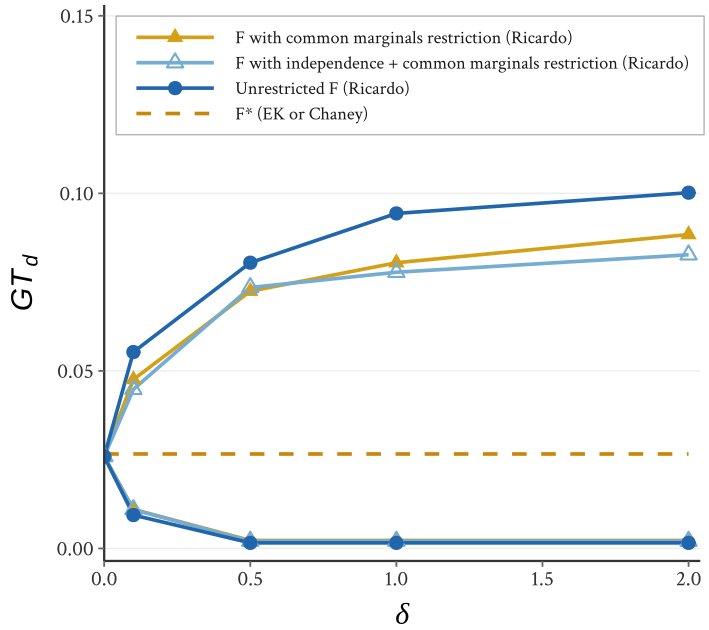


Figure 5: Gains from Trade with Independence and Common Marginals

but otherwise unrestricted marginals in Figure 5. That is, at least as far as gains from trade are concerned, there are important ways to deviate from F^*, even in the small neighborhood of $\delta = 2$ and while matching all available data, by perturbing either the common marginal that countries share or the extent to which countries draw their efficiencies in dependent ways.

6 Concluding Remarks

Three modern revolutions have taken place in the study of international trade. The first is the desire to answer counterfactual questions using quantitative simulations. The second is the incorporation of micro-level production heterogeneity across goods and countries, either in Ricardian models a la Eaton and Kortum (2002) or or monopolistically competitive models a la Melitz (2003). And the 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 the trade elasticity, 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. It is also true 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.

These findings leave us with two directions for future work. One is to ask: what would reduce the sensitivity of quantitative trade models? We have seen how high-level restrictions will not reduce this sensitivity. But are there particular data features that might? Are there augmented empirical facts, beyond the baseline value of each trade flows, that would help to condition model predictions?

The other direction is to augment our baseline economic environment. We have omitted features such as domestic input-output linkages, international supply chains, multiple sectors, multiple factors, and variable markups. These features typically widen the gains from trade in quantitative models with fixed assumptions about production heterogeneity, but it remains an open question whether this remains true despite the presence of a flexible approach to production heterogeneity.

A final direction we are pursuing in ongoing work considers a nesting model that allows us to think of the competitive (Ricardian) and monopolistically competitive (Melitz-style) models as points on the spectrum of unified hybrid models with flexible modes of competition.

References

Allen, Treb and Costas Arkolakis, "Quantitative regional economics," *Handbook of Regional and Urban Economics*, 2025, *6*, 1.

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

Caliendo, Lorenzo and Fernando Parro, "Trade policy," *Handbook of International Economics*, 2022, *5*, 219–295.

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," *Handbook of International Economics*, 2014, *4*, 197–261.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Unbalanced Trade," *NBER working paper 13035*, 2007.

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

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

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

Redding, Stephen J, "Quantitative urban economics," *Handbook of Regional and Urban Economics*, 2025, *6* (1), 73–141.

A Theory Appendix

Proposition 2. *Consider the monopolistic competition model. Gains from trade under any F satisfy $0 \leq GT_d \leq 1 - \lambda_{dd}^{\frac{1}{\sigma-1}}$. Moreover, there exist F^- and F^+ that attain the lower and upper bound, respectively.*

Proof. Consider first the upper bound. The definition of the trade share implies

$$\lambda_{dd} = \frac{N_d \mathbb{E}_F[p_{dd}(\omega)^{1-\sigma} 1\{X_{dd}(\omega) > \sigma f_{dd} w_d\}]}{\gamma_d}. \tag{31}$$

The entry condition is equivalent to a productivity cutoff:

$$z_d(\omega) > \left(\frac{\sigma f_{dd} w_d \gamma_d}{E_d}\right)^{\frac{1}{\sigma-1}} \frac{w_d(1+t_{dd})}{A_{dd}} \equiv z_d^*.$$

The only term that can change between the baseline and counterfactual equilibrium is γ_d, so it is easier to write $z_d^* = \Xi_{od} \gamma_d^{1/(\sigma-1)}$, where Ξ_{od} contains the remaining terms.[25] This cutoff will decrease when d moves to autarky. Formally, suppose the cutoff was greater, i.e., $(z_d^*)' \geq z_d^*$. Then, the equivalent of Equation (31) in autarky implies

$$\gamma_d' = N_o \mathbb{E}_F[p_{dd}(\omega)^{\sigma-1} 1\{z_d(\omega) > (z_d^*)'\}] \leq N_o \mathbb{E}_F[p_{dd}(\omega)^{\sigma-1} 1\{z_d(\omega) > z_d^*\}] = \lambda_{dd} \gamma_d.$$

Then, assuming that any potential trade deficit is a fixed share of GDP, we can write

$$(z_d^*)' = \Xi_{od} (\gamma_d')^{\frac{1}{\sigma-1}} \leq \Xi_{od} (\lambda_{dd} \gamma_d)^{\frac{1}{\sigma-1}} = \Xi_{od} (\lambda_{dd} \gamma_d)^{\frac{1}{\sigma-1}} = \lambda_{dd}^{\frac{1}{\sigma-1}} z_d^* < z_d^*,$$

which is a contradiction. Thus, it must be that $(z_d^*)' < z_d^*$. Again using Equation (31), it then immediately follows that $\gamma_d' \geq \lambda_{dd} \gamma_d$, in which case the expression $GT_d = 1 - (\gamma_d'/\gamma_d)^{1/(\sigma-1)}$ is maximized when $\gamma_d' = \lambda_{dd} \gamma_d$, i.e., at $1 - \lambda_{dd}^{\frac{1}{\sigma-1}}$. This is achieved when the mass of firms between the new domestic cutoff and the old domestic cutoff is zero–i.e., the consumer does not replace the lost imports with any new varieties. Any distribution with this property will generate these maximal gains from trade. As with the Ricardian case, the model becomes effectively ex post Armington because it is as if foreign varieties cannot be produced domestically.

[25] Total expenditure, E_d, does not change because we can normalize $w_d = w_d' = 1$, deficits are a constant share of GDP, and profits are zero due to free entry.

For the lower bound, note that the definition of γ_d is

$$\gamma_d = \sum_o N_o \mathbb{E}_F \left[p_{od}(\omega)^{1-\sigma} 1\{z_o(\omega) > z^*_{od}\} \right].$$

This implies that we can write

$$GT_d = 1 - \left(\frac{\mathbb{E}_F[N_d p_{dd}(\omega)^{1-\sigma} 1\{z_d(\omega) > (z^*_{dd})'\}]}{\mathbb{E}_F[\sum_o N_o p_{od}(\omega)^{1-\sigma} 1\{z_o(\omega) > z^*_{od}\}]} \right)^{\frac{1}{\sigma-1}}.$$

Since the denominator is weakly greater than the numerator, the expression in parentheses is maximized at 1 and thus GT_d is minimized at 0. To construct a distribution that achieves these bounds, note again that the contribution of domestic varieties produced at baseline to the price index is $N_d \mathbb{E}_F[p_{dd}(\omega)^{1-\sigma} 1\{z_d(\omega) > z^*_{dd}\}] = \lambda_{dd}\gamma_d$. The varieties from importing that are lost under autarky contribute $(1 - \lambda_{dd})\gamma_d$, so we need to ensure that the additional domestic firms producing in autarky (those above the counterfactual cutoff, $(z^*_{dd})'$) generate exactly $(1 - \lambda_{dd})\gamma_d$ in when aggregated. We can achieve this by placing mass in the interval $[(z^*_{dd})', z^*_{dd}]$ such that $N_d \mathbb{E}_F[p_{dd}(\omega)^{1-\sigma} 1\{(z^*_{dd})' < z_d(\omega) < z^*_{dd}\}] = (1 - \lambda_{dd})\gamma_d$. Intuitively, there is a latent mass of domestic firms that do not produce at baseline but do produce in autarky. Even if foreign firms are much more productive on average than domestic firms, the fact that there are two separate margins (the share of firms that produce and the average productivity of such firms) means we can make this latent mass of domestic autarky-only producers sufficiently large so as replace them. □