Gravity with Granularity*

Holger Breinlich† Harald Fadinger‡ Volker Nocke§ Nicolas Schutz¶

May 31, 2023

Abstract

We demonstrate that the estimation of gravity equations of trade flows suffers from an omitted variable bias when firms are granular and behave oligopolistically. We show how to correct for this bias in the estimation of both firm- and industry-level gravity. Using French and Chinese export data, we find that the oligopoly bias leads to a substantial underestimation of the effects of distance on trade flows. In a calibrated version of the model, the welfare gains from a trade liberalization are found to be almost twice as large under oligopoly as under monopolistic competition.

Keywords: Gravity Equation, Oligopoly, Trade Liberalization, Trade Elasticity

Journal of Economic Literature Classification: F12, F14, L13

*We thank seminar participants at Berkeley, DICE, Essex, Geneva, LSE, Munich, Nanjing, Nottingham, PSE, SSE, and Surrey, and conference participants at WIEN 2022, ENTER Jamboree 2022, Bayreuth Trade and IO Workshop 2022, EARIE 2019, ETSG 2019 and the international economics conference of the German Economic association in 2019 for helpful suggestions. We gratefully acknowledge financial support from the German Research Foundation (DFG) through CRC TR 224 (Projects B03 and B06). Access to French confidential data, on which some of this work is based, has been made possible within a secure environment offered by CASD – Centre d'accès sécurisé aux données (Ref. 10.34724/CASD).

†University of Surrey, CEP and CEPR. Email: h.breinlich@surrey.ac.uk.

‡Department of Economics and MaCCI, University of Mannheim. Also affiliated with CEPR. Email: harald.fadinger@uni-mannheim.de.

§Department of Economics and MaCCI, University of Mannheim, and School of Economics, University of Surrey. Also affiliated with CEPR. Email: volker.nocke@gmail.com.

¶Department of Economics and MaCCI, University of Mannheim. Also affiliated with CEPR. Email: schutz@uni-mannheim.de.

1 Introduction

Gravity equations have been the predominant tool for analyzing the determinants of bilateral trade flows since their introduction by Tinbergen (1962) over 60 years ago. In their most basic form, gravity equations predict that trade between countries is a log-linear function of the economic mass of the two trading partners and bilateral frictions such as distance or tariffs. Even in this simple form, gravity equations have substantial explanatory power, often explaining in excess of 70-80% of the variation in the trade flows between countries. Starting with Anderson (1979), researchers have shown that gravity equations can be derived from a number of mainstream theoretical frameworks, allowing a tight link to economic welfare analysis. Not surprisingly then, gravity equations have become the workhorse tool for evaluating trade-related economic policies, such as tariffs, trade agreements or WTO membership.

Despite the rapid progress that research on gravity equations has made over the past decades, existing approaches remain at odds with a key stylized fact about international trade, however: Much of world trade is dominated by a small number of large firms. The classic example is the market for wide-bodied passenger aircraft which comprises just two firms (Airbus and Boeing); but the markets of many other tradable goods such as cars, mobile phones or television sets are also dominated by a handful of large producers. That is, in the language of Gaubert and Itskhoki (2021), trade flows are "granular". Given their size, it seems likely that such "granular" firms enjoy substantial market power and have incentives to internalize the effects of their actions on aggregate market outcomes. In this paper, we evaluate the consequences of oligopolistic behavior for the estimation of gravity equations.

Under oligopoly, standard approaches to gravity estimation deliver inconsistent estimates of key parameters, such as the trade elasticity with respect to distance. The reason is that markups co-vary systematically with bilateral variable trade costs (e.g., distance or tariffs) and are contained in the error term of the gravity equation. The key intuition is that firms selling in destinations with higher bilateral trade costs face higher marginal costs, and that these higher marginal costs are incompletely passed through under oligopoly.[1] This induces a classical omitted variable bias, which leads to an under-estimation of the trade elasticity: The value of exports does not fall as much with variable trade costs as it would fall if markups were held constant, because firms systematically reduce markups when selling to destinations

[1] For evidence on incomplete cost pass-through in the industrial organization and international trade literatures, see Feenstra (1989), Nakamura and Zerom (2010), Burstein and Gopinath (2014), Ganapati, Shapiro, and Walker (2020), and Genakos and Pagliero (2022).

with higher trade costs.

We derive firm- and industry-level gravity equations from a rich heterogeneous-firm model with oligopolistic competition and product differentiation based on Atkeson and Burstein (2008). Consumers have CES preferences with industry-specific demand elasticities, firms face industry-specific returns to scale and draw idiosyncratic productivity and quality shocks. We show how to consistently estimate gravity-equation parameters in this model using firm- and industry-level trade flows.

Specifically, we show how to eliminate the oligopoly bias by constructing a correction term that purges observed trade flows from oligopolistic market power effects. At the firm level, the correction term uses information on firms' market shares and demand elasticities and returns to scale.[2] At the industry level, the correction term takes the form of an origin-destination-level Herfindahl index (HHI) of exporters multiplied by the exporting country's aggregate market share in the destination market. This is intuitive: Exporters' markups are high if exports are concentrated in a small number of firms that have a large aggregate market share in the destination.

In our empirical applications, we use firm- and industry-level data on exports of French and Chinese firms to European countries, and therefore focus on distance as the only bilateral trade cost variable.[3] We show that failing to account for oligopoly leads to a substantial underestimation of the distance elasticity of trade flows.[4] At the firm level, the average oligopoly bias is around 50%. At the industry level, the bias is around 10% for the average industry but it is substantially larger in a significant minority of industries, in which exports tend to be highly concentrated.

To confirm the validity of our empirical approach, we perform a detailed Monte Carlo study. We calibrate our rich heterogeneous-firm CES oligopoly model to match key statistics of the French and Chinese micro-level trade data, and use it to generate a simulated data-set. We then run firm- and industry-level gravity regressions on that data-set, and find that our oligopoly corrections do very well in recovering the distance coefficient. By contrast, without the oligopoly correction, we obtain a bias of similar magnitude to that in the regressions run on the actual data.

---

[2] As a secondary contribution, we also show how to obtain model-consistent estimates of demand and supply parameters by industry, building on Feenstra (1994) and Broda and Weinstein (2006).

[3] Because of the restriction to European destinations (for reasons outlined below), there is insufficient variation to include other common gravity variables. However, our methodology naturally applies also to policy-relevant variables such as tariffs and regional trade agreements.

[4] As a result, the estimated distance elasticity of trade costs is biased downwards. Using price data for shipments and estimates of cost pass-through to account for variable markups, Atkin and Donaldson (2015) also find such a downward oligopoly bias.

Finally, we use our calibrated model to evaluate the welfare effects of a 10% trade cost reduction. We find that the resulting welfare gains are almost twice as high under oligopoly as under monopolistic competition. This is driven both by the larger estimated distance coefficient under oligopoly compared to monopolistic competition, and by additional pro-competitive gains from trade due to reduced markups.

Related literature. Our paper builds on the literature deriving theory-consistent gravity equations. Anderson (1979), Eaton and Kortum (2002), Anderson and van Wincoop (2003), Chaney (2008), Melitz and Ottaviano (2008), Arkolakis, Costinot, and Rodriguez-Clare (2012), Arkolakis, Costinot, Donaldson, and Rodriguez-Clare (2019) and Allen, Arkolakis, and Takahashi (2020) show how to obtain aggregate/industry-level gravity equations from a variety of theoretical frameworks. We contribute to this literature by deriving theory-consistent gravity equations at the firm- and industry level under oligopoly.

Another strand of the literature, surveyed by Head and Mayer (2014), is concerned with the estimation of gravity equations. We contribute to this literature by proposing methods to estimate gravity equations when firms have market power. Anderson and van Wincoop (2003) highlight the importance of controlling for 'multilateral resistance' (i.e., the price index in the destination). Harrigan (1996) was the first to do so using destination fixed effects, an approach that has been followed in most subsequent studies, including the present paper. Santos Silva and Tenreyro (2006) advocate the use of the Poisson Pseudo Maximum Likelihood (PPML) estimator to address a potential bias arising from heteroscedasticity in log-linearized models—an approach that we also follow. An important problem for the estimation of gravity estimations arises from firms self-selecting into export markets. At the firm level, Bas, Mayer, and Thoenig (2017) propose to focus on top exporters that are present in most destinations.[5] At the industry/aggregate level, Helpman, Melitz, and Rubinstein (2008) propose a two-step estimation procedure, which in addition to the standard Heckman correction also controls for the extensive margin of exports. We adopt both approaches in this paper.

Our paper is among the first to use gravity estimates to evaluate the welfare effects of trade policies under oligopoly. Arkolakis, Costinot, and Rodriguez-Clare (2012) identify a class of models with monopolistic or perfect competition in which the trade elasticity is constant and constitutes (in conjunction with countries' trade shares) a sufficient statistic for the welfare gains from trade.[6] Arkolakis, Costinot, Donaldson, and Rodriguez-Clare (2019) extend this

[5]By focusing on the largest exporters, however, this approach is likely to exacerbate the oligopoly bias, thus making it even more important to control for market power.

[6]See, however, Melitz and Redding (2015) who emphasize the role of micro structure for the gains from

approach to a class of monopolistic competition models with variable markups, assuming that productivities are Pareto-distributed. They show empirically that gains from trade are lower when markups are variable rather than constant. While our model does not admit a sufficient statistic, we ask a related but different question: Are the gains from trade higher or lower under oligopoly, which features variable markups, compared to monopolistic competition with constant markups? We find significantly larger welfare gains under oligopoly. This is in line with Edmond, Midrigan, and Xu (2015) who calibrate a two-country version of the oligopoly model of Atkeson and Burstein (2008) to assess the gains from trade.[7]

In the last decade, there has been a revived interest in integrating oligopoly into models of international trade, partially building on earlier contributions by the strategic trade policy literature (see Brander, 1995). The by-now dominant framework, which we also adopt, was proposed by Neary (2003) and further developed by Atkeson and Burstein (2008). It features a continuum of oligopolistic industries, implying that firms have market power in their own industry but not in the aggregate. Quantitative papers that build on this framework include Edmond, Midrigan, and Xu (2015) and Gaubert and Itskhoki (2021).[8]

The rest of the paper is organized as follows. In Section 2, we present our theoretical framework and derive oligopoly correction terms for firm- and industry-level gravity equations. In Section 3, we describe the data sources, discuss estimation challenges, and present the empirical results from our firm-level gravity estimations. In Section 4, we repeat these steps for our industry-level gravity estimations. In Section 5, we provide Monte Carlo simulations to evaluate the performance of our estimation procedures. In Section 6, we use a calibrated version of our model to study the welfare gains from trade-cost reductions. Finally, we conclude in Section 7.

2 Gravity Equations under Oligopoly

In this section, we first present the oligopoly model underlying our approach to gravity with granular firms. Next, we derive gravity equations at both the firm and industry level.

trade in models that do not fit the assumptions of Arkolakis, Costinot, and Rodriguez-Clare (2012).

[7] Heid and Staehler (2020) propose an extension of Arkolakis, Costinot, and Rodriguez-Clare (2012)'s formula to evaluate the gains from trade under oligopoly. To recover the necessary parameters, they derive a firm-level gravity equation in oligopoly. However, they estimate it from aggregate trade data, assuming the economy consists of a large number of identical industries, each of which hosts only one firm per country. They also find that the welfare gains from trade liberalization are substantially larger under oligopoly.

[8] Other papers introducing oligopoly into international trade models include Eckel and Neary (2010), Parenti (2018), Head and Mayer (2019), and Breinlich, Nocke, and Schutz (2020).

2.1 Theoretical Framework

We consider a multi-country world with a continuum of industries, indexed by z.[9] The representative consumer in country n maximizes

$$U_n = \int_{z \in Z} \alpha_n(z) \log \left[\sum_{j \in \mathcal{J}_n(z)} a_{jn}(z)^{\frac{1}{\sigma(z)}} q_{jn}(z)^{\frac{\sigma(z)-1}{\sigma(z)}} \right]^{\frac{\sigma(z)}{\sigma(z)-1}} dz,$$

where $\alpha_n(z)$ denotes the industry-z expenditure share in country n, $\mathcal{J}_n(z)$ is the set of products available in industry z and country n, and $\sigma(z)$ denotes the elasticity of substitution between products in industry z. Consumption of product j in country n is denoted $q_{jn}(z)$. The utility shifter $a_{jn}(z)$ captures quality differences or other factors such as brand appeal. The resulting direct and inverse demands for product $i \in \mathcal{J}_n(z)$ in country n are given by

$$q_{in}(z) = a_{in}(z) p_{in}(z)^{-\sigma(z)} P_n(z)^{\sigma(z)-1} \alpha_n(z) E_n \tag{1}$$

and

$$p_{in}(z) = a_{in}(z)^{\frac{1}{\sigma(z)}} q_{in}(z)^{-\frac{1}{\sigma(z)}} Q_n(z)^{-\frac{\sigma(z)-1}{\sigma(z)}} \alpha_n(z) E_n,$$

where E_n is total expenditure in country n, and

$$P_n(z) \equiv \left[\sum_{j \in \mathcal{J}_n(z)} a_{jn}(z) p_{jn}(z)^{1-\sigma(z)} \right]^{\frac{1}{1-\sigma(z)}} \quad \text{and} \quad Q_n(z) \equiv \left[\sum_{j \in \mathcal{J}_n(z)} a_{jn}(z)^{\frac{1}{\sigma(z)}} q_{jn}(z)^{\frac{\sigma(z)-1}{\sigma(z)}} \right]^{\frac{\sigma(z)}{\sigma(z)-1}}$$

are the industry-z CES price index and composite commodity in country n, respectively. From now on, we focus on a single industry and drop the index z.

Each product $i \in \mathcal{J}_n$ is offered by a different firm, which may be either a domestic or foreign producer. Firms compete in quantities in each market n, being able to segment markets perfectly.[10] The profit of the firm offering product i from selling in destination n is $\pi_{in} = p_{in} q_{in} - C_{in}(q_{in})$, where $C_{in}(q_{in})$ is the firm's cost of producing and selling output q_{in}.

[9]The framework we lay out here can be viewed as being general equilibrium. However, as we focus on a given equilibrium and do not conduct comparative statics at the aggregate level, we refrain from explicitly closing the model by writing down factor market-clearing conditions and endogenizing consumer income. Closing the model would be straightforward—for example, by assigning a labor endowment to each country, assuming that all costs are incurred in terms of origin-country labor, choosing labor in a reference country as the numeraire, and assuming that profits and tariff revenues are distributed lump sum to domestic consumers.

[10]We focus on quantity competition here and present results for price competition in Appendix C.

We allow for variable returns to scale and use the following functional form:

$$C_{in}(q_{in}) = \frac{1}{1+\gamma} c_{in} (\tilde{\tau}_{in} q_{in})^{1+\gamma} = \frac{1}{1+\gamma} c_{in} \tau_{in} q_{in}^{1+\gamma},$$

where c_{in} is a firm-destination cost shifter and $\tilde{\tau}_{in}$ a firm-destination trade cost of the iceberg type.[11] We assume throughout that the returns-to-scale parameter γ satisfies $\gamma > -1/\sigma$, which means that the marginal cost of production should not decrease too fast with output. This (weak) assumption guarantees that all the profit functions we consider are unimodal.

Under oligopoly, firms take into account the impact of their quantity choices on the CES-composite, Q_n. For what follows, it is useful to generalize further the degree of strategic interaction between firms by introducing a conduct parameter, λ (see Bresnahan, 1989): When firm i increases its output q_{in} by an infinitesimal amount, it perceives the induced effect on Q_n to be equal to $\lambda \partial Q_n / \partial q_{in}$. Under monopolistic competition, the conduct parameter λ takes the value of zero, whereas it is equal to one under Cournot competition. The first-order condition of profit maximization of firm i in destination n is given by

$$\begin{aligned} 0 = \frac{\partial \pi_{in}}{\partial q_{in}} &= \frac{\alpha_n E_n}{Q_n^{\frac{\sigma-1}{\sigma}}} a_{in}^{\frac{1}{\sigma}} \frac{\sigma-1}{\sigma} q_{in}^{-\frac{1}{\sigma}} - \frac{\sigma-1}{\sigma} \lambda \frac{\partial Q_n}{\partial q_{in}} \frac{\alpha_n E_n a_{in}^{\frac{1}{\sigma}} q_{in}^{\frac{\sigma-1}{\sigma}}}{Q_n^{\frac{\sigma-1}{\sigma}+1}} - C'_{in}(q_{in}) \\ &= \frac{\sigma-1}{\sigma} p_{in} \left(1 - \lambda s_{in}\right) - C'_{in}(q_{in}), \end{aligned} \tag{2}$$

where

$$s_{in} \equiv \frac{a_{in}^{\frac{1}{\sigma}} q_{in}^{\frac{\sigma-1}{\sigma}}}{\sum_{j \in \mathcal{J}_n} a_{jn}^{\frac{1}{\sigma}} q_{jn}^{\frac{\sigma-1}{\sigma}}} \tag{3}$$

is the market share of firm i in destination n.

Rearranging terms in equation (2) yields firm i's optimal markup in destination n:

$$\mu_{in} = \frac{1}{\sigma} + \lambda \frac{\sigma-1}{\sigma} s_{in}, \tag{4}$$

where $\mu_{in} \equiv \left(p_{in} - C'_{in}(q_{in})\right) / p_{in}$ is the Lerner index. Under monopolistic-competition conduct ($\lambda = 0$), the usual constant markup $1/\sigma$ obtains. If instead $\lambda > 0$, then markups are no longer constant and depend positively on market shares. We will make use of the additional flexibility afforded by the conduct parameter λ in Section 2.3, but for now, we assume Cournot conduct and set $\lambda = 1$.

[11] For one unit of the output to arrive in destination n, the firm needs to ship $\tilde{\tau}_{in}$. Note that we define $\tau_{in} = \tilde{\tau}_{in}^{1+\gamma}$ to ease the subsequent notation.

2.2 Firm-Level Gravity in Oligopoly

From the definition of the Lerner index, firm i's price in market n is $p_{in} = c_{in}\tau_{in}q_{in}^{\gamma}/(1-\mu_{in})$. Using equation (1), the value of its sales can be written as

$$r_{in} = p_{in}q_{in} = \left(\frac{c_{in}\tau_{in}}{1-\mu_{in}}\right)^{\frac{1-\sigma}{1+\sigma\gamma}} \left(a_{in}P_n^{\sigma-1}E_n\right)^{\frac{1+\gamma}{1+\sigma\gamma}}. \quad (5)$$

We log-linearly decompose the quality- and cost-shock terms as $\log a_{in} = \varepsilon_i^a + \varepsilon_n^a + \varepsilon_{in}^a$ and $\log c_{in} = \varepsilon_i^c + \varepsilon_n^c + \varepsilon_{in}^c$, respectively. We further decompose trade costs as $\log \tau_{in} = \beta X_{in} + \varepsilon_i^{\tau} + \varepsilon_n^{\tau} + \varepsilon_{in}^{\tau}$ where X_{in} includes variables with bilateral variation such as (log) distance, common language, or dummies for the presence of trade agreements or currency unions. Obtaining consistent estimates of the coefficients on these bilateral terms (β) is a key objective of much of the gravity literature.

Taking the logarithm of equation (5) yields the firm-level gravity equation

$$\log r_{in} = \xi_n + \zeta_i + \beta\frac{1-\sigma}{1+\sigma\gamma}X_{in} + \frac{\sigma-1}{1+\sigma\gamma}\log\left(1-\mu_{in}\right) + \varepsilon_{in}, \quad (6)$$

where ξ_n and ζ_i summarize destination- and firm-specific terms, and

$$\varepsilon_{in} = \frac{1}{1+\sigma\gamma}\left[\left(1+\gamma\right)\varepsilon_{in}^a + \left(1-\sigma\right)\left(\varepsilon_{in}^c + \varepsilon_{in}^{\tau}\right)\right].$$

Note that under the assumption of monopolistic competition, the markup term involving μ_{in} would be constant and could be subsumed in ζ_i. In that case, estimation of (6) would yield consistent estimates of the coefficient on X_{in}, provided that firm and destination fixed effects (ζ_i and ξ_n) are controlled for and that the usual identifying assumptions made in the gravity literature hold.[12]

Under oligopoly, however, the markup term will depend on firms' market shares and will thus be correlated with the regressors of interest, X_{in}, resulting in an omitted variable bias. For example, to the extent that firms face larger variable trade costs in more-distant markets, their market shares are lower there, ceteris paribus. Hence, firms charge lower markups in

[12] For least-squares estimation of the log-linearized gravity equation, the identifying assumption is $\mathbb{E}\left(\varepsilon_{in}|X_{in},\xi_n,\zeta_i\right) = 0$. This assumption does not rule out correlations between the bilateral variables and taste, production and trade cost shocks working through the firm- and destination-level components (ε_i^a, ε_n^a, ε_i^c, ε_n^c, ε_i^{τ} and ε_n^{τ}). Such correlations are not a problem, as these components can be controlled for through firm and destination fixed effects. If the data contain a time dimension, one can also allow for time-invariant bilateral elements in the error term which can be captured through bilateral fixed effects—as is standard, e.g., in the literature on the trade effects of preferential trade agreements (e.g., Baier and Bergstrand, 2007).

such destinations, implying a positive correlation between distance and the omitted variable, $\log(1-\mu_{in})$. Importantly, as markups vary at the firm-destination level, their variation cannot be controlled for by firm and destination fixed effects.[13]

We propose to solve the omitted variable problem by constructing a proxy for the markup term in (6). Specifically, with estimates for σ and γ and data for s_{in}, we can compute

$$\widehat{\mu}_{in} = \frac{1}{\widehat{\sigma}} + \frac{\widehat{\sigma}-1}{\widehat{\sigma}} s_{in}$$

and estimate

$$\log \widetilde{r}_{in} = \xi_n + \zeta_i + \beta \frac{1-\sigma}{1+\sigma\gamma} X_{in} + \varepsilon_{in}, \tag{7}$$

where

$$\log \widetilde{r}_{in} \equiv \log r_{in} - \frac{\widehat{\sigma}-1}{1+\widehat{\sigma}\widehat{\gamma}} \log\left(1-\widehat{\mu}_{in}\right). \tag{8}$$

The right-hand side of equation (7) is the same as in the standard firm-level gravity equation. The novelty is that the trade flows on the left-hand side have been purged from oligopolistic market power effects, as detailed in equation (8). Given our earlier identifying assumptions, using $\log \widetilde{r}_{in}$ instead of $\log r_{in}$ as the dependent variable yields a consistent estimate of $\beta \frac{1-\sigma}{1+\sigma\gamma}$. Using again the estimates for σ and γ then allows recovering the parameter of interest, β.[14]

2.3 Industry-Level Gravity in Oligopoly

We now turn to gravity at the industry level. We first analyze the equilibrium in a given market using an aggregative games approach (Nocke and Schutz, 2018b; Anderson, Erkal, and Piccinin, 2020). We then leverage Nocke and Schutz (2018a)'s approximation techniques to derive an industry-level gravity equation that accounts for oligopolistic behavior.

An aggregative games approach to industry equilibrium. Consider industry z in destination n. Dropping reference to z to ease notation, we define the market-level aggregator H_n as

$$H_n \equiv Q_n^{\frac{\sigma-1}{\sigma}} = \sum_{j\in\mathcal{J}_n} a_{jn}^{\frac{1}{\sigma}} q_{jn}^{\frac{\sigma-1}{\sigma}}$$

[13]The inclusion of firm-destination fixed effects would make it impossible to identify separately the effect of key regressors of interest such as distance, tariffs or dummy variables for trade agreements. Having a time dimension in the data would not help either because markups would then vary by firm, destination, and time.

[14]Note the parallel to the literature on trade and quality which uses a similar approach to correct export values or quantities (e.g., Khandelwal, Schott, and Wei, 2013).

and firm i's type T_{in}, a measure of quality-adjusted productivity, as

$$T_{in} \equiv a_{in}^{\frac{1}{\sigma}} \left(\frac{\alpha E}{c_{in}\tau_{in}} \frac{\sigma - 1}{\sigma} \right)^{\frac{\sigma-1}{\sigma(1+\gamma)}}. \tag{9}$$

Plugging these definitions into equation (2), using equation (3), and rearranging, we obtain:

$$1 - \lambda s_{in} = s_{in}^{\frac{1+\sigma\gamma}{\sigma-1}} \left(\frac{H}{T_{in}} \right)^{\frac{\sigma(1+\gamma)}{\sigma-1}}, \tag{10}$$

where λ is the conduct parameter introduced in Section 2.2. As the left-hand side is non-increasing in s_{in} and the right-hand side is strictly increasing in s_{in}, the equation has a unique solution in s_{in}, denoted $S(T_{in}/H_n, \lambda)$—the *market-share fitting-in function*. It is easily verified that $S(\cdot, \cdot)$ is strictly increasing in its first argument and strictly decreasing in its second.

The equilibrium level of the aggregator, $H^*(\lambda)$, is pinned down by market shares having to add up to unity:

$$\sum_{i \in \mathcal{J}_n} S\left(\frac{T_{in}}{H_n}, \lambda \right) = 1. \tag{11}$$

The uniqueness of the solution follows by the strict monotonicity of the market-share fitting-in function.

To summarize:

Proposition 1. *In each destination market n, and for any conduct parameter λ, there exists a unique equilibrium in quantities. The equilibrium aggregator level $H_n^*(\lambda)$ is the unique solution to equation* (11)*. Each firm i's equilibrium market share is $s_{in}^*(\lambda) = S(T_{in}/H_n^*(\lambda), \lambda)$, where $S(T_{in}/H_n^*(\lambda), \lambda)$ is the unique solution to equation* (10).

Proof. See Appendix A.1 □

The first-order approach to industry-level gravity. Let $\mathcal{J}_{on} \subsetneq \mathcal{J}_n$ denote the subset of exporters from country o that sell in the destination market n. Their aggregate exports to market n are given by $s_{on}^* \alpha E_n$, where $s_{on}^*(\lambda) \equiv \sum_{i \in \mathcal{J}_{on}} s_{in}^*(\lambda)$. We are interested in these aggregate exports when firms compete in a Cournot fashion, i.e., when $\lambda = 1$. Unfortunately, there is no closed-form solution to $s_{on}^*(1)$. Our approach therefore entails approximating it.

As we show in the following, the approximation relies on two Herfindahl indices, namely the HHI of all firms selling in the destination market n,

$$\mathrm{HHI}_n(\lambda) \equiv \sum_{j \in \mathcal{J}_n} \left(s^*_{jn}(\lambda)\right)^2,$$

and the HHI among the exporters in country o that sell in the destination market n,

$$\mathrm{HHI}_{on}(\lambda) \equiv \sum_{j \in \mathcal{J}_{on}} \left(\frac{s^*_{jn}(\lambda)}{s^*_{on}(\lambda)}\right)^2.$$

We obtain:

Proposition 2. *At the first order, in the neighborhood of $\lambda = 0$ (monopolistic-competition conduct), the logged joint market share in destination n of the firms from origin o is given by*

$$\log s^*_{on}(\lambda) = \log s^*_{on}(0) + \frac{\sigma - 1}{1 + \sigma\gamma}\Big[\mathrm{HHI}_n(\lambda) - s^*_{on}(\lambda)\,\mathrm{HHI}_{on}(\lambda)\Big]\lambda + o(\lambda).$$

Proof. See Appendix A.2. □

The proposition shows that the joint market share of the exporters from country o differs from the one that would obtain under monopolistic competition by a market-power term that takes account of both the overall concentration in the destination market and the concentration among the country-o exporters.

This result motivates the following approximation:

$$\log s^*_{on}(0) \simeq \log s^*_{on}(1) - \frac{\sigma - 1}{1 + \sigma\gamma}\Big[\mathrm{HHI}_n(1) - s^*_{on}(1)\,\mathrm{HHI}_{on}(1)\Big].$$

Thus, the export flow that would obtain under monopolistic competition is approximately given by

$$\underbrace{\log(\alpha_n E_n) + \log s^*_{on}(1)}_{\log r_{on}} - \frac{\sigma - 1}{1 + \sigma\gamma}\Big[\mathrm{HHI}_n(1) - s^*_{on}(1)\,\mathrm{HHI}_{on}(1)\Big],$$

where r_{on} is the (actually observed) export flow under oligopoly. As the HHI_n term will be subsumed in the destination fixed effect, we define

$$\log \widetilde{r}_{on} \equiv \log r_{on} + \frac{\sigma - 1}{1 + \sigma\gamma} s_{on}\,\mathrm{HHI}_{on} \tag{12}$$

as the value of the export flow from o to n purged from market-power effects, which can be computed with data on r_{on}, s_{on} and HHI_{on}, and estimates of σ and γ.

Next, we derive a gravity equation for oligopoly-corrected trade flows $\log \widetilde{r}_{on}$. To do so, we impose the following structure on the shocks to quality, marginal costs and trade costs:

$$\log a_{in} = \log a_i + \varepsilon_o^a + \varepsilon_n^a + \varepsilon_{on}^a,$$
$$\log c_{in} = \log c_i + \varepsilon_o^c + \varepsilon_n^c + \varepsilon_{on}^c,$$
$$\log \tau_{in} = \beta X_{on} + \varepsilon_o^\tau + \varepsilon_n^\tau + \varepsilon_{on}^\tau.$$

Combining this with equation (5) (with $\mu_{in} = 1/\sigma$), adding up over all the exporters from o to n, and taking the logarithm, yields a gravity equation of the following form:

$$\log \widetilde{r}_{on} = \xi_o + \zeta_n + \beta \frac{1-\sigma}{1+\sigma\gamma} X_{on} + \phi_{on} + \varepsilon_{on}, \tag{13}$$

where ξ_o is an origin fixed effect, ζ_n is a destination fixed effect,

$$\phi_{on} \equiv \log \sum_{j \in \mathcal{J}_{on}} a_j^{\frac{1+\gamma}{1+\sigma\gamma}} c_j^{\frac{1-\sigma}{1+\sigma\gamma}},$$

and

$$\varepsilon_{on} \equiv \frac{1+\gamma}{1+\sigma\gamma} \varepsilon_{on}^a + \frac{1-\sigma}{1+\sigma\gamma} (\varepsilon_{on}^c + \varepsilon_{on}^\tau).$$

If the set of exporters from origin o were the same in all destinations n (i.e., if $\mathcal{J}_{on}$ were independent of n), then the term ϕ_{on} would be subsumed into the origin fixed effect. In that case, regressing $\log \widetilde{r}_{on}$ on origin and destination fixed effects, and the bilateral variables X_{on} would yield a consistent estimate of $\beta(1-\sigma)/(1+\sigma\gamma)$, provided the usual identifying assumption $\mathbb{E}(\varepsilon_{on} | X_{on}, \xi_o, \zeta_n) = 0$ holds.

If instead, the set $\mathcal{E}_{on}$ does depend on n because of self-selection into export destinations, then ϕ_{on} is no longer absorbed by the origin fixed effect, and is likely to be correlated with X_{on}. In Section 4 below, we discuss how to address such self-selection issues and still obtain a consistent estimate of the "intensive-margin effect", $\beta(1-\sigma)/(1+\sigma\gamma)$.[15]

[15] We choose to focus on the intensive-margin effect for the following reasons. First, this is also the effect of interest at the firm level. Second, the overall effect of bilateral trade costs on trade flows, which would include the extensive-margin effect through self-selection into exporting, is generally not constant in this model, due to oligopoly and productivities not being Pareto-distributed.

3 Empirical Implementation: Firm-Level Gravity

In this section, we show how to empirically implement our gravity-estimation approach at the firm level. Our empirical specification is

$$\log \widetilde{r}_{inzt} = \xi_{nzt} + \zeta_{izt} + \beta \frac{1-\sigma_z}{1+\sigma_z \gamma_z} X_{in} + \varepsilon_{inzt}, \tag{14}$$

where

$$\log \widetilde{r}_{inzt} \equiv \log r_{inzt} - \frac{\widehat{\sigma}_z - 1}{1+\widehat{\sigma}_z \widehat{\gamma}_z} \log\left(1 - \widehat{\mu}_{inzt}\right), \tag{15}$$

corresponding to equations (7) and (8) in Section 2.2 above, except that we have made here the industry (z) and time (t) dimensions explicit. Due to data limitations explained below, we focus on distance as our only gravity variable, so that X_{in} boils down to the scalar $\log(\text{dist}_{on})$, where dist_{on} is the distance between firm i's origin o and destination n. In the following, we discuss estimation challenges, present our data, run gravity regressions with and without oligopoly correction, and investigate under what circumstances ignoring oligopolistic behavior leads to quantitatively important biases.

Estimation Challenges. A first issue is how to control for destination fixed effects ξ_{nzt} in a setting with firm-level export data. With export data from a single origin country, we would not be able to separate the impact of bilateral variables from the destination fixed effects.[16] To address this issue, we follow Bas, Mayer, and Thoenig (2017) by combining two data-sets on the exports of French and Chinese firms, respectively.

Secondly, we have to address self-selection issues, as most firms export only to a subset of possible destinations. When estimating equation (14), observations with zero trade flows drop out. In the presence of export fixed cost, there is selection into exporting in our model: Firms selling in more distant foreign markets will be more likely to have received a favorable taste, productivity, or trade-cost shock for that destination, allowing them to operate in this more difficult environment. As a consequence, the conditional expectation $\mathbb{E}\left(\varepsilon_{inzt} | X_{in}, r_{inzt} > 0\right)$ is likely to be different from zero. To address this issue, we adapt an approach proposed by Bas, Mayer, and Thoenig (2017) and restrict our estimation sample to the largest three French and Chinese firms in each industry as measured by total industry-level exports, added up over all destinations. As those firms are generally very productive, produce high-quality products (low ε^c_{izt} and/or high ε^a_{izt}), or use low-cost market-access technologies (low ε^τ_{izt}),

[16] For example, if we used data on the exports of French firms only, we would not be able to distinguish whether firms' exports to a given destination are high because France and the country in question are close to each other or because of other destination-specific factors such as a high price index or expenditure level.

they are likely to serve most destinations, so that the destination-specific shocks (ε^c_{inzt}, ε^a_{inzt}, and ε^τ_{inzt}) do not play an important role in their market entry decisions. We acknowledge that this is an imperfect solution but our simulation evidence presented in Section 5 shows that focusing on top exporters does indeed substantially reduce selection bias. Moreover, we show that our results are very similar when using only the top exporter or top-5 exporters from France and China; see footnote 23 below.

Third, as shown by Santos Silva and Tenreyro (2006), in the presence of heteroscedasticity the log-linearized gravity equation yields inconsistent estimates of $\mathbb{E}(\widetilde{r}_{inzt}|X_{in}, \xi_{nzt}, \zeta_{izt})$, where

$$\widetilde{r}_{inzt} = \exp(\xi_{nzt} + \zeta_{izt} + \beta\frac{1-\sigma_z}{1+\sigma_z\gamma_z}X_{in})\exp(\varepsilon_{inzt}). \tag{16}$$

To see this, note that if $\text{Var}(\exp(\varepsilon_{inzt})|X_{in})$ depends on X_{in}, then so does $\mathbb{E}(\varepsilon_{inzt}|X_{in})$. A solution to this problem is to estimate the gravity equation (16) by PPML in multiplicative form, which also allows us to include zero trade flows in our estimation sample. Recent computational advances in PPML estimation (e.g., Correia, Guimaraes, and Zylkin, 2019) make it possible to include the large number of fixed effects required in our setting.

Finally, the oligopoly correction term for firm-level gravity (see equation (15)) requires estimates of σ and γ. For our main empirical specifications, we follow the literature and assume constant returns to scale ($\gamma = 0$) and set σ equal to 5, a typical value from the literature. However, we will also obtain estimates of σ and γ by adapting the estimation procedure of Feenstra (1994) and Broda and Weinstein (2006) to firm-level data and oligopolistic competition; see Appendix B for details.[17]

Data. We use annual firm-level export data for French and Chinese exporters provided by the two countries' customs authorities for the years 2000–2010. In each data-set, we observe all the products a firm exports and all the destinations it serves, and the quantity and value of the underlying flows. Although both data-sets record export data at the 8-digit level, we need to aggregate this information up to the 6-digit level of the Harmonised System (HS), which is the most disaggregate level at which the two national classifications are comparable. Using export values and quantities, we compute unit values—a commonly used proxy for prices in the trade literature.

To compute market shares, defined as the ratio of export value to absorption, we combine our firm-level data with absorption data at the HS 6-digit level (or close to it) from Eurostat's

[17] If our data exhibited tariff variation by destination, we could instead adopt the approach of Head and Mayer (2019) to estimate σ directly from the gravity equation.

PRODCOM database.[18] The downside of using PRODCOM is that absorption data is only available for European countries. As a result, there is insufficient variation to include, in addition to distance, regressors such as dummies for common language or policy-related variables (e.g., membership in a free-trade agreement and bilateral tariffs).[19]

After combining our data sources, we end up with information on export values, export quantities, and market shares for 31 European destinations, 1,864 industries and around 250,000 exporters for the period 2000–2010.[20,21] We source information on bilateral distance from CEPII.

Descriptive Statistics. The key determinants of our oligopoly correction term are firm-level market shares, and estimates for demand elasticities (σ) and returns to scale (γ). Column (1) of Table 1 presents information on the firm-level market shares for the French and Chinese exporters. The average market share across the approximately 14 million firm-destination-industry-year combinations in our data is small (0.4%) and the median is even smaller (around 0.01%). At the 95th percentile, the firm-level market share is 1.12%. Clearly, the typical firm in our data does not enjoy much market power.

However, this does not imply that correcting firm-level exports for oligopoly forces will not matter quantitatively, as estimation results could be substantially biased by a small number of exporters with high market shares. Columns (2)–(4) focus on such firms. Column (2) shows descriptive statistics for the top exporters (i.e., the French and Chinese firms with the largest total export values for a given 6-digit industry and year). The average top-exporter

[18] Absorption, defined as domestic production + imports − exports, is the counterpart to $\alpha_n(z)E_{nt}$ in our model. In principle, this information is available at the HS 6-digit level but issues such as classification changes over time often require aggregation to higher levels. The original classification of the PRODCOM data is the 8-digit CN classification, which changes almost every year. We apply the procedure developed by Van Beveren, Bernard, and Vandenbussche (2012) to map the CN classification to an artificial HS classification, "HS 6-digit plus", that is comparable over time and compatible with the 6-digit HS classification. The idea is to aggregate both trade and PRODCOM data as little as possible and as much as required to guarantee a one-to-one mapping between them. See their paper for an in-depth discussion of the procedure.

[19] The destination countries in our sample were either EU member states or had implemented free-trade agreements (FTAs) with the EU before 2000 and therefore had no tariffs on EU imports. By contrast, China did not have any FTAs with countries in our sample and EU external tariffs for imports from China only had variation across industries. Thus, all variation in the FTA dummy or in tariffs would be absorbed by our firm-industry-year fixed effects. Likewise, there is insufficient variation to include an indicator for common language.

[20] The export destinations are: Austria, Belgium, Bulgaria, Croatia, Cyprus, Czechia, Denmark, Estonia, Finland, Germany, Greece, Hungary, Iceland, Ireland, Italy, Latvia, Lithuania, Luxembourg, Malta, the Netherlands, Norway, Poland, Portugal, Romania, Serbia, Slovakia, Slovenia, Spain, Sweden, Turkey, and the United Kingdom.

[21] Possibly because of measurement issues in PRODCOM, we occasionally observe instances where absorption is smaller than a firm's export value, resulting in market shares larger than one; in such cases, we winsorize market shares to 0.95.

market share is around 6%, substantially larger than the average exporter's market share. Moreover, at the 95th percentile the top firm enjoys a market share of almost 30%. Column (3) presents information on the market shares of the top-3 exporters. The mean market share in this sample is around 3.9%; at the 95th percentile, it is around 18%. Column (4) shows the cumulative market shares of the top-3 exporters. The average cumulative market share is equal to 7.30%; at the 95th percentile, it is 33.4%. In sum, in a significant minority of destination markets, the largest French and Chinese exporters enjoy substantial market shares.

Table 1: Summary Statistics for the Market Shares of French and Chinese Exporters

	(1) All Exporters	(2) Top Exporters	(3) Top 3 Exporters	(4) Top 3 Exporters (Cumulative)
Mean	0.40%	6.00%	3.88%	7.30%
5th pctile	0.00007%	0.01%	0.006%	0.03%
10th pctile	0.0004	0.03	0.02%	0.09%
Median	0.01%	1.21%	0.65%	2.05%
90th pctile	0.44%	15.72%	9.20%	19.36%
95th pctile	1.12%	28.96%	18.04%	33.44%
Observations	14,009,005	276,718	708,409	708,409

Notes: Table shows summary statistics for strictly positive market shares of French and Chinese exporters for the years 2000–2010. The unit of observation is at the firm-destination-industry-year level.

Table 2 shows descriptive statistics for our estimates of σ and γ. We constrain coefficients to be identical within 2-digit HS sectors to guarantee a sufficient number of observations underlying each estimate. For the average and median sector, we estimate mildly decreasing returns to scale of $\gamma = 0.34$ and $\gamma = 0.19$, respectively. For our price elasticity estimates, we find a mean of $\sigma = 5.39$ and a median of $\sigma = 3.74$. These numbers are similar to estimates at comparable levels of aggregation in the literature (e.g., Broda and Weinstein, 2006).

Estimation Results. All regressions are run on the sample of the top-3 French and Chinese exporters for each 6-digit HS industry. As a first step, we pool all industries and years and estimate equation (14) by OLS and equation (16) by PPML. We do so both on actual export flows and on oligopoly-corrected export flows, as defined in equation (15).

For our main empirical specification, we assume constant returns to scale ($\gamma = 0$) and set $\sigma = 5$ to construct the oligopoly correction term.[22] The estimation results are reported in

[22] $\sigma = 5$ is a "conventional" value in the international trade literature (see, e.g., Gaubert and Itskhoki, 2021). Estimates of σ are often in the range from 4 to 6; for example, Gaubert and Itskhoki (2021) obtain

Table 2: Price Elasticities and Returns-to-Scale Estimates

	σ	γ
Mean	5.39	0.34
25th Percentile	2.22	0.03
Median	3.74	0.10
75th Percentile	7.50	0.30
Min	1.01	-0.13
Max	26.07	4.46
Standard Deviation	4.07	0.69
Observations	78	78

Notes: Table shows descriptive statistics for estimates of σ and γ across 2-digit HS sectors. Estimates computed using 6-digit HS firm-level information but constrained to be identical within 2-digit HS sectors.

Table 3: Firm-Level Gravity Estimates, $\sigma = 5$ and $\gamma = 0$.

	(1)	(2)	(3)	(4)
Method	PPML w/o corr	PPML w/ corr	OLS w/o corr	OLS w/ corr
log distance	-0.874***	-1.518***	-0.232***	-0.275***
	(0.221)	(0.020)	(0.014)	(0.013)
$\hat{\beta}_{distance}$	0.218	0.379	0.058	0.069
Observations	11,955,786	11,955,786	708,392	708,386
(Pseudo) R-squared	0.14	0.28	0.06	0.05
Firm-industry-year FE	YES	YES	YES	YES
Industry-dest.-year FE	YES	YES	YES	YES

Notes: Firm-level data, pooled across industries and years. Results for top 3 exporters. Oligopoly correction with $\sigma = 5$ and $\gamma = 0$. Standard errors in parentheses, clustered at the destination-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Firm-Level Gravity Estimates, $\sigma = 5.39$ and $\gamma = 0.34$.

	(1)	(2)	(3)	(4)
Method	PPML w/o corr	PPML w/ corr	OLS w/o corr	OLS w/ corr
log distance	-0.874***	-1.201***	-0.232***	-0.248***
	(0.021)	(0.118)	(0.014)	(0.014)
$\hat{\beta}_{distance}$	0.577	0.792	0.149	0.160
Observations	11,955,786	11,955,786	708,392	708,386
(Pseudo) R-squared	0.13	0.19	0.06	0.05
Firm-industry-year FE	YES	YES	YES	YES
Industry-dest.-year FE	YES	YES	YES	YES

Notes: Firm-level data, pooled across industries and years. Results for top 3 exporters. Oligopoly correction with mean of estimated σ and γ. Standard errors in parentheses, clustered at the destination-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3. Columns (1) and (2) present the PPML estimates for the specification without and with the oligopoly correction term, respectively. Without oligopoly correction, the distance coefficient is strongly biased towards zero: The point estimate is -1.52 in column (2) but only -0.87 in column (1). This confirms our theoretical insight that the trade elasticity with respect to distance suffers from a substantial attenuation bias because firms systematically reduce their markups in markets where they face higher variable trade costs and thus have lower market shares. As a consequence, export values decrease by less than they would have decreased under constant markups. The coefficient on distance corresponds to $\beta(1-\sigma)/(1+\sigma\gamma)$. With $\sigma = 5$ and $\gamma = 0$, the implied estimated values of the distance elasticity of trade costs, $\widehat{\beta}$, are 0.218 in column (1) and 0.359 in column (2), implying a downward bias of around 44%.

Columns (3) and (4) report the OLS estimates for the specification without and with the oligopoly correction term, respectively. In both cases, the absolute value of the estimated distance coefficient is significantly lower than with PPML, suggesting a substantial heteroscedasticity bias. The presence of such a bias is confirmed in our Monte Carlo simulations in Section 5; see Table 12.

Table 4 is the analog of Table 3, but using the mean estimated σ and γ for the construction of the oligopoly correction term. In our preferred specification (column (2), PPML with oligopoly correction), the absolute value of the distance coefficient is now slightly smaller (1.2), but still substantially larger than without oligopoly correction, implying an oligopoly bias of 30%. The implied distance elasticity of trade costs, $\widehat{\beta}$, is now larger than in Table 3 (0.792 vs. 0.359), despite the slightly smaller estimated distance coefficient; this is entirely driven by the fact that the estimated returns-to-scale parameter γ is substantially larger than zero.[23]

We now turn to estimating gravity equations separately for each of the 78 HS 2-digit sectors, pooling observations across 6-digit industries within a given 2-digit sector. (Recall that, in the specifications in which σ and γ are estimated, those parameters are allowed to vary at the sector level.) Table 5 reports summary statistics on the distribution of the distance coefficient estimate (with and without oligopoly correction) and oligopoly bias across sectors. We employ three alternative specifications: (1) $\sigma = 5$ and $\gamma = 0$; (2) σ estimated and $\gamma = 0$; and (3) (σ, γ) estimated. In all specifications, the median point estimate on

an "imprecisely estimated" $\sigma = 4.927$, while Breinlich, Nocke, and Schutz (2020) calibrate $\sigma = 5.163$ in the median 5-digit industry.

[23] As shown in Table 19 in Appendix D, very similar results obtain when using (i) only the top French and Chinese exporter and (ii) the top-5 French and Chinese exporters.

distance is much larger when including the correction term. For our baseline specification with $\sigma = 5$ and $\gamma = 0$, the median point estimate of the distance coefficient is -1.35, which is close to the pooled estimate. By contrast, without the oligopoly correction the median distance coefficient is only -0.51. There is substantial variation in the absolute percentage oligopoly bias:[24] For a sector at the 10th percentile, this bias is around 10%, but it increases to 160% for a sector at the 90th percentile. Thus, in some industries, the oligopoly bias is much larger than the pooled estimates would suggest.

Table 5: Firm-level Gravity Estimates by 2-digit Sector

Median est coefficient	$\sigma = 5, \gamma = 0$	σ est, $\gamma = 0$	σ, γ est
log distance w/o corr	-0.508	-0.065	-0.081
log distance w/ corr	-1.347	-1.796	-0.740
$\hat{\beta}_{distance}$ w/o corr	0.127	0.013	0.045
$\hat{\beta}_{distance}$ w/ corr	0.337	0.779	0.445
abs. pct. bias (10th pctile)	10%	37.7%	8%
abs. pct. bias (median)	95.6%	100%	95.8%
abs. pct. bias (90th pctile)	160%	122%	165%

Notes: Firm-level data. Table shows summary statistics on the distribution of estimated coefficients by 2-digit HS sector for top-3 exporters.

4 Empirical Implementation: Industry-Level Gravity

We now show how to empirically implement our gravity-estimation approach at the industry level. Our empirical specification is

$$\log \widetilde{r}_{onz} = \xi_{oz} + \zeta_{nz} + \beta \frac{1 - \sigma_z}{1 + \sigma_z \gamma_z} \log(\text{dist}_{on}) + \eta_{onz}, \tag{17}$$

where

$$\log \widetilde{r}_{onz} = \log r_{onz} + \frac{\widehat{\sigma}_z - 1}{1 + \widehat{\sigma}_z \widehat{\gamma}_z} s_{onz} \, \text{HHI}_{onz} \,. \tag{18}$$

This corresponds to equations (13) and (12) above, with the industry (z) dimension being made explicit, $\log(\text{dist}_{on})$ being our only gravity variable (for the reasons explained in Section 3), and $\eta_{onz} \equiv \phi_{onz} + \varepsilon_{onz}$. We refrain from introducing a time index t because, in the estimations below, we confine attention to data from the year 2010 for computational reasons.[25]

[24] The absolute percentage bias is defined as the absolute value of $(\widehat{\beta}_{\text{w/o corr}} - \widehat{\beta}_{\text{w/ corr}})/\widehat{\beta}_{\text{w/ corr}}$.
[25] Similar results are obtained for other years.

As in the previous section, we now turn to a discussion of the estimation challenges. We then briefly describe our data, run gravity regressions with and without oligopoly correction, and investigate under what circumstances ignoring oligopolistic behavior leads to quantitatively important biases.

Estimation Challenges. The self-selection of firms into export markets again poses problems for a consistent estimation of the intensive-margin elasticity of trade with respect to distance, $\beta(1-\sigma)/(1+\sigma\gamma)$. If in each origin country o there were a single firm choosing whether to enter any given destination country n, we would have the same sample selection problem as at the firm level. While the ϕ_{onz}-term would be subsumed into the origin fixed effect, the conditional expectation $\mathbb{E}(\varepsilon_{onz}|\widetilde{r}_{onz} > 0, \xi_{oz}, \zeta_{nz}, \text{dist}_{on})$ would depend on dist_{on}: The observation that a firm is exporting to a remote market is likely to be the result of that firm having received a favorable ε_{onz}-shock. With multiple potential exporters, this problem remains. In addition to sample selection, however, a potential extensive-margin bias arises whenever the set of exporters from o varies with n, so that the ϕ_{onz}-term can no longer be subsumed into the origin fixed effect. In particular, the conditional expectation $\mathbb{E}(\phi_{onz}|\widetilde{r}_{onz} > 0, \xi_{oz}, \zeta_{nz}, \text{dist}_{on})$ is likely to depend on dist_{on}, as a larger number of firms would presumably find it profitable to export to nearby destinations. Summarizing, while the sample-selection bias tends to lead to an underestimation of the effect of distance on trade, the extensive-margin bias tends to result in an overestimation.

To alleviate both biases, we apply the two-step procedure developed by Helpman, Melitz, and Rubinstein (2008) (henceforth, HMR) to the oligopoly-corrected trade flows. The first step consists in estimating a Probit model of whether positive trade flows between o and n are observed. The regressors are origin and destination fixed effects, $\log \text{dist}_{on}$, and bilateral variables that are likely to affect the fixed export cost but not variable trade costs. For the latter variables, we follow HMR and use: (i) A dummy equal to one if business startup time is above the median in both o and n; and (ii) a dummy equal to one if business startup cost is above the median in both o and n. The estimated conditional probability of observing positive trade flows is denoted $\widehat{\rho}_{onz}$.

The second step consists in estimating the following gravity equation:

$$\log \widetilde{r}_{onz} = \zeta_{oz} + \xi_{nz} + \beta \frac{1-\sigma_z}{1+\sigma_z\gamma_z} X_{onz} + P(\log \hat{Z}_{onz}) + \omega \widehat{\lambda}_{onz} + \eta_{onz}, \tag{19}$$

where $\widehat{\lambda}_{onz}$ is the inverse Mills ratio from the first step, $\log \hat{Z}_{onz}$ is the $\widehat{\rho}_{onz}$-quantile of the standard normal distribution, and $P(\log \hat{Z}_{onz})$ is a polynomial in $\log \hat{Z}_{onz}$. The role of $\widehat{\lambda}_{onz}$

is to correct for sample selection, while $P(\log \hat{Z}_{onz})$ addresses the extensive-margin bias by non-parametrically controlling for ϕ_{onz}.

A downside of the HMR method is that it does not account for potential heteroscedasticity. We therefore also run PPML regressions, which however address neither sample selection nor the extensive-margin bias. The Monte Carlo simulations in Section 5 below indicate that self-selection issues are more severe than problems arising from potential heteroscedasticity.

Data and Descriptive Statistics. To make the estimation sample consistent with our firm-level regressions, we construct our industry-level data by aggregating our firm-level data (described in Section 3 above) to the 6-digit HS level.[26] We end up with information on export values, market shares and HHIs at the origin-destination-industry level for two exporting countries (France and China), 31 European destinations, and 1,864 industries for the year 2010. Data on business startup times and costs are sourced from the Worldbank's Doing Business Database.

Table 6 presents summary statistics on exporter HHIs and aggregate market shares of French and Chinese firms. It confirms that aggregate exports are concentrated among a small number of firms: The mean exporter HHI (which corresponds to HHI_{on} in equation (18)) is 0.55; at the 90th percentile, a single firm accounts for the total market share of each country. Moreover, the mean aggregate market share of French and Chinese firms in each destination (s_{on} in equation (18)) is around 9%; at the 90th percentile, that market share reaches 24%. Thus, in many markets, these exporters have substantial market power.

Estimation Results. Before presenting our results, it is worth pointing out that, as our industry-level data are constructed from firm-level data, we should expect to find estimates for the distance coefficient similar to those at the firm level.

We first present results for the pooled regressions where we constrain β to be the same across industries. Table 7 reports regression results with and without oligopoly correction, but without controlling for selection. The table focuses on the baseline case $\sigma = 5$ and $\gamma = 0$.

[26] Researchers may not generally have firm-level data-sets for several countries available. We have also experimented with using PRODCOM absorption data, combined with product-level trade data from Eurostat's COMEXT database. For the exporter HHIs, we have used the World Bank's Exporter Dynamics Database (EDD) which provides Herfindahl indices at the destination level, computed from firm-level export data for 48 exporting countries at the HS 2-digit level. Given that this is a relatively high degree of aggregation (90 aggregated manufacturing products), we have also experimented with computing 'pseudo-HHIs' based on 8-digit import data from COMEXT. While these data are at the product level rather than the firm level, they are highly disaggregated (ca. 9,600 different products). We thus need to make the assumption that each origin-destination-product observation originates from a single firm, allowing us to compute HHIs based on these data. The results with these alternative data-sets were similar to the ones based on aggregated firm-level information. (Results are available upon request).

Table 6: Summary Statistics for Industry-Level Market Shares and Exporter HHIs

	Exporter HHI	Destination Market Share
Mean	0.53	10%
5th pctile	0.07	0.01%
10th pctile	0.12	0.06%
Median	0.48	3%
90th pctile	1	27%
95th pctile	1	48%

Notes: Industry-level data. Table shows summary statistics on the distribution of exporter aggregate market share and HHI. The unit of observation is at the origin-destination-industry level. Sample for year 2010.

The PPML estimates of the distance coefficient in columns (1) and (2) are much smaller in magnitude than the ones obtained at the firm level, and not even statistically significant. Thus, self-selection issues appear to lead to a strong attenuation bias in the PPML estimates. By contrast, the OLS results presented in columns (3) and (4) look more reasonable: The distance coefficient is -1.12 without, and -1.26 with, oligopoly correction. Thus, the oligopoly bias is still at work at the industry level, albeit of a smaller magnitude (around 10%) than at the firm level. The finding that these estimates are somewhat smaller than those at the firm level (-1.52 in our preferred firm-level specification; see Table 3) suggests that the sample-selection bias is stronger than the extensive-margin bias.

Table 7: Industry-level Gravity Estimates without Controlling for Selection

	(1)	(2)	(3)	(4)
Method	PPML w/o corr	PPML w/ corr	OLS w/o corr	OLS w/ corr
log distance	-0.263	-0.186	-1.128***	-1.260***
	(0.168)	(0.147)	(0.195)	(0.216)
Observations	107,064	107,064	66,563	66,563
(Pseudo) R-squared	0.49	0.36	0.31	0.29
Industry-origin FE	YES	YES	YES	YES
Industry-dest. FE	YES	YES	YES	YES

Notes: Industry-level data. Oligopoly correction with $\sigma = 5$ and $\gamma = 0$. Standard errors clustered at destination level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

We now apply the HMR approach to correct for self-selection into exporting. Table 8 reports results from the step-one Probit estimation of the propensity to export.[27] As expected,

[27] We include 2-digit sector-origin and 2-digit sector-destination fixed effects, as using 6-digit industry-origin and industry-destination fixed effects is computationally infeasible with the Probit model. However, results using a linear probability model indicate hardly any changes in the point estimates when adding these more detailed fixed effects.

the dummies for high business-startup cost and long business-startup time are negatively and significantly associated with the propensity to export. Unsurprisingly, distance is also negatively related to export propensity, providing evidence for sample selection.

Table 8: First Step of HMR Procedure (Export Propensity)

	Export > 0
log distance	-0.420*
	(0.234)
high startup cost	-1.340***
	(0.421)
long startup time	-2.102***
	(0.337)
Observations	107,064
Sector-origin FE	YES
Sector-dest. FE	YES

Notes: Industry-level data. HMR step-one Probit regression of propensity to export. Standard errors clustered at destination level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 9 reports results from the step-two HMR regression, using $\sigma = 5$ and $\gamma = 0$ in the construction of the oligopoly correction term. The specifications in columns (1) and (2) only include the inverse Mills ratio, and thus correct for sample selection but not for the extensive-margin effect. Columns (3)–(6) show the results from the full HMR procedure. Columns (3) and (4) include a quadratic function of $\log \hat{Z}_{onz}$, while columns (5) and (6) include a third-order polynomial. Throughout, the coefficient on distance in the specifications with oligopoly correction is about 10% larger in magnitude than in those without. While the oligopoly-corrected estimated coefficients are still slightly smaller than those estimated at the firm level, they are larger than those from the OLS regression. This suggests that the sample-selection bias is slightly stronger than the extensive-margin bias. As expected, the specifications that only include the inverse Mills ratio yield the highest coefficient estimates, as they do not control for the extensive-margin effect.

Table 10 reports the results for the same specifications, but using the mean estimates of σ and γ for the oligopoly correction. While the point estimates on log distance hardly change, the estimated fundamental distance coefficient ($\widehat{\beta}$) becomes significantly larger, as it did at the firm level.

Finally, we run the HMR procedure as before, but now separately by 2-digit sectors (which is also the level at which σ and γ are allowed to vary; see Table 2). Table 11 reports summary statistics across sectors on the distribution of the estimated coefficients and the magnitude of the oligopoly bias. While the oligopoly bias is relatively small in the median

Table 9: Industry-level Gravity Estimates, $\sigma = 5$ and $\gamma = 0$

	(1)	(2)	(3)	(4)	(5)	(6)
Method	Heck w/o	Heck w/	HMR^2 w/o	HMR^2 w/	HMR^3 w/o	HMR^3 w/
log distance	-1.189***	-1.327***	-1.151***	-1.284***	-1.150***	-1.284***
	(0.198)	(0.220)	(0.190)	(0.209)	(0.193)	(0.212)
inv. Mills	-0.121	-0.130	0.678***	0.799***	0.639**	0.805**
	(0.165)	(0.182)	(0.169)	(0.192)	(0.309)	(0.344)
log $\hat{Z}$			0.907***	1.067***	0.736	1.097
			(0.265)	(0.293)	(1.305)	(1.396)
log $\hat{Z}^2$			-0.103*	-0.125**	-0.0297	-0.137
			(0.0565)	(0.0621)	(0.534)	(0.567)
log $\hat{Z}^3$					-0.0102	0.00176
					(0.0693)	(0.0732)
$\hat{\beta}_{distance}$	0.297	0.331	0.288	0.32	0.288	0.321
Observations	60,662	60,662	60,662	60,662	60,662	60,662
R-squared	0.30	0.28	0.30	0.28	0.30	0.28
Or.-ind. FE	YES	YES	YES	YES	YES	YES
Dest.-ind. FE	YES	YES	YES	YES	YES	YES

Notes: Industry-level data, pooled across industries, second step of HMR procedure. Oligopoly correction with $\sigma = 5$ and $\gamma = 0$. Standard errors clustered at destination level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 10: Industry-level Gravity Estimates, $\sigma = 5.39$ and $\gamma = 0.34$

	(1)	(2)	(3)	(4)	(5)	(6)
Method	Heck w/o	Heck w/	HMR^2 w/o	HMR^2 w/	HMR^3 w/o	HMR^3 w/
log distance	-1.189***	-1.243***	-1.151***	-1.202***	-1.150***	-1.202***
	(0.198)	(0.206)	(0.190)	(0.197)	(0.193)	(0.200)
inv mills	-0.121	-0.124	0.678***	0.725***	0.639**	0.703**
	(0.165)	(0.171)	(0.169)	(0.178)	(0.309)	(0.322)
log $\hat{Z}$			0.907***	0.969***	0.736	0.876
			(0.265)	(0.275)	(1.305)	(1.338)
log $\hat{Z}^2$			-0.103*	-0.112*	-0.0297	-0.0714
			(0.0565)	(0.0586)	(0.534)	(0.546))
log $\hat{Z}^3$					-0.0102	-0.00556
					(0.0693)	(0.0707)
$\hat{\beta}_{distance}$	0.767	0.802	0.743	0.776	0.743	0.776
Observations	60,662	60,662	60,662	60,662	60,662	60,662
R-squared	0.30	0.29	0.30	0.29	0.30	0.29
Or.-ind. FE	YES	YES	YES	YES	YES	YES
Dest.-ind. FE	YES	YES	YES	YES	YES	YES

Notes: Industry-level data, pooled across industries, second step of HMR procedure. Oligopoly correction with mean of estimated σ and γ. Standard errors clustered at destination level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

sector, it is substantial in a significant minority of sectors. Moreover, the absolute oligopoly bias is positively correlated with the product of the average (across origins, destinations, and industries) Herfindahl index and the average exporting country's aggregate market share, with a correlation coefficient of 0.16, suggesting a larger bias in more-concentrated sectors.

Table 11: Industry-Level Gravity Estimates by 2-digit Sector

Median est coefficient	$\sigma = 5, \gamma = 0$	σ est, $\gamma = 0$	σ, γ est
log distance w/o corr	-0.880	-0.873	-0.872
log distance w/ corr	-1.040	-1.023	-0.930
$\hat{\beta}_{distance}$ w/ corr	0.260	0.294	0.622
$\hat{\beta}_{distance}$ w/o corr	0.220	0.278	0.585
abs. pct. bias (10th pctile)	4%	1.4%	0.01%
abs. pct. bias (median)	13.7 %	11.2%	5.2%
abs. pct. bias (90th pctile)	39.9%	46.4%	19.4%

Notes: Industry-level data. Table shows summary statistics for the distribution of estimated coefficients from the HMR procedure by 2-digit HS sector.

5 Monte Carlo Simulations

In this section, we perform Monte Carlo simulations to evaluate the merits of our oligopoly correction terms. To this end, we develop and calibrate a model in which firms first self-select into export destinations and then compete in quantities. Using the calibrated model, we generate a Monte Carlo data-set to which we then apply our firm- and industry-level estimation procedures. We confirm that our oligopoly correction significantly improves the accuracy of our estimates.

Setup. The model is as described in Section 2, with $\lambda = 1$ (Cournot-Nash conduct). To avoid general-equilibrium effects (and, in the next section, to obtain a money-metric measure of social welfare), we assume that the representative consumer in each country has quasi-linear preferences:[28]

$$U_n = q_{0n} + E_n \int_{z\in[0,1]} \log \left(\sum_{j\in\mathcal{J}_n(z)} a_{jn}(z)^{\frac{1}{\sigma}} q_{jn}(z)^{\frac{\sigma-1}{\sigma}} \right)^{\frac{\sigma}{\sigma-1}} dz,$$

[28]This utility function is the quasi-linear version of the one in Section 2, with $\alpha_n(z) \equiv 1$. The E_n term in front of the integral ensures that the representative consumer's expenditures on the differentiated products are E_n.

where q_{0n} denotes consumption of the outside good. For simplicity, we assume that parameters (such as the elasticity of substitution, or various technology parameters that are described in more detail below) do not vary across industries. Industries will still be heterogeneous due to different realizations of random variables such as productivity draws.

Each country has a fixed labor endowment. The outside good is freely traded and produced using only labor with a constant-returns-to-scale technology that is the same in all countries. We assume that parameters are such that it is produced in positive amount everywhere, so that its price is the same in all countries. We further choose that good as the numeraire, which pins down the wage rate in all countries. In what follows, all costs should be understood as being incurred in terms of labor.

We focus on an industry $z \in [0, 1]$ and drop the industry index to ease notation. We now put more structure on the distribution of cost and quality shocks, and on how firms make entry decisions into export destinations.

Recall from Section 2 that the cost for firm i of producing and selling q_{in} units in market n is $C_{in}(q_{in}) = \frac{1}{1+\gamma} c_{in} \tau_{in} q_{in}^{1+\gamma}$. We decompose c_{in} log-linearly as $\log c_{in} = \varepsilon_i^c + \varepsilon_{in}^c$, where ε_i^c and ε_{in}^c are independent draws from normal distributions with mean zero and variance υ^2 and θ^2, respectively. The iceberg-type trade cost τ_{in} is set equal to 1 if firm i is based in country n and otherwise to $\tau_{on} \equiv T \times (\text{dist}_{on})^\beta$, where o denotes the country in which firm i is located, and T and β are parameters. Finally, we set a_{in} (the quality of product i in market n) equal to 1 for every i and n.[29]

A country-o firm that wants to sell in country $n \neq o$ must pay a fixed cost $f_{on} \equiv F \times \phi_{on}^o \times \phi_{on}^u$, where F is a parameter and ϕ_{on}^o and ϕ_{on}^u are i.i.d. draws from a standard log-normal distribution. The reason for this decomposition is that we will later assume that ϕ_{on}^o is observable to the econometrician whereas ϕ_{on}^u is not, so that ϕ_{on}^o can be used as an excluded first-stage variable when applying the HMR procedure. We set $f_{oo} = 0$ for every country o, so that a firm is always active in its home market.

We consider a two-stage game of complete information in which firms first simultaneously decide which markets to enter, and then compete in quantities in each market. Under oligopoly, this game is likely to have multiple subgame-perfect equilibria. If there were no

[29] Thus, using the notation of Section 2.2, we are setting

$$\varepsilon_n^c = \varepsilon_i^a = \varepsilon_n^a = \varepsilon_{in}^a = \varepsilon_i^\tau = \varepsilon_n^\tau = \varepsilon_{in}^\tau = 0.$$

The assumption that there is no destination-specific shock ($\varepsilon_n^c = \varepsilon_n^a = \varepsilon_n^\tau = 0$) is without loss of generality: Since such shocks would affect all firms symmetrically, they would have no impact on equilibrium market shares and profits given CES demand. As for the firm and firm-destination quality and trade-cost shocks, we could alternatively assume that they are drawn i.i.d. from normal distributions and obtain an observationally equivalent model, since the resulting firm types would still be log-normally distributed.

fixed-cost heterogeneity, it would be possible to rank firms from highest to lowest (destination-specific) type and construct a subgame-perfect equilibrium in which high-type firms enter first. With fixed-cost heterogeneity (in addition to type heterogeneity), there is no such natural ranking of firms and constructing a subgame-perfect equilibrium is a non-trivial combinatorial problem. We therefore make the following simplifying behavioral assumption: When making entry decisions, firms believe that they will receive monopolistic-competition profits (given the set of firms that entered). We can then follow Spence (1976) and rank firms according to their survival coefficients, $(c_{in}\tau_{on})^{\frac{1-\sigma}{1+\sigma\gamma}} / f_{on}$, in each market n. This pins down a natural "equilibrium" entry sequence in market n, in which firms with a higher survival coefficient enter first.

Calibration. We choose parameter values to generate a Monte Carlo data-set broadly similar to our firm-level data-set. We use the same set of countries as in the empirical implementation and take the bilateral distance matrix dist_{on} directly from the data. Market size in country n, E_n, is set equal to an amount proportional to that country's GDP in the data. We follow Chaney (2008) in assuming that the number of firms based in each country is proportional to its GDP. The proportionality coefficient is chosen so that the total number of firms is 220, which is similar to the number of firms in the average industry in our data-set. The elasticity of substitution σ and the returns-to-scale parameter γ are set to 5 and 0, respectively, as in our baseline empirical specification. Finally, we set $\beta = 0.38$, which is our baseline empirical estimate of the distance coefficient (see Table 3).

We still require values for the following four parameters: F, the intercept of the fixed-cost function; T, the intercept of the trade-cost function; υ, the standard deviation of firm baseline productivity draws; and θ, the standard deviation of firm-destination productivity shocks. We calibrate those parameters to match the following empirical moments (computed using the French and Chinese firm-level data): 1. The fraction of firm-destination-industry-year observations with zero trade flows (92%); 2. the mean (by destination-industry-year) aggregate combined market share of French and Chinese firms (13.9%); 3. the median (by origin-industry-year) 90/10 ratio of firm-level total exports (451); and 4. the median (by origin-destination-industry-year) 90/10 ratio of firm-destination exports (220).

The fact that each of the moments has a natural parameter counterpart gives rise to the following informal identification argument. Intuitively, we expect F to have a strong and negative effect on the first moment, T to have a strong and negative effect on the second moment, υ to have a strong and positive effect on the third moment, and θ to have a strong and positive effect on the fourth moment. In practice, we adjust the vector of parameters (F, T, υ, θ) to minimize the sum of the squared Davis-Haltiwanger deviations between theoretical and empirical moments.[30]

We approximate the theoretical moments using Monte Carlo integration. For each parameter vector, we perform 10 Monte Carlo runs.[31] For each run, we randomly draw vectors and matrices of firm-level baseline costs (ε_i^c), firm-destination cost shocks (ε_{in}^c), and fixed-cost shocks (ϕ_{on}^o) and (ϕ_{on}^u). For each destination within a run, we then compute the equilibrium of the entry stage and, using a variant of Nocke and Schutz (2018b)'s nested fixed-point algorithm, the equilibrium of the quantity-setting stage. Having done that for all ten runs, we compute arithmetic means (for moments 1 and 2) and medians (for moments 3 and 4) to obtain Monte Carlo approximations to our theoretical moments.

Our calibration algorithm converges to $F = 3.62 \times 10^{-9}$ (times total world expenditures in the industry, which we normalized to unity), $T = 0.827$, $\upsilon = 0.394$, and $\theta = 1.23$. We obtain nearly perfect matches for the second, third, and fourth moments (0.140, 449, and 220, respectively, vs. 0.139, 451, and 220 in the data), and we slightly under-predict the fraction of zeros in the firm-level export matrix (82.8% vs. 92% in the data). The resulting sum of squared deviations is 0.011.

Data generation and results. Using the calibrated parameters, we generate the Monte Carlo data-set. We perform 200 Monte Carlo runs. Each run features different realizations of the random vectors and matrices of firm-level baseline costs (ε_i^c), firm-destination cost shocks (ε_{in}^c), and fixed-cost shocks (ϕ_{on}^o and ϕ_{on}^u), and can thus be thought of as a different industry or a different time period. For each run, we compute the equilibrium of the entry model and of the quantity-setting game in all markets, and we store firm-level sales and market shares, origin, destination, firm, and run indicators, and bilateral distance and observable fixed-cost shocks. To make the data-set comparable to the one used in our empirical applications, we keep observations only for firms based in the countries corresponding to France and China. We thus obtain a firm-level data-set, which we also aggregate up to construct an industry-level data-set. On these, we run our firm- and industry-level regressions to evaluate the

[30]The Davis-Haltiwanger deviation (Davis, Haltiwanger, and Schuh, 1996) is defined as the difference between the theoretical and empirical moments, divided by the arithmetic average of the theoretical and empirical moments. This residual converges to the percentage deviation when the theoretical moment tends to the empirical moment. The advantage of using this residual for our calibration procedure is that, in contrast to the percentage deviation, it always remains bounded and gives rise to symmetric punishments for positive and negative deviations.

[31]Note that, while the number of Monte Carlo runs is small, each run generates data for about 100 firms and 33 destinations, so that there is relatively little variation in theoretical moments across runs. Increasing the number of runs beyond 10 would only make a small difference, but would substantially increase computational requirements.

performance of our oligopoly correction terms.

Table 12 reports the results from our firm-level regressions. A first observation is that all OLS estimates are strongly biased towards zero, consistent with the heteroskedasticity bias discussed in Section 3 and with the empirical results in Table 3. Focusing now on PPML estimates, we see that the specification without oligopoly correction significantly underestimates the absolute value of the distance coefficient. The PPML specification with oligopoly correction delivers an estimate that is very similar to the true distance coefficient (-1.52). Interestingly, the (biased) PPML estimate without oligopoly correction is very close to the empirical estimate in Table 3 (-0.851 vs. 0.874). Thus, our Monte Carlo data-set generates an oligopoly bias almost identical in size to the one obtained in our empirical analysis.[32]

Table 12: Monte Carlo: Firm-Level Results

	(1)	(2)	(3)	(4)
Method	PPML w/o corr	PPML w/ corr	OLS w/o corr	OLS w/ corr
log distance	-0.851***	-1.403***	-0.454**	-0.472**
	(0.135)	(0.365)	(0.202)	(0.214)
Observations	30,198	30,198	9,529	9,529
(Pseudo) R-squared	0.21	0.74	0.48	0.47
Firm-run FE	YES	YES	YES	YES
Destination-run FE	YES	YES	YES	YES

Notes: Monte Carlo data-set, firm-level data, pooled across Monte Carlo runs. Results for top 3 exporters. Cournot model with $\sigma = 5$ and $\gamma = 0$. Standard errors clustered at the destination level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. True log-distance coefficient is -1.52.

Table 13 reports the estimation results from industry-level regressions. In all specifications, our oligopoly correction improves the accuracy of the distance-coefficient estimate. Specifications that account for self-selection (columns 'Heck' and 'HMR'), when combined with our correction term, deliver estimates that are very close to the true value of -1.52. Interestingly, the PPML estimator delivers by far the worst results, as it did in our empirical application in Section 4.

6 Counterfactual Simulations

In this section, we turn to the welfare effects of a trade liberalization, and quantitatively assess the importance of accounting for oligopolistic behavior. We do so by calibrating two

[32]As shown in Table 20 in Appendix D, very similar results obtain when using (i) only the top French and Chinese exporter and (ii) the top-5 French and Chinese exporters.

Table 13: Monte Carlo: Industry-Level Results

Method	(1) PPML w/o	(2) PPML w/	(3) OLS w/o	(4) OLS w/	(5) Heck w/o	(6) Heck w/	(7) HMR w/o	(8) HMR w/
log distance	-0.849***	-1.044***	-1.272***	-1.383***	-1.282***	-1.399***	-1.298***	-1.418***
	(0.0549)	(0.0760)	(0.0658)	(0.0728)	(0.0957)	(0.108)	(0.0990)	(0.111)
inv. Mills					0.795	0.909	-8.251	-8.456
					(0.523)	(0.610)	(5.890)	(7.124)
log $\hat{Z}$							-16.73	-17.16
							(14.29)	(17.25)
log $\hat{Z}^2$							5.983	6.119
							(5.813)	(7.013)
log $\hat{Z}^3$							-0.727	-0.743
							(0.810)	(0.975)
Observations	12,132	12,132	11,094	11,094	8,296	8,296	8,296	8,296
R-squared	0.15	0.29	0.64	0.61	0.64	0.62	0.64	0.62
Or.-run FE	YES	YES	YES	YES	YES	YES	YES	YES
Dest.-run FE	YES	YES	YES	YES	YES	YES	YES	YES

Notes: Monte Carlo data-set, industry-level data, pooled across Monte Carlo runs. Cournot model with $\sigma = 5$ and $\gamma = 0$. Standard errors clustered at the destination level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. True log-distance coefficient is -1.52.

versions of the model of Section 5: the oligopoly ('*oli*') version, in which $\lambda = 1$ (Cournot-Nash conduct) and the distance coefficient is set equal to our baseline empirical estimate *with* the oligopoly correction term; and the monopolistic competition ('*mc*') version in which $\lambda = 0$ (monopolistic-competition conduct) and the distance coefficient is our baseline estimate *without* the correction term. We calibrate both versions by matching the same moments in the data, and then use them to simulate a 10% trade-cost reduction and compute the induced welfare effects.

Setup. The setup is as described in Section 5, with the following amendments. First, at the entry stage of the *oli* version, firms now correctly expect to earn oligopoly profits. As discussed above, in the *oli* version such correct conjectures make it infeasible to solve for a subgame-perfect equilibrium under fixed-cost heterogeneity. We thus, second, assume that the fixed export cost, $f > 0$, does not vary across origin-destination pairs. This allows us to rank firms from highest to lowest type and construct an equilibrium of the entry game in which firms with a higher type enter first. Finally, we increase the number of firms based in each country by 1 to ensure that each market is always served by at least two firms, so that consumer surplus is always finite.[33]

Calibration. The set of countries, the bilateral distance matrix, each country's market size, the coefficient that determines the number of firms in each country, the elasticity of substitution, and the returns-to-scale parameter are as in Section 5. The distance coefficient, β, is

[33]Under CES demand, a monopolist would set an infinite price, resulting in consumer surplus being equal to minus infinity.

set to 0.38 in the *oli* version and to 0.22 in the *mc* version, which corresponds to our baseline empirical estimates with and without oligopoly correction (see Table 3). The remaining parameters (f, T, υ, and θ) are chosen to match the same moments as in the previous section. The theoretical moments are again approximated using Monte Carlo integration with 10 iterations.

For the *oli* version, the calibration algorithm converges to $f = 1.36 \times 10^{-8}$, $T = 0.379$, $\upsilon = 0.406$, and $\theta = 1.15$. The calibrated model does a very good job of matching the 90/10 dispersion moments (218 and 456 for firm-destination and firm-level exports, respectively, vs. 220 and 451 in the data) and the mean aggregate share of French and Chinese firms (13.6% vs. 13.9% in the data) but tends to under-predict the fraction of zeros in the firm-level export matrix (79.8% vs. 92% in the data). This results in a sum of squared deviations of 0.0207. The fit of the *mc* calibration is almost as good, with a sum of squared residuals of 0.0262. The calibrated model continues to provide a good match for the 90/10 and aggregate-share moments (221, 455, and 14.2%, respectively) but still under-predicts the fraction of zeros (78.4%). The values of the productivity parameters υ and θ are close to the *oli* calibration ($\upsilon = 0.257$ and $\theta = 1.15$). Productivities are slightly less dispersed in the *mc* calibration, which is intuitive since the sales distribution tends to be more compressed under oligopoly due to incomplete passthrough. As profits tend to be lower under monopolistic competition, the calibrated fixed cost ($f = 5.67 \times 10^{-9}$) is lower than in the *oli* calibration. Finally, the fact that the distance coefficient β is significantly lower in the *mc* calibration mechanically reduces trade costs to all destinations. This results in the intercept of the trade cost function ($T = 1.62$) being higher than in the *oli* calibration, so as not to overpredict the exports of French and Chinese firms.

Computing social welfare. Plugging country n's budget constraint into the representative consumer's utility function, we obtain an expression for social welfare in that country (up to an additive constant):

$$W_n = \int_{z\in[0,1]} \left(E_n \left[\frac{\sigma}{\sigma - 1} \log \left(\sum_{j\in\mathcal{J}_n(z)} q_{jn}(z)^{\frac{\sigma-1}{\sigma}} \right) - 1 \right] + \Pi_n(z) \right) dz, \tag{20}$$

where $\Pi_n(z)$ represents the total profits made by firms based in country n.[34] To report the values of our welfare measures in U.S. dollars, we set E_n equal to country n's GDP share in our data-set multiplied by the value added in manufacturing, added up over all 33 countries.

[34] We are thus assuming that firms are owned by the residents of their country of origin. The results are very similar when assuming instead that consumers own an internationally diversified portfolio.

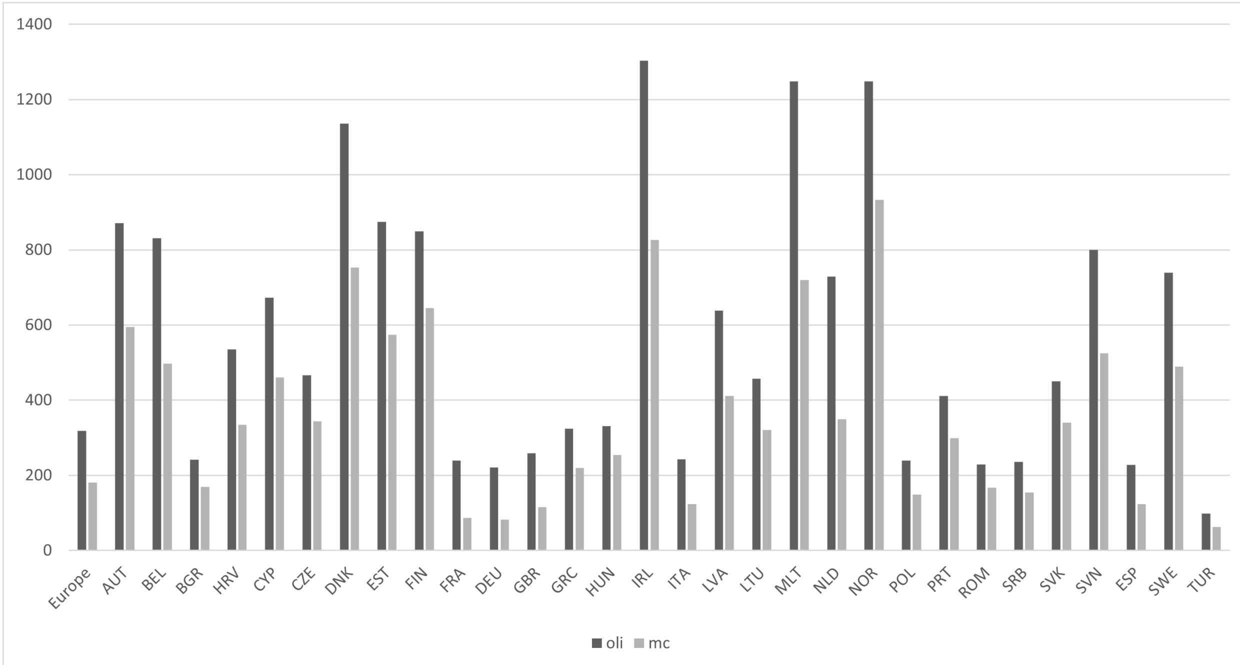

Figure 1: Welfare Effects of a 10% Trade Cost Reduction (in USD per Capita)

Notes: Figure shows the effects of a 10% trade cost reduction on per-capita welfare (measured in USD), by country, under oligopoly (*oli*) and monopolistic competition (*mc*).

As in Section 5, the integral in equation (20) is approximated using Monte Carlo integration with 200 iterations (i.e., 200 industries).

Results. We simulate the equilibrium effects of a 10% reduction in variable trade costs using the *oli* and *mc* versions of the model. Figure 1 reports the resulting changes in social welfare per capita.[35] According to our simulations, the welfare gains from a trade liberalization are substantially higher under oligopoly, with the average European consumer experiencing a utility gain of USD 319 in the *oli* version and USD 181 in the *mc* version. A similar picture emerges when looking at individual European countries, with most countries experiencing gains from trade that are at least 30% higher in the *oli* calibration than in the *mc* calibration. In large, central countries such as France or Germany, the gains from trade under oligopoly are almost three times as high as in the *mc* calibration.

To better understand what drives the difference between the *oli* and the *mc* predictions, we decompose the welfare effects of the trade liberalization into: 1. a trade-cost component (the marginal costs of all exporters decrease by 10%, holding fixed all markups and the

[35]To improve the figure's readability, we have dropped two outliers, Iceland and Luxemburg, for which the gains from trade under oligopoly significantly exceed USD 1,500 per capita. We have also dropped China, which, in this model, benefits very little from trade liberalization, due to it being a remote market to which very few firms find it profitable to export.

set of exporters); 2. a domestic-markups component (due to increased competitive pressure, domestic firms lower their markups); 3. a foreign-markups component (exporters, whose market shares have increased, raise their markups); and 4. an extensive-margin component (the set of exporters adjust).

We now report on the magnitude of these components for European social welfare per capita; the general picture is similar when looking at individual European countries. In our simulations, the extensive-margin component is negligible under both oligopoly and monopolistic competition. The domestic-markups component raises per-capita welfare by USD 67 in the *oli* version, while the foreign-markups component lowers it by USD 20; both components are of course inoperative under monopolistic competition. Finally, the trade-cost component raises per-capita welfare by USD 272 in the *oli* version, and by USD 181 in the *mc* version. Thus, around two-thirds of the gap between the gains from trade under oligopoly and monopolistic competition can be attributed to the fact that the *oli* calibration results in different trade cost parameters, with the remaining third being explained by markup adjustments.[36] This highlights the importance of obtaining a reliable estimate of the distance coefficient β.

7 Conclusion

We have shown that the standard approach to gravity estimation of trade flows suffers from an omitted variable bias when firms behave oligopolistically. We have proposed methods to purge the observed trade flows from market-power effects and thus obtain consistent estimates of gravity parameters. Using French and Chinese export data and Monte Carlo simulations, we have shown that accounting for oligopoly is quantitatively important. When estimating gravity at the firm level, the elasticity of trade flows with respect to distance is twice as large when correcting for market power. While the magnitude of the oligopoly bias is smaller when estimating gravity at the industry level, it is still substantial in a significant minority of industries, in which exports tend to be highly concentrated. In a calibrated version of our model, the welfare gains from a trade liberalization are almost twice as large under oligopoly as under monopolistic competition. These findings reinforce the view that market power effects matter in international trade.

[36] Our results thus contrast with those of Arkolakis, Costinot, Donaldson, and Rodriguez-Clare (2019) in two dimensions. First, while they find that the welfare gains from trade liberalization are lower under variable markups, we find larger gains under oligopoly than under monopolistic competition with constant markups. Second, while they report that the (negative) foreign-markups component more than outweighs the (positive) domestic-markups component, we obtain the opposite.

Appendix

A Proofs

A.1 Proof of Proposition 1

Proof. To complete the proof of the proposition, we need to: (a) Show that the function S is well defined, and study its monotonicity properties and its limits; (b) show that the equilibrium condition (11) has a unique solution; (c) show that, at $\lambda = 1$, the first-order conditions of profit maximization are sufficient for global optimality. We do so below. In the following, we drop the destination index (n) to ease notation.

(a) As $1 + \sigma\gamma > 0$, the right-hand of equation (10) is strictly increasing in s_i, whereas the left-hand side is non-increasing in s_i. It follows that equation (10) has at most one solution. As s_i tends to 0, the left-hand side of that equation tends to 1, whereas the right-hand side tends to 0. As s_i tends to ∞, the left-hand side tends to 1 or $-\infty$, and the right-hand side tends to $+\infty$. The equation therefore has a unique solution, $S(T_i/H, \lambda) \in (0, 1/\lambda)$, where $1/\lambda \equiv \infty$ when $\lambda = 0$.

It is easily checked that $S(\cdot, \cdot)$ is strictly increasing in its first argument and strictly decreasing in its second argument. By monotonicity, $S(\cdot, \lambda)$ has limits at 0 and ∞. Clearly, those limits are equal to 0 and $1/\lambda$, respectively.

(b) The results in part (a) of the proof imply that the left-hand side of equation (11) is strictly decreasing in H, and has limits 0 and $|\mathcal{J}|/\lambda$ as H tends to ∞ and 0, respectively. It follows that equation (11) has a unique solution, $H^*(\lambda)$.

(c) Rewriting equation (2) with $\lambda = 1$ and rearranging terms yields:

$$\frac{\partial \pi_i}{\partial q_i} = q_i^{\gamma} \left[\frac{\sigma - 1}{\sigma} \alpha E \frac{a_i^{\frac{1}{\sigma}} q_i^{-\frac{1+\sigma\gamma}{\sigma}}}{\sum_{j \in \mathcal{J}} a_j^{\frac{1}{\sigma}} q_j^{\frac{\sigma-1}{\sigma}}} \left(1 - \frac{a_i^{\frac{1}{\sigma}} q_i^{\frac{\sigma-1}{\sigma}}}{\sum_{j \in \mathcal{J}} a_j^{\frac{1}{\sigma}} q_j^{\frac{\sigma-1}{\sigma}}} \right) - c_i \tau_i \right].$$

As $1 + \sigma\gamma > 0$, the term inside square brackets is strictly decreasing in q_i. Moreover, that terms tends to $+\infty$ and $-\tau_i c_i$ as q_i tends to 0 and $+\infty$, respectively. It follows that q_i maximizes firm i's profit if and only if firm i's first-order condition holds at q_i. □

A.2 Proof of Proposition 2

Proof. To apply Taylor's theorem, we require the value of $s_{on}^{*\prime}(0)$. This requires computing the partial derivatives of $S(\cdot,\cdot)$ at $\lambda = 0$ and $H_n^{*\prime}(0)$. Differentiating equation (10) with respect to s_{in}, λ, and $t_{in} \equiv T_{in}/H_n$ at $\lambda = 0$ yields

$$-s_{in} d\lambda = \frac{1+\sigma\gamma}{\sigma-1}\frac{ds_{in}}{s_{in}} - \frac{\sigma(1+\gamma)}{\sigma-1}\frac{dt_{in}}{t_{in}}.$$

It follows that[37]

$$t_{in}\partial_1 \log S(t_{in},0) = \frac{\sigma(1+\gamma)}{1+\sigma\gamma} \quad \text{and} \quad \partial_2 \log S(t_{in},0) = -\frac{\sigma-1}{1+\sigma\gamma}S(t_{in},0).$$

Next, we differentiate equation (11) with respect to λ and H_n:

$$\sum_{j\in\mathcal{J}_n}\left[-\frac{T_{jn}}{H}\partial_1 S\left(\frac{T_{jn}}{H},\lambda\right)\frac{dH_n}{H_n} + \partial_2 S\left(\frac{T_{jn}}{H_n},\lambda\right)d\lambda\right] = 0.$$

Setting $\lambda = 0$ and plugging in the values of the partial derivatives of S, we obtain:

$$\sum_{j\in\mathcal{J}_n}\left[-\frac{\sigma(1+\gamma)}{1+\sigma\gamma}s_{jn}^*(0)\frac{dH_n}{H_n} - \frac{\sigma-1}{1+\sigma\gamma}\left(s_{jn}^*(0)\right)^2 d\lambda\right] = 0.$$

Making use of the definition of $\text{HHI}_n(0)$ and of the fact that market shares add up to unity, we obtain:

$$\frac{H_n^{*\prime}(0)}{H_n^*(0)} = -\frac{\sigma-1}{\sigma(1+\gamma)}\,\text{HHI}_n(0).$$

We can now compute $s_{in}^{*\prime}(0)$:

$$\begin{aligned} s_{in}^{*\prime}(0) &= \frac{\partial}{\partial\lambda}S\left(\frac{T_{in}}{H_n^*(\lambda)},\lambda\right)\Bigg|_{\lambda=0} \\ &= -\frac{T_{in}}{H_n^*(0)}\partial_1 S\left(\frac{T_{in}}{H_n^*(0)},0\right)\frac{H_n^{*\prime}(0)}{H_n^*(0)} + \partial_2 S\left(\frac{T_{in}}{H_n^*(0)},0\right) \\ &= \frac{\sigma-1}{1+\sigma\gamma}\left[s_{in}^*(0)\,\text{HHI}_n(0) - (s_{in}^*(0))^2\right]. \end{aligned}$$

[37] Notation: $\partial_k S$ is the partial derivative of S with respect to its kth argument.

It follows that

$$\begin{aligned}
\frac{s^{*\prime}_{on}(0)}{s^*_{on}(0)} &= \frac{\sigma-1}{1+\sigma\gamma}\frac{1}{s^*_{on}(0)}\sum_{j\in\mathcal{J}_{on}}\left[s^*_{jn}(0)\,\mathrm{HHI}_n(0)-\left(s^*_{jn}(0)\right)^2\right] \\
&= \frac{\sigma-1}{1+\sigma\gamma}\left[\mathrm{HHI}_n(0)-s^*_{on}(0)\sum_{j\in\mathcal{J}_{on}}\left(\frac{s^*_{jn}(0)}{s^*_{on}(0)}\right)^2\right] \\
&= \frac{\sigma-1}{1+\sigma\gamma}\left[\mathrm{HHI}_n(0)-s^*_{on}(0)\,\mathrm{HHI}_{on}(0)\right].
\end{aligned}$$

Applying Taylor's theorem at the first order in the neighborhood of $\lambda = 0$ yields:

$$\begin{aligned}
\log s^*_{on}(\lambda) &= \log s^*_{on}(0) + \frac{d}{d\lambda}\log s^*_{on}(\lambda)\bigg|_{\lambda=0}\lambda + o(\lambda) \\
&= \log s^*_{on}(0) + \frac{\sigma-1}{1+\sigma\gamma}\left[\mathrm{HHI}_n(0)-s^*_{on}(0)\,\mathrm{HHI}_{on}(0)\right]\lambda + o(\lambda) \\
&= \log s^*_{on}(0) + \frac{\sigma-1}{1+\sigma\gamma}\left[\mathrm{HHI}_n(\lambda)-s^*_{on}(\lambda)\,\mathrm{HHI}_{on}(\lambda)\right]\lambda + o(\lambda),
\end{aligned}$$

where the last line follows from the fact that $\mathrm{HHI}_n(\lambda) - \mathrm{HHI}_n(0)$ and $s^*_{on}(\lambda)\,\mathrm{HHI}_{on}(\lambda) - s^*_{on}(0)\,\mathrm{HHI}_{on}(0)$ are at most first order. □

B Estimation of Supply and Demand Elasticities

Feenstra (1994) and Broda and Weinstein (2006) propose estimators for the elasticity of substitution, σ, based on the key identifying assumption that shocks over time to import demand and export supply for a given product are uncorrelated. The equivalent condition in our context is that $\mathbb{E}\left(\varepsilon^a_{in}\varepsilon^{\tau c}_{i'n'}\right) = 0$ for all i, i' and n, n', where $\varepsilon^{\tau c}_{in} = \varepsilon^{\tau}_{in} + \varepsilon^{c}_{in}$. That is, we assume that the firm-destination-level elements of taste and cost shocks are uncorrelated across firms and markets.

Note that this assumption is consistent with non-zero correlations between overall taste and cost shocks (i.e., $\mathbb{E}\left(a_{in}c_{in}\right) \neq 0$ is allowed). In particular, our method allows for a positive correlation between firm-level costs and quality (ε^a_i and ε^c_i) which is to be expected if the production costs of firms producing high-quality products are higher. Likewise, our results are robust to a positive correlation between destination market quality and cost shocks (ε^a_n and ε^c_n). For example, such a correlation could arise if firms sell higher-quality goods to high-income markets and incur positive costs of doing so.

We start our derivation by expressing firm-level revenues of firm i in market n in terms of expenditure shares. From equation (1),

$$\log s_{in} = \log\left(\frac{p_{in}q_{in}}{E_n}\right) = \log a_{in} + (1-\sigma)\, p_{in} + (\sigma-1)\log P_n.$$

Now assume that we observe another firm i' selling to the same market n. We can then subtract the logged market share of that firm to eliminate the price index:[38]

$$\Delta^f \log s_{in} = \log s_{in} - \log s_{i'n} = \log a_{in} - \log a_{i'n} + (1-\sigma)\left(\log p_{in} - \log p_{i'n}\right)$$

If we observe the same two firms in another destination n', we can compute a double difference across the two markets as

$$\Delta^d\Delta^f \log s_{in} = (1-\sigma)\,\Delta^d\Delta^f \log p_{in} + \Delta^d\Delta^f \log a_{in},$$

where Δ^f and Δ^d denote log differences across firms and destinations, respectively. Note that double differencing only leaves the firm-destination-level parts of the taste shocks:

$$\Delta^d\Delta^f \log a_{in} = \left(\varepsilon^a_{in} - \varepsilon^a_{i'n}\right) - \left(\varepsilon^a_{in'} - \varepsilon^a_{i'n'}\right).$$

We next derive a similar supply-side equation. We start by rewriting firm i's price in market n as $p_{in}^{1+\gamma} = \left(\frac{c_{in}\tau_{in}}{1-\mu_{in}}\right)(s_{in}E_n)^\gamma$. Taking logs yields

$$(1+\gamma)\log p_{in} = \log\left(c_{in}\tau_{in}\right) - \log\left(1-\mu_{in}\right) + \gamma\log s_{in} + \gamma\log E_n.$$

Double differencing across firms and markets as above, we obtain

$$(1+\gamma)\,\Delta^d\Delta^f \log p_{in} = \Delta^d\Delta^f \log\left(c_{in}\tau_{in}\right) - \Delta^d\Delta^f \log\left(1-\mu_{in}\right) + \gamma\Delta^d\Delta^f \log s_{in},$$

where the double-differenced cost shock again only contains the parts of production and trade costs that are at the firm-destination level:

$$\Delta^d\Delta^f \log\left(c_{in}\tau_{in}\right) = \left(\varepsilon^{\tau c}_{in} - \varepsilon^{\tau c}_{i'n}\right) - \left(\varepsilon^{\tau c}_{in'} - \varepsilon^{\tau c}_{i'n'}\right).$$

Note that as per our identifying assumption, the double-differenced cost and taste shocks are

[38] In principle, we could also subtract the average across all firms active in market n. However, we will argue below that taking differences across individual firms with high market shares is better suited to dealing with selection problems.

uncorrelated, yielding the following moment condition:

$$\mathbb{E}\left(\Delta^d \Delta^f \log a_{in} \times \Delta^d \Delta^f \log\left(c_{in}\tau_{in}\right)\right) = 0.$$

For given σ and γ, we can construct the sample analogues from data on export prices and market shares:

$$\Delta^d \Delta^f \widehat{\log\left(c_{in}\tau_{in}\right)} = (1+\gamma)\, \Delta^d \Delta^f \log p_{in} + \Delta^d \Delta^f \log\left(1-\mu_{in}\right) - \gamma \Delta^d \Delta^f \log s_{in}$$

and

$$\Delta^d \widehat{\Delta^f \log} a_{in} = \Delta^d \Delta^f \log s_{in} - (1-\sigma)\, \Delta^d \Delta^f \log p_{in}.$$

The sample analogue of our moment condition is then given by

$$\Psi\left(\sigma,\gamma\right) = \frac{1}{\left|\mathcal{J}_{nn'}\right|} \sum_{j \in \mathcal{J}_{nn'}} \Delta^d \widehat{\Delta^f \log} a_{in} \times \Delta^d \Delta^f \widehat{\log\left(c_{in}\tau_{in}\right)},$$

where $\mathcal{J}_{nn'}$ denotes the set of firms active in the same two markets. Notice that we obtain one moment condition per country pair. Stacking these up allows to implement a standard GMM estimator of σ and γ.[39]

This still leaves us with a potential selection problem in our GMM estimation procedure for σ and γ. As a solution, we focus again on the top 3 French and Chinese exporters (in terms of their overall exports) for any given 6-digit HS industry. Finally, to obtain a sufficiently large number of observations for the computation of moments in our GMM estimation, we restrict the estimates of σ and γ to be identical within 2-digit HS sectors.

C Price Competition

C.1 Theoretical Results

Under price competition, the profit of firm i when selling in destination n is:

$$\pi_{in} = p_{in} a_{in} p_{in}^{-\sigma} P_n^{\sigma-1} \alpha_n E_n - C_{in}\left(a_{in} p_{in}^{-\sigma} P_n^{\sigma-1} \alpha_n E_n\right),$$

where we have dropped the industry index z for ease of notation.

[39] In practice, this means that we need to observe a sufficiently large number of firms selling in the same industry in at least three different markets.

The degree of strategic interactions between firms continues to be governed by the conduct parameter $\lambda \in [0, 1]$: When firm i increases its price by an infinitesimal amount, it perceives the induced effect on P_n to be equal to $\lambda \partial P_n / \partial p_{in}$. It is still the case that monopolistic competition arises when $\lambda = 0$, whereas Bertrand competition arises when $\lambda = 1$. The first-order condition of profit maximization of firm i in destination n is given by

$$0 = \frac{\partial \pi_{in}}{\partial p_{in}} = a_{in} p_{in}^{-\sigma} P_n^{\sigma-1} \alpha_n E_n + (p_{in} - C'_{in}(q_{in})) \left[-\frac{\sigma}{p_{in}} + \frac{\sigma - 1}{P_n} \lambda \frac{\partial P_n}{\partial p_{in}} \right] \alpha_n E_n a_{in} p_{in}^{-\sigma} P_n^{\sigma-1}$$

$$= q_{in} \left(1 - \frac{p_{in} - C'_{in}(q_{in})}{p_{in}} \left[\sigma - \lambda(\sigma - 1) s_{in} \right] \right), \tag{21}$$

where

$$s_{in} \equiv \frac{a_{in} p_{in}^{1-\sigma}}{\sum_{j \in \mathcal{J}} a_{jn} p_{jn}^{1-\sigma}} \tag{22}$$

continues to be the market share of firm i in destination n.

Equation (21) pins down firm i's optimal markup under price competition:

$$\mu_{in} = \frac{1}{\sigma - \lambda \left(\sigma - 1 \right) s_{in}},$$

where $\mu_{in} = \frac{p_{in} - C'_{in}(q_{in})}{p_{in}}$ is firm i's Lerner index. Apart from this change in the expression for the firm's optimal markup, all other firm-level results go through as before.

We now turn our attention to the industry-level results. As in Section 2.3, we begin by employing an aggregative games approach to analyze the equilibrium in a given market, dropping the market subscript n to ease notation. The market-level aggregator H is now defined as

$$H \equiv P^{1-\sigma} = \sum_{j \in \mathcal{J}} a_j p_j^{1-\sigma}$$

and firm i's type as

$$T_i \equiv a_i \left(\alpha E \right)^{\frac{\gamma(1-\sigma)}{1+\gamma}} \left(c_i \tau_i \right)^{\frac{1-\sigma}{1+\gamma}}.$$

Plugging these definitions into equation (21), making use of equation (22), and rearranging, we obtain:

$$\left(1 - s_i^{\frac{1+\sigma\gamma}{\sigma-1}} \left(\frac{H}{T_i} \right)^{\frac{1+\gamma}{\sigma-1}} \right) \left(\sigma - \lambda(\sigma - 1) s_i \right) = 1. \tag{23}$$

Note that the left-hand side of equation (23) is strictly decreasing on the interval

$$\left(0, \min\left\{\frac{\sigma}{\lambda(\sigma-1)}, \left(\frac{T_i}{H}\right)^{\frac{1+\gamma}{1+\sigma\gamma}}\right\}\right)$$

and tends to σ and 0 as s_i tends to the lower and upper endpoints of that interval, respectively. Equation (23) therefore has a unique solution on the above interval, denoted $S(t_i, \lambda)$ with $t_i \equiv T_i/H$. (Solutions outside that interval necessarily give rise to strictly negative markups and are thus suboptimal.)

It is easily checked that S is strictly increasing in its first argument, strictly decreasing in its second argument, and tends to 0 and $1/\lambda$ as t_i tends to 0 and ∞, respectively.

As before, the equilibrium condition is that market shares must add up to unity:

$$\sum_{j\in\mathcal{J}} S\left(\frac{T_i}{H}, \lambda\right) = 1. \tag{24}$$

The properties of the function S, described above, imply that this equation has a unique solution, $H^*(\lambda)$.

To summarize:

Proposition A. *In each destination market, and for any conduct parameter λ, there exists a unique equilibrium in prices. The equilibrium aggregator level $H^*(\lambda)$ is the unique solution to equation* (24)*. Each firm i's equilibrium market share is $s_i^*(\lambda) = S(T_i/H^*(\lambda), \lambda)$, where $S(T_i/H^*(\lambda), \lambda)$ is the unique solution to equation* (23)*.*

Proof. All that is left to do is check that first-order conditions are sufficient for optimality when $\lambda = 1$. Combining equations (21) and (23) yields:

$$\frac{\partial \pi_i}{\partial p_i} = q_i \left[1 - \chi(p_i)\phi(p_i)\right],$$

where

$$\chi(p_i) \equiv 1 - \left(\frac{a_i p_i^{1-\sigma}}{\sum_j a_j p_j^{1-\sigma}}\right)^{\frac{1+\sigma\gamma}{\sigma-1}} \left(\frac{\sum_j a_j p_j^{1-\sigma}}{T_i}\right)^{\frac{1+\gamma}{\sigma-1}} \text{ and } \phi(p_i) \equiv \sigma - (\sigma-1)\frac{a_i p_i^{1-\sigma}}{\sum_j a_j p_j^{1-\sigma}}.$$

As $1+\sigma\gamma > 0$, the functions χ and ϕ are strictly increasing. Moreover, $\phi(p_i) > 0$ for every p_i, whereas there exists $\widetilde{p}_i > 0$ such that $\chi(p_i) > 0$ if $p_i > \widetilde{p}_i$ and $\chi(p_i) < 0$ if $p_i < \widetilde{p}_i$. Hence, π_i is strictly increasing on the interval $(0, \widetilde{p}_i)$, and firm i's first-order condition holds nowhere on that interval. The fact that $\lim_{p_i \to \infty} \chi(p_i) = 1$ and $\lim_{p_i \to \infty} \phi(p_i) = \sigma$ and the monotonicity properties of χ and ϕ on $(\widetilde{p}_i, \infty)$ imply the existence of a unique $\widehat{p}_i$ at which firm i's first-order condition holds. Moreover, π_i is strictly increasing on $(\widetilde{p}_i, \widehat{p}_i)$ and strictly decreasing on $(\widehat{p}_i, \infty)$. First-order conditions are therefore sufficient for optimality. □

Having characterized the equilibrium in a given destination, we now adapt the first-order approach to industry-level gravity to the case of price competition. As in Section 2.3, let $\mathcal{E} \subsetneq \mathcal{J}$ denote the subset of exporters in country e that sell in the destination market n. The combined market share of those exporters in market n is given by

$$s_e^*(\lambda) \equiv \sum_{i \in \mathcal{E}} s_i^*(\lambda).$$

As before, we approximate $s_e^*(1)$ at the first order. The definitions of HHI and HHI_e are as in Section 2.3.

We obtain:

Proposition B. *At the first order, in the neighborhood of $\lambda = 0$, the logged joint market share in destination n of the firms from export country e is given by*

$$\log s_e^*(\lambda) = \log s_e^*(0) + \frac{\sigma - 1}{\sigma(1 + \sigma\gamma)} \left[\text{HHI}(\lambda) - s_e^*(\lambda) \, \text{HHI}_e(\lambda) \right] \lambda + o(\lambda).$$

Proof. The proof follows the same developments as the proof of Proposition 2. We begin by computing the partial derivatives of S at $\lambda = 0$. It is useful to rewrite first equation (23) as

$$s_i = t_i^{\frac{1+\gamma}{1+\sigma\gamma}} \left(1 - \frac{1}{\sigma - \lambda(\sigma - 1)s_i} \right)^{\frac{\sigma-1}{1+\sigma\gamma}}. \tag{25}$$

Taking the logarithm and totally differentiating the equation at $\lambda = 0$ yields:

$$\frac{ds_i}{s_i} = \frac{1 + \gamma}{1 + \sigma\gamma} \frac{dt_i}{t_i} - \frac{\sigma - 1}{\sigma(1 + \sigma\gamma)} s_i d\lambda.$$

The partial derivatives of S are thus given by

$$t_i \partial_1 \log S(t_i, 0) = \frac{1 + \gamma}{1 + \sigma\gamma} \quad \text{and} \quad \partial_2 \log S(t_i, 0) = -\frac{\sigma - 1}{\sigma(1 + \sigma\gamma)} S(t_i, 0).$$

To obtain $H^{*\prime}(0)$, we differentiate equation (24):

$$\sum_{j\in\mathcal{J}}\left[-\frac{T_j}{H}\partial_1 S\left(\frac{T_j}{H},\lambda\right)\frac{dH}{H}+\partial_2 S\left(\frac{T_j}{H},\lambda\right)d\lambda\right]=0.$$

Setting $\lambda = 0$, plugging in the values of the partial derivatives of S, and using the fact that market shares add up to unity, we obtain:

$$\frac{H^{*\prime}(0)}{H^*(0)}=-\frac{\sigma-1}{\sigma(1+\gamma)}\,\mathrm{HHI}(0).$$

Next, we compute $s_i^{*\prime}(0)$:

$$\begin{aligned}s_i^{*\prime}(0)&=-\frac{T_i}{H^*(0)}\partial_1 S\left(\frac{T_i}{H^*(0)},0\right)\frac{H^{*\prime}(0)}{H^*(0)}+\partial_2 S\left(\frac{T_i}{H^*(0)},0\right)\\&=\frac{\sigma-1}{\sigma(1+\sigma\gamma)}\left[s_i^*(0)\,\mathrm{HHI}(0)-(s_i^*(0))^2\right].\end{aligned}$$

Adding up and dividing by $s_e^*(0)$ yields:

$$s_e^{*\prime}(0)=\frac{\sigma-1}{\sigma(1+\sigma\gamma)}\left[\mathrm{HHI}(0)-s_e^*(0)\,\mathrm{HHI}_e(0)\right].$$

As in the proof of Proposition 2, we can then apply Taylor's theorem to obtain the result. □

Proposition B motivates the following approximation:

$$\log s_e^*(1)\simeq\log s_e^*(0)+\frac{\sigma-1}{\sigma(1+\sigma\gamma)}\left[\mathrm{HHI}(1)-s_e^*(1)\,\mathrm{HHI}_e(1)\right].$$

As in Section 2.3, this approximation can then be used to derive the industry-level gravity regression

$$\log\widetilde{r}_{en}=\zeta_e+\xi_n+\beta\frac{1-\sigma}{1+\sigma\gamma}X_{en}+\eta_{en}$$

where

$$\log\widetilde{r}_{en}\equiv\log r_{en}+\frac{\sigma-1}{\sigma(1+\sigma\gamma)}s_{en}\,\mathrm{HHI}_{en}$$

is the value of export flows from e to n, purged from oligopolistic market power effects. Note that the correction term under price competition is equal to the one under quantity competition divided by σ.

C.2 Empirical Results

Table 14 presents results for our estimates of σ and γ using the estimation procedure from Section B but replacing the Cournot markup formula with its Bertrand equivalent. This only leads to minor changes in coefficient estimates.

Table 14: Price Elasticities and Returns-to-Scale Estimates – Price Competition

	σ	γ
Mean	4.96	0.31
25th Percentile	2.06	0.02
Median	3.27	0.10
75th Percentile	5.22	0.28
Min	1.01	-0.11
Max	26.03	4.5
Standard Deviation	4.89	0.67
Observations	78	78

Notes: Table shows descriptive statistics for estimates of σ and γ. Estimates computed using 6-digit HS firm-level information but constrained to be identical within 2-digit HS sectors.

Tables 15–16 show results for the pooled firm-level regressions. In all specifications, the point estimates on the distance coefficient are much larger in absolute magnitude when correcting for oligopoly bias. The absolute value of the distance coefficient is slightly smaller than with Cournot competition.

Tables 17–18 show results for the pooled industry-level regressions. Again, the distance coefficient becomes larger in absolute magnitude when including the oligopoly correction term. Like in the case of Cournot competition, the absolute differences in coefficient magnitudes between the estimates with and without correction are smaller than with the firm-level estimates.

Table 15: Firm-level Gravity Estimates – Bertrand competition, $\sigma = 5$, $\gamma = 0$.

	(1)	(2)	(3)	(4)
Method	PPML w/o corr	PPML w/ corr	OLS w/o corr	OLS w/ corr
log distance	-0.874***	-1.418***	-0.232***	-0.246***
	(0.021)	(0.190)	(0.014)	(0.014)
$\hat{\beta}_{distance}$	0.219	0.355	0.058	0.062
Observations	11,955,786	11,955,786	708,386	708,386
(Pseudo) R-squared	0.14	0.26	0.06	0.05
Firm-ind.-year FE	YES	YES	YES	YES
Dest.-ind.-year FE	YES	YES	YES	YES

Notes: Firm-level data, pooled across industries. Results for top 3 exporters. Bertrand model with $\sigma = 5$ and $\gamma = 0$. Standard errors in parentheses, clustered at the destination-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 16: Firm-level Gravity Estimates – Bertrand competition, $\sigma = 4.96$, $\gamma = 0.31$.

	(1)	(2)	(3)	(4)
Method	PPML w/o corr	PPML w/ corr	OLS w/o corr	OLS w/ corr
log distance	-0.874***	-0.957***	-0.232***	-0.237***
	(0.021)	(0.040)	(0.014)	(0.014)
$\hat{\beta}_{distance}$	0.613	0.560	0.058	0.059
Observations	11,955,786	11,955,786	708,392	708,386
(Pseudo) R-squared	0.14	0.15	0.06	0.06
Firm-ind.-year FE	YES	YES	YES	YES
Dest.-ind.-year FE	YES	YES	YES	YES

Notes: Firm-level data, pooled across industries. Results for top 3 exporters. Bertrand model with mean of estimated σ and γ. Standard errors in parentheses, clustered at the destination-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 17: Industry-level Gravity Estimates – Bertrand competition, $\sigma = 5$, $\gamma = 0$

	(1)	(2)	(3)	(4)	(5)	(6)
Method	Heck w/o	Heck w/	HMR^2 w/o	HMR^2 w/	HMR^3 w/o	HMR^3 w/
log distance	-1.189***	-1.217***	-1.151***	-1.177***	-1.150***	-1.177***
	(0.198)	(0.202)	(0.190)	(0.194)	(0.193)	(0.197)
inv mills	-0.121	-0.123	0.678***	0.702***	0.639**	0.672**
	(0.165)	(0.168)	(0.169)	(0.174)	(0.309)	(0.315)
log $\hat{Z}$			0.907***	0.939***	0.736	0.809
			(0.265)	(0.270)	(1.305)	(1.322)
log $\hat{Z}^2$			-0.103*	-0.108*	-0.0297	-0.0513
			(0.0565)	(0.0575)	(0.534)	(0.540)
log $\hat{Z}^3$					-0.0102	-0.00780
					(0.0693)	(0.0700)
$\hat{\beta}_{distance}$						
Observations	60,662	60,662	60,662	60,662	60,662	60,662
R-squared	0.30	0.30	0.30	0.30	0.30	0.30
Or.-ind. FE	YES	YES	YES	YES	YES	YES
Dest.-ind. FE	YES	YES	YES	YES	YES	YES

Notes: Industry-level data. Bertrand model with $\sigma = 5$ and $\gamma = 0$. Standard errors clustered at destination level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 18: Industry-level Gravity Estimates – Bertrand competition, $\sigma = 4.96$, $\gamma = 0.31$

	(1)	(2)	(3)	(4)	(5)	(6)
Method	Heck w/o	Heck w/	HMR^2 w/o	HMR^2 w/	HMR^3 w/o	HMR^3 w/
log distance	-1.189***	-1.200***	-1.151***	-1.161***	-1.150***	-1.161***
	(0.198)	(0.199)	(0.190)	(0.192)	(0.193)	(0.194)
inv. Mills	-0.121	-0.122	0.678***	0.687***	0.639**	0.652**
	(0.165)	(0.166)	(0.169)	(0.171)	(0.309)	(0.311)
log $\hat{Z}$			0.907***	0.919***	0.736	0.765
			(0.265)	(0.267)	(1.305)	(1.311)
log $\hat{Z}^2$			-0.103*	-0.105*	-0.0297	-0.0382
			(0.0565)	(0.0569)	(0.534)	(0.536)
log $\hat{Z}^3$					-0.0102	-0.00925
					(0.0693)	(0.0696)
$\hat{\beta}_{distance}$						
Observations	60,662	60,662	60,662	60,662	60,662	60,662
R-squared	0.30	0.30	0.30	0.30	0.30	0.30
Or.-ind. FE	YES	YES	YES	YES	YES	YES
Dest.-ind. FE	YES	YES	YES	YES	YES	YES

Notes: Industry-level data. Bertrand model with mean of of estimated σ and γ . Standard errors clustered at destination level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

D Robustness Checks

Table 19: Firm-level Gravity Estimates. Robustness on Sample of Firms

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	Top 1	Top 1	Top 3	Top 3	Top 5	Top 5
Method	PPML w/o	PPML w/	PPML w/o	PPML w/	PPML w/o	PPML w/
dist., $\gamma = 0$, $\sigma = 5$	-0.978***	-1.257***	-0.874***	-1.518***	-0.793***	-1.532***
	(0.018)	(0.083)	(0.021)	(0.220)	(0.021)	(0.232)
dist., $\gamma = 0.34$, $\sigma = 5.39$	-0.978***	-1.148***	-0.874***	-1.201***	-0.793***	-1.139***
	(0.0177)	(0.0582)	(0.0210)	(0.118)	(0.0212)	(0.115)
Observations	3,690,099	3,690,099	11,955,786	11,955,786	20,265,693	20,265,693
Firm-ind.-year FE	YES	YES	YES	YES	YES	YES
Dest.-ind.-year FE	YES	YES	YES	YES	YES	YES

Notes: Firm-level data, pooled across industries. Results for top 1 exporters (columns 1-2), top 3 exporters (columns 3-4), top 5 exporters (columns 5-6). Cournot model with $\sigma = 5$ and $\gamma = 0$. Standard errors in parentheses, clustered at the destination-year level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 20: Monte Carlo: Firm-level Results. Robustness on Sample of Firms

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	Top 1	Top 1	Top 3	Top 3	Top 5	Top 5
Method	PPML w/o	PPML w/	PPML w/o	PPML w/	PPML w/o	PPML w/
dist., $\gamma = 0$, $\sigma = 5$	-1.093***	-2.480***	-0.851***	-1.403***	-0.805***	-1.421***
	(0.267)	(0.703)	(0.135)	(0.365)	(0.104)	(0.310)
Observations	6,340	6,340	30,198	30,198	55,680	55,680
Firm-run FE	YES	YES	YES	YES	YES	YES
Destination-run FE	YES	YES	YES	YES	YES	YES

Notes: Monte Carlo data-set, firm-level data, pooled across Monte Carlo runs. Results for top 1 exporters (columns 1-2), top 3 exporters (columns 3-4), top 5 exporters (columns 5-6). Cournot model with $\sigma = 5$ and $\gamma = 0$. Standard errors in parentheses, clustered at the destination level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

References

Allen, T., C. Arkolakis, and Y. Takahashi (2020): "Universal Gravity," *Journal of Political Economy*, 128(2), 393–433.

Anderson, J. (1979): "A theoretical foundation for the gravity equation," *American Economic Review*, 69(1), 106–116.

Anderson, J., and E. van Wincoop (2003): "Gravity with Gravitas: A Solution to the Border Puzzle," *American Economic Review*, 93(1), 170–192.

Anderson, S. P., N. Erkal, and D. Piccinin (2020): "Aggregative games and oligopoly theory: short-run and long-run analysis," *RAND Journal of Economics*, 51(2), 470–495.

Arkolakis, C., A. Costinot, D. Donaldson, and A. Rodriguez-Clare (2019): "The Elusive Pro-Competitive Effects of Trade," *The Review of Economic Studies*, 86(1), 46–80.

Arkolakis, C., A. Costinot, and A. Rodriguez-Clare (2012): "New Trade Models, Same Old Gains?," *American Economic Review*, 102(1), 94–130.

Atkeson, A., and A. Burstein (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98(5), 1998–2031.

Atkin, D., and D. Donaldson (2015): "Who's Getting Globalized? The Size and Implications of Intra-national Trade Costs," NBER Working Papers 21439, National Bureau of Economic Research, Inc.

Baier, S. L., and J. H. Bergstrand (2007): "Do free trade agreements actually increase members' international trade?," *Journal of International Economics*, 71(1), 72–95.

Bas, M., T. Mayer, and M. Thoenig (2017): "From micro to macro: Demand, supply, and heterogeneity in the trade elasticity," *Journal of International Economics*, 108(C), 1–19.

Brander, J. A. (1995): "Chapter 27: Strategic Trade Policy," in *Handbook of International Economics*, ed. by G. Grossman, and K. Rogoff, vol. 3, pp. 1395 – 1455. Elsevier.

Breinlich, H., V. Nocke, and N. Schutz (2020): "Merger Policy in a Quantitative Model of International Trade," *The Economic Journal*, 130(626), 393–421.

Bresnahan, T. F. (1989): "Chapter 17: Empirical Studies of Industries with Market Power," in *Handbook of Industrial Organization*, ed. by R. Schmalensee, and R. Willig, vol. 2, pp. 1011–1057. Elsevier.

Broda, C., and D. E. Weinstein (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121(2), 541–588.

Burstein, A., and G. Gopinath (2014): "International Prices and Exchange Rates," in *Handbook of International Economics*, vol. 4, chap. Chapter 7, pp. 391–451. Elsevier.

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

Correia, S., P. Guimaraes, and T. Zylkin (2019): "PPMLHDFE: Stata module for Poisson pseudo-likelihood regression with multiple levels of fixed effects," Statistical Software Components, Boston College Department of Economics.

Davis, S. J., J. Haltiwanger, and S. Schuh (1996): *Job Creation and Destruction*. Cambridge, MA: MIT Press.

Eaton, J., and S. Kortum (2002): "Technology, Geography, And Trade," *Econometrica*, 70(5), 1741–1779.

Eckel, C., and P. Neary (2010): "Multi-Product Firms and Flexible Manufacturing in the Global Economy," *Review of Economic Studies*, 77(1), 188–217.

Edmond, C., V. Midrigan, and D. Y. Xu (2015): "Competition, Markups, and the Gains from International Trade," *American Economic Review*, 105(10), 3183–3221.

Feenstra, R. C. (1989): "Symmetric pass-through of tariffs and exchange rates under imperfect competition: An empirical test," *Journal of International Economics*, 27(1-2), 25–45.

——— (1994): "New Product Varieties and the Measurement of International Prices.," *American Economic Review*, 84(1), 157–177.

Ganapati, S., J. S. Shapiro, and R. Walker (2020): "Energy Cost Pass-Through in US Manufacturing: Estimates and Implications for Carbon Taxes," *American Economic Journal: Applied Economics*, 12(2), 303–342.

Gaubert, C., and O. Itskhoki (2021): "Granular Comparative Advantage," *Journal of Political Economy*, 121(3).

Genakos, C., and M. Pagliero (2022): "Competition and Pass-Through: Evidence from Isolated Markets," *American Economic Journal: Applied Economics*, 14(4), 35–57.

Harrigan, J. (1996): "Openness to trade in manufactures in the OECD," *Journal of International Economics*, 40(1-2), 23–39.

Head, K., and T. Mayer (2014): "Gravity Equations: Workhorse,Toolkit, and Cookbook," in *Handbook of International Economics*, vol. 4, chap. Chapter 3, pp. 131–195. Elsevier.

——— (2019): "Poor Substitutes? Counterfactual methods in IO and Trade compared," Manuscript, SciencesPo.

Heid, B., and F. Staehler (2020): "Structural Gravity and the Gains from Trade under Imperfect Competition," Cesifo Working Papers 8121, Cesifo Working Papers.

Helpman, E., M. Melitz, and Y. Rubinstein (2008): "Estimating Trade Flows: Trading Partners and Trading Volumes," *Quarterly Journal of Economics*, 123(2), 441–487.

Khandelwal, A. K., P. K. Schott, and S.-J. Wei (2013): "Trade Liberalization and Embedded Institutional Reform: Evidence from Chinese Exporters," *American Economic Review*, 103(6), 2169–2195.

Melitz, M., and G. Ottaviano (2008): "Market Size, Trade, and Productivity," *Review of Economic Studies*, 75, 295–316.

Melitz, M., and S. Redding (2015): "New Trade Models, New Welfare Implications," *American Economic Review*, 105(3), 1105–46.

Nakamura, E., and D. Zerom (2010): "Accounting for Incomplete Pass-Through," *Review of Economic Studies*, 77(3), 1192–1230.

Neary, J. P. (2003): "Presidential Address: Globalization and Market Structure," *Journal of the European Economic Association*, 1(2-3), 245–271.

Nocke, V., and N. Schutz (2018a): "An Aggregative Games Approach to Merger Analysis in Multiproduct-Firm Oligopoly," NBER Working Papers 24578, National Bureau of Economic Research, Inc.

——— (2018b): "Multiproduct-Firm Oligopoly: An Aggregative Games Approach," *Econometrica*, 86(2), 523–557.

Parenti, M. (2018): "Large and small firms in a global market: David vs. Goliath," *Journal of International Economics*, 110, 103–118.

Santos Silva, J. M. C., and S. Tenreyro (2006): "The Log of Gravity," *The Review of Economics and Statistics*, 88(4), 641–658.

Spence, A. M. (1976): "Product Selection, Fixed Costs, and Monopolistic Competition," *Review of Economic Studies*, 43(2), 217–235.

Tinbergen, J. (1962): *Shaping the World Economy: Suggestions for an International Economic Policy.* Twentieth Century Fund, New-York.

Van Beveren, I., A. B. Bernard, and H. Vandenbussche (2012): "Concording EU Trade and Production Data over Time," NBER Working Papers 18604, National Bureau of Economic Research, Inc.