

Gravity with Granularity*

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

September 15, 2022

Abstract

We evaluate the consequences of oligopolistic behavior for the estimation of gravity equations for trade flows. With oligopolistic competition, firm-level gravity equations based on a standard CES demand framework need to be augmented by markup terms that are functions of firms' market shares. At the aggregate level, the additional term takes the form of the exporting country's market share in the destination country multiplied by an exporter-destination-specific Herfindahl-Hirschman index. For both cases, we show how to construct appropriate correction terms that can be used to avoid problems of omitted variable bias. We demonstrate the quantitative importance of our results for combined French and Chinese firm-level export data as well as for a sample of product-level imports by European countries. Our results show that correcting for oligopoly bias can lead to substantial changes in the coefficients on standard gravity regressors.

Keywords: Gravity Equation, Oligopoly, CES Demand, Aggregative Game

Journal of Economic Literature Classification: F12, F14, L13

*We thank seminar participants at Berkeley, Essex, LSE, Munich, Nottingham, PSE 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).

[†]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 50 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 trade agreements, WTO membership or currency unions.

Despite the rapid progress 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 (forthcoming), 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 and propose modifications to existing frameworks necessary to reconcile the two.

We start our analysis by deriving firm-level gravity equations from a standard CES demand framework. Instead of the usual assumption of monopolistically competitive firms, however, we introduce oligopolistic competition.¹ This leads to the inclusion of variable markup terms in firm-level gravity equations that depend on firms’ market shares. Market shares in turn are determined by both firm- and destination-specific variables as well as bilateral trade frictions, so that omitting oligopoly markup terms will bias coefficient estimates on all other included regressors. Since the relevant markup variation is at the firm-product-destination-time level, standard approaches such as the inclusion of combinations of fixed effects are not feasible.

¹We focus on quantity competition in the main body of the paper and present results for price competition in a separate appendix.

Instead, we show how to adjust trade flows by a suitable correction term to eliminate markup bias, requiring only information on market shares and demand and supply elasticities. To estimate the latter, we propose an extension of the well-known Feenstra-Broda-Weinstein estimation procedure (see Feenstra, 1994; Broda and Weinstein, 2006) that accounts for market power and can be implemented using data on firm-level market shares and unit values.

We also analyze the consequences of oligopolistic behavior for the estimation of gravity equations at more aggregate levels, such as on sector-level trade flows.² We show that the presence of firm-level markup terms leads to bias here as well but that a suitable correction term can again be constructed. This time, the correction takes the form of a destination-specific Herfindahl-Hirschman index (HHI) for exports multiplied by the aggregate market share of the exporting country. Intuitively, what matters for the market power of the exporting country is both its overall market share and how that share is distributed among individual exporters. For example, oligopolistic behavior will be much more pronounced if the overall market share is accounted for by just one firm rather than a large number of small producers.

Having shown theoretically how the presence of oligopolistic competition impacts the estimation of gravity equations, we next analyze the quantitative importance of the resulting bias using both firm- and sector-level data. We pool French and Chinese firm-level export data to allow the separate identification of bilateral variables such as distance from destination-specific multilateral resistance and absorption terms. To measure market shares, we combine our firm-level export data with information on product-level absorption for European countries from the PRODCOM database. We demonstrate that adjusting trade flows by our correction term leads to substantial changes in the coefficient on standard bilateral gravity regressors such as distance, particularly in settings where individual French and Chinese exporters have significant market shares.

Moving on to the estimation using sectoral data, we to construct theory-consistent correction terms using data for French and Chinese exports to European markets. Again, correcting trade flows by these terms can lead to significant changes in gravity equation estimates, although the effects are less pronounced than at the firm level.

Our paper contributes to several strands of the literature on gravity equations. Anderson (1979), Eaton and Kortum (2002), Anderson and van Wincoop (2003), Chaney (2008) and

²Throughout the paper, we use the terms ‘sector’ and ‘product’ interchangeably. While our empirical analysis is based on product-level data classified according to the Harmonized System (HS), all of our theoretical results apply to both sector- and product-level data.

Melitz and Ottaviano (2008) show how to derive gravity equations from a number of different theoretical frameworks. For example, Chaney (2008) derives firm and aggregate gravity equations from a CES demand framework under monopolistic competition. We demonstrate how allowing for oligopolistic competition adds additional markup-based terms to otherwise identical equations. Melitz and Ottaviano (2008) use a setting with monopolistic competition and linear demand to generate an aggregate-level gravity equation. Similar to our approach, their framework generates variable destination-specific markups although it does not yield a firm-level gravity equation. We prefer to work with the more standard CES demand system as this allows a clean separation of the effects of oligopolistic behavior from markup variability arising from the shape of the demand function.³

We also contribute to part of the gravity literature that is concerned with obtaining consistent estimates of parameters of interest, such as distance elasticities. For example, Anderson and van Wincoop (2003) point out the need to control for multilateral resistance terms in gravity equations and Redding and Venables (2004) propose to include exporter and importer fixed effects to this end. Santos Silva and Tenreyro (2006) advocate the use of Poisson pseudo maximum likelihood (PPML) estimation techniques to address bias arising from heteroscedasticity in log-linearized models and to allow the inclusion of zero trade flows. Helpman, Melitz, and Rubinstein (2008) show how to account for the self-selection of firms into export markets when estimating aggregate gravity equations. We contribute to this literature by showing how to correct parameter bias arising from oligopolistic behavior by exporting firms at different levels of aggregation.

Third, within the last decade there has been revived interest in integrating oligopolistic competition into models of international trade, partially building on earlier contributions by the strategic trade policy literature (see Brander, 1995). For example, Edmond, Midrigan, and Xu (2015) study the gains from trade in the oligopolistic competition model of Atkeson and Burstein (2008). Eckel and Neary (2010) model the consequences of market integration in a setting with Cournot competition between multi-product firms. Parenti (2018) looks at the impact of trade liberalization in a model of imperfect competition where a few oligopolistic firms coexist with a monopolistically competitive fringe. Head and Mayer (2019) compare counterfactual predictions for the effects of freer trade across a number of modeling frameworks, including CES demand with monopolistic and oligopolistic competition and random coefficients discrete choice models with oligopolistic price setting. None of these papers investigates the consequences of oligopolistic behavior for the estimation of gravity equation,

³CES demand also seems a more natural starting point as it generates gravity for both individual firms and at more aggregate levels.

which is our key contribution.

Finally, in work concurrent to and independent of ours, 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 consistently estimate parameters necessary for the quantification of their model, they derive and estimate an aggregate gravity equation in oligopoly under the assumption that all industries are symmetric and each country hosts one firm per industry. Moreover, they have to take key parameters (such as price elasticities) from the existing literature, although the underlying estimation procedures are inconsistent with oligopolistic competition. By contrast, the firm- and industry-level gravity equations that we derive and estimate allow industries to differ in an arbitrary way and each country to host multiple (heterogeneous) firms. Moreover, we propose an adaptation of existing estimation procedures to obtain key parameter estimates in a way consistent with oligopolistic behavior.

The rest of this paper is organized as follows. In Section 2, we derive a firm-level gravity equation from a CES-demand framework with oligopolistic quantity competition. We also discuss how to deal with selection and heteroscedasticity in estimating our oligopolistic firm-level gravity equation. Next, we show in Section 3 how to modify the Feenstra-Broda-Weinstein estimation procedure to account for oligopolistic behavior and obtain demand and supply elasticity estimates. In Section 4, we derive our correction term for aggregate product-level trade flows. We also discuss how to adapt the methodology developed by Helpman, Melitz, and Rubinstein (2008) so as to deal with selection in the estimation of sector-level gravity under oligopoly. In Section 5, we describe the data sources and present the empirical results from our firm- and sector-level gravity estimations. In Section 6, we provide Monte Carlo simulations to evaluate the performance of our oligopoly correction term for sector-level regressions and that of our methods to deal with heteroscedasticity and selection. Finally, we conclude in Section 7. Appendix A collects proofs of our theoretical results. Results obtained when assuming price instead of quantity competition are presented in Appendix B. Appendix C contains lists of the countries present in our datasets.

2 Firm-Level Gravity in Oligopoly

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

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

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

Given these preferences, the direct and inverse demands for product $i \in \mathcal{J}_n(z)$ in country n are given by:

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

where E_n is total expenditure in country n , and

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

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

Each product $j \in \mathcal{J}_n$ 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.⁴ 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} . We allow for variable returns to scale and assume a functional form for costs of

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

⁴We focus on quantity competition here and present results for price competition in Appendix B.

where c_{in} is a firm-destination-specific cost shifter and $\tilde{\tau}_{in}$ a firm-destination-specific trade cost that takes the usual iceberg form.⁵

We assume throughout that the returns-to-scale parameters 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 will be unimodal.

Unlike in monopolistically competitive markets, firms take into account the impact of their actions on the CES-composite, Q_n , when setting quantities. 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

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

where

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

is the market share of firm i in destination n .

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

$$\mu_{in} = \frac{1}{\sigma} + \lambda \frac{\sigma-1}{\sigma} s_{in} \quad (4)$$

where $\mu_{in} \equiv (p_{in} - C'_{in}(q_{in})) / p_{in}$ is the Lerner index of product i in country n . 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 4, but for now, we assume Cournot conduct and set $\lambda = 1$.

Given the optimal markup in equation (4), firm i 's price is $p_{in} = c_{in} \tau_{in} q_{in}^{\gamma} / (1 - \mu_{in})$ and

⁵For 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.

the value of its sales in market n can be written as:

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

So far, we have not imposed any structure on trade costs, τ_{in} or the taste and cost shock terms, a_{in} and c_{in} . For comparison with the existing literature and to facilitate the exposition of our identifying assumptions, we now assume that the two shock terms can be decomposed log-linearly 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 assume that trade costs can be decomposed as $\log \tau_{in} = \beta X_{in} + \varepsilon_{in}^\tau$ where the X_{in} include 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 gravity equation-based research.⁶ Finally, we again assume a three-way decomposition of the trade cost error term, $\varepsilon_{in}^\tau = \varepsilon_i^\tau + \varepsilon_n^\tau + \eta_{in}^\tau$.

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

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

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

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

Note that under the assumption of monopolistic competition, the markup term involving μ 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 we control for firm and destination fixed effects (ζ_i and ξ_n) and that the usual orthogonality assumptions (explicitly or implicitly) made in the gravity literature hold.⁷

In the presence of strategic interaction between firms, however, the markup term will

⁶See, for example, Baier and Bergstrand (2007) and Rose (2000) on the effects of free trade agreements and currency unions, respectively, on trade flows.

⁷Specifically, for least-squares estimation of the log-linearized gravity equation, the orthogonality conditions are $\mathbb{E}(\eta_{in}^\tau | X_{in}) = \mathbb{E}(\varepsilon_{in}^a | X_{in}) = \mathbb{E}(\varepsilon_{in}^c | X_{in}) = 0$. Note that these assumptions allow for non-zero correlations between the bilateral variables and taste, production and trade cost shocks working through the firm- and destination-level-specific components (ε_i^a , ε_n^a , ε_i^c , ε_n^c , ε_i^τ and ε_n^τ). This is not a problem for consistent estimation as these components can be controlled for through firm and destination fixed effects. If the data used to estimate equation (6) contain a time dimension, it is also possible to allow for time-invariant bilateral elements in the error term which can be captured through bilateral fixed effects as is standard, for example, in the literature on the trade effects of preferential trade agreements (e.g., Baier and Bergstrand, 2007)

depend on firms' market shares and will thus be correlated with the regressors of interest, X_{in} ; not including this term will lead to an omitted variable bias. For example, we would expect firms to have lower market shares in more distant markets, *ceteris paribus*, and hence to charge lower markups there. This implies that $\log(1 - \mu_{in})$ will be higher in such markets, leading to a positive correlation between distance and the omitted variable.

Note that this problem is qualitatively different from those arising from other hard-to-observe gravity components such as expenditure (E_n), price indices (P_n) or firm-level marginal costs because these components can be controlled for by firm or destination fixed effects. By contrast, markups vary at the firm-destination level and the inclusion of bilateral 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 agreement.⁸

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

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

and estimate

$$\log \tilde{r}_{in} \equiv \log r_{in} - \frac{\hat{\sigma} - 1}{1 + \hat{\sigma}\hat{\gamma}} \log(1 - \hat{\mu}_{in}) = \xi_n + \zeta_i + \beta \frac{1 - \sigma}{1 + \sigma\gamma} X_{in} + \varepsilon_{in}. \quad (7)$$

Given our earlier orthogonality assumptions, using $\log \tilde{r}_{in}$ instead of $\log r_{in}$ as the dependent variable would yield a consistent estimate of $\beta \frac{1 - \sigma}{1 + \sigma\gamma}$. Using our estimates for σ and γ would then allow recovering the parameter of interest, β .⁹ This approach raises the question of how to estimate σ and γ . In the next section, we show how to adapt the estimation procedure by Feenstra (1994) and Broda and Weinstein (2006) to our setting with firm-level data and oligopolistic competition.

2.1 Estimation Challenges for Firm-level Gravity

Recall that our aim is to obtain consistent estimates of the coefficients on bilateral variables using either firm or sector level data. Above we showed that after subtracting a markup correction term from firm export values, we could estimate a standard gravity equation with a set of firm-product-year and destination-product-year fixed effects as well as the bilateral

⁸Having a time dimension in the data would not help either because markups would then vary by firm, destination and time.

⁹Note 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).

variables of interest.¹⁰

A first issue that arises is how to control for destination-specific fixed effects in a setting with firm-level export data. If we only have data for exports from a single country, it is immediately clear that we can no longer separate the impact of bilateral variables from the fixed effects. For example, if we use information on the exports of French firms only, standard bilateral variables such as common language become destination-specific as France is the only origin country in our data. Intuitively, we will not be able to distinguish whether firms' exports to a given destination are high because France and the country in question share a common language or because of other destination-specific factors such as a high price index or expenditure level. In order to address this issue, we follow Bas, Mayer, and Thoenig (2017) by combining two datasets on the exports of French and Chinese firms, respectively. This ensures that there is within-destination variation in the bilateral regressors of interest, enabling the use of destination fixed effects.

Secondly, we have so far ignored selection issues. In practice, most firms only export to a small subset of possible destinations for any given product. When estimating (7) in log-linear form, firm-product-destination observations with zero trade flows drop out. In the presence of export fixed cost $f_{on} > 0$ there is selection into exporting in our model: firms will be more likely to export positive amounts to a given destination if they experience a positive taste, production or trade cost shock for that destination, potentially creating a non-zero correlation with the regressors of interest. For example, firms selling in more distant foreign markets will be more likely to have received a positive shock, allowing them to operate in this more difficult environment. As consequence,

$$\mathbb{E}(\varepsilon_{in}^c | X_{in}, r_{in} > 0) \neq 0, \quad \mathbb{E}(\varepsilon_{in}^a | X_{in}, r_{in} > 0) \neq 0$$

Here, we adapt an approach proposed by Bas, Mayer, and Thoenig (2017) and restrict our estimation sample to the largest three French and Chinese firm in each product category as measured by *overall* product-specific exports. The basic idea is that these firms have high overall exports because they are very productive, produce high-quality products in general (high ε_i^a or ε_i^c) or have access to low-cost market access technologies (low ε_i^r). Such firms will tend to serve all or at least most available markets, making the destination-specific shocks less important for market entry decisions. We acknowledge that this is an imperfect solution

¹⁰Recall that we dropped the sector/product index (z) for most of our derivations and also ignored the time dimension to ease exposition. But these dimensions are of course present in our data, and hence price indices and expenditure levels will vary by destination, product and year, requiring the use of fixed effects at that level.

but simulation evidence by Bas, Mayer, and Thoenig (2017) shows that focusing on top exporters does indeed substantially reduce selection bias

Third, in the presence of heteroscedasticity, the log-linear gravity equation provides inconsistent coefficient estimates (Santos Silva and Tenreyro, 2006). In particular, we seek a consistent estimate of $\mathbb{E}(\tilde{r}_{in}|X_{in})$. Recall that $\tilde{r}_{in} = \exp(\xi_n + \zeta_i + \beta \frac{1-\sigma}{1+\sigma\gamma} X_{in}) \exp(\varepsilon_{in})$. Suppose that $\text{Var}(\exp(\varepsilon_{in}))$ depends on X_{in} . Then $\mathbb{E}(\varepsilon_{id})$ is a function of X_{id} and thus the error term is correlated with the control variables. A solution to this problem is to include zeros in our left-hand side variable and estimate (7) in multiplicative form:

$$\mathbb{E}(\tilde{r}_{in}|X_{in}) = \exp(\xi_n + \zeta_i + \beta \frac{1-\sigma}{1+\sigma\gamma} X_{in})$$

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.¹¹

3 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 $\mathbb{E}(\varepsilon_{in}^a \varepsilon_{i'n'}^c) = 0$ for all i, i' and n, n' , where $\varepsilon_{in}^c = \varepsilon_{in}^r + \varepsilon_{in}^c$. 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., $\mathbb{E}(a_{in} c_{in}) \neq 0$ is allowed). In particular, our method allows for a positive correlation between firm-level costs and quality (ε_i^a and ε_i^c) 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_n^a and ε_n^c). 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),

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

¹¹We include zero trade flows when estimating (7) on a product-by-product basis in Section 5.3

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:¹²

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

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-specific parts of the taste shocks:

$$\Delta^d \Delta^f \log a_{in} = (\varepsilon_{in}^a - \varepsilon_{i'n}^a) - (\varepsilon_{in'}^a - \varepsilon_{i'n'}^a).$$

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

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

Double differencing across firms and markets as above, we obtain

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

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

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

Note that as per our identifying assumption, the double-differenced cost and taste shocks are uncorrelated, yielding the following moment condition:

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

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

¹²In 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.

market shares:

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

and

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

The sample analogue of our moment condition is then given by

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

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 .¹³

Finally, 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 Chinese and French exporters (in terms of their overall exports) for any given 6-digit HS product. Finally, in order 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 products.

4 Sector-Level Gravity in Oligopoly

In this section, we study sector-level trade flows in the oligopoly model of Section 2. 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 a sector-level gravity equation that accounts for oligopolistic behavior.

Oligopoly analysis in a given destination market. Consider sector z in destination n . Dropping reference to both z and n to ease notation, we define the market-level aggregator H as

$$H \equiv Q^{\frac{\sigma-1}{\sigma}} = \sum_{j \in \mathcal{J}} a_j^{\frac{1}{\sigma}} q_j^{\frac{\sigma-1}{\sigma}}$$

¹³In practice, this means that we need to observe a sufficiently large number of firms selling in the same sector in at least three different markets.

and firm i 's type T_i as

$$T_i \equiv a_i^{\frac{1}{\sigma}} \left(\frac{\alpha E}{c_i \tau_i} \frac{\sigma - 1}{\sigma} \right)^{\frac{\sigma-1}{\sigma(1+\gamma)}}. \quad (8)$$

Plugging these definitions into equation (2), making use of equation (3), and rearranging, we obtain:

$$1 - \lambda s_i = s_i^{\frac{1+\sigma\gamma}{\sigma-1}} \left(\frac{H}{T_i} \right)^{\frac{\sigma(1+\gamma)}{\sigma-1}}, \quad (9)$$

where λ is the conduct parameter introduced in Section 2. As the left-hand side is strictly decreasing in s_i and the right-hand side is strictly increasing in s_i , the equation has a unique solution in s_i , which we denote $S(T_i/H, \lambda)$ —the *market-share fitting-in function*. It can easily be 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:

$$\sum_{i \in \mathcal{I}} S\left(\frac{T_i}{H}, \lambda\right) = 1. \quad (10)$$

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 $H^*(\lambda)$ is the unique solution to equation (10). 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 (9). From equation (3), firm i 's equilibrium output is given by*

$$q_i^*(\lambda) = a_i^{-\frac{1}{\sigma-1}} (s_i^*(\lambda) H^*(\lambda))^{\frac{\sigma}{\sigma-1}}.$$

Proof. See Appendix A.1 □

The first-order approach to sector-level gravity. Let $\mathcal{E} \subsetneq \mathcal{I}$ denote the subset of exporters from country o that sell in the destination market n . The aggregate exports of those firms to market n are given by

$$\underbrace{\sum_{i \in \mathcal{E}} s_i^*(\lambda)}_{\equiv s_o^*(\lambda)} \times \alpha E.$$

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 $s_o^*(1)$. Our approach therefore entails approximating $s_o^*(1)$ at the first order.

As we show in the following, the approximation relies on two versions of the Herfindahl-Hirschman index (HHI), namely the HHI of all firms selling in the destination market n ,

$$\text{HHI}_n(\lambda) \equiv \sum_{j \in \mathcal{J}} (s_{jn}^*(\lambda))^2,$$

and the (normalized) HHI of all those exporters in country o that sell in the destination market n ,

$$\text{HHI}_{on}(\lambda) \equiv \sum_{j \in \mathcal{E}} \left(\frac{s_{jn}^*(\lambda)}{s_{on}^*(\lambda)} \right)^2.$$

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 export country o is given by*

$$\log s_{on}^*(\lambda) = \log s_{on}^*(0) + \frac{\sigma - 1}{1 + \sigma\gamma} \left[\text{HHI}_n(\lambda) - s_{on}^*(\lambda) \text{HHI}_{on}(\lambda) \right] \lambda + o(\lambda).$$

Proof. See Appendix A.2. □

The proposition shows that the logged 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 as well as the concentration among the country- o exporters.

This result motivates the following approximation:

$$\log s_{on}^*(1) \simeq \log s_{on}^*(0) + \frac{\sigma - 1}{1 + \sigma\gamma} \left[\text{HHI}_n(1) - s_{on}^*(1) \text{HHI}_{on}(1) \right]. \quad (11)$$

In Cournot oligopoly, the logged sector-level exports from country o to destination market n are given by

$$\begin{aligned} \log r_{on} &= \log(\alpha_n E_n) + \log s_{on}^*(1) \\ &\simeq \log(\alpha_n E_n) + \log s_{on}^*(0) + \frac{\sigma - 1}{1 + \sigma\gamma} \left[\text{HHI}_n(1) - s_{on}^*(1) \text{HHI}_{on}(1) \right], \end{aligned} \quad (12)$$

where the second line follows from the approximation in equation (11).

4.1 Estimation Challenges for Sector-level Gravity

In order to consistently estimate gravity equations at the sector-level, we need to correct for selection of high-quality, low-cost firms into high-trade-cost destinations. For estimating firm-level gravity, we followed Bas, Mayer, and Thoenig (2017) in focusing on the top 3 exporters from the origin country. This is, of course, not possible when using sector-level data. Here, we instead adapt the methodology developed by Helpman, Melitz, and Rubinstein (2008) (from now on HMR) for gravity estimation with heterogeneous firms and constant markups to oligopoly. To avoid multiplicity of equilibria, we assume that, at the entry stage, firms behave as under monopolistic competition—and thus take the aggregate price index as given.

Under this hypothesis, firm i from origin o enters market n as long as

$$\pi_{in} = \frac{r_{in}}{\sigma} \geq f_{on}$$

where

$$r_{in} = \left(\frac{\tau_{in}}{1 - \mu} \right)^{\frac{1-\sigma}{1+\sigma\gamma}} \varphi_i^{\frac{\sigma-1}{1+\sigma\gamma}} (P_n^{\sigma-1} \alpha_n E_n)^{\frac{1+\gamma}{1+\sigma\gamma}}$$

and $\varphi_i \equiv \frac{a_i^{\frac{1+\gamma}{\sigma-1}}}{c_i}$ is firm productivity, summarizing the effect of firm quality and costs, which is now constrained to be firm-specific (instead of firm-destination specific). We define the ratio of variable export profits to fixed export costs as $Z_{on} = \Pi_{on}(\varphi_i^{\max})/f_{on}$ for the highest-productivity firm from o . Note that we observe positive aggregate exports from o to n as long as $Z_{on} \geq 1$. Moreover, Z_{on} is decreasing in fixed trade costs f_{on} and variable trade costs τ_{on} .

Taking logs, we obtain:

$$\begin{aligned} \log Z_{on} &= \frac{\sigma-1}{1+\sigma\gamma} \log \varphi_o^{\max} + \frac{1+\gamma}{1+\sigma\gamma} \log(P_n^{\sigma-1} \alpha_n E_n) + \\ &\quad + \frac{1-\sigma}{1+\sigma\gamma} \log \tau_{on} - \log f_{on} + \frac{1-\sigma}{1+\sigma\gamma} \log(1-\sigma) \end{aligned}$$

We now put some more structure on the fixed and variable export costs. We assume that $\log f_{on} = f_o + f_n + \kappa \chi_{on} + \epsilon_{on}^f$ with $\epsilon_{on}^f \sim N(0, \sigma_f^2)$ and $\log \tau_{on} = \tau_o + \tau_n + \gamma \psi_{on} + \epsilon_{on}^\tau$ with $\epsilon_{on}^\tau \sim N(0, \sigma_\tau^2)$. Here χ_{on} and ψ_{on} are observable trade costs, (f_o, τ_o) are origin-specific constants and (f_n, τ_n) are destination-specific constants. We denote the vector of observable bilateral fixed and variable trade cost by X_{on} and we define $\rho_{od} = \mathbb{E}(\log Z_{od}|X_{on}) = \Pr(\text{Exports}_{on} > 0|X_{on})$.

Then, we obtain the empirical specification for the extensive-margin decision with sector-level export data as the following Probit model:

$$\rho_{on} = \Phi(\xi_o + \xi_n + \beta X_{on}),$$

where ξ_o is an origin fixed effect and ξ_n is a destination fixed effect. This corresponds to the first stage of the HMR estimation procedure.

In the second stage, firms decide about export quantities and here we maintain again the assumption of oligopolistic behavior. As explained above, we approximate the aggregate market share of country o in destination n under oligopoly around the one under monopolistic competition conduct ($S_{od}^*(0)$). This variable is given by:

$$s_{on}^*(0) = \begin{cases} \left(\frac{\sigma-1}{\sigma} \alpha_n E_n \right)^{\frac{\sigma-1}{1+\sigma\gamma}} P_n^{\frac{-\sigma(1+\gamma)}{1+\sigma\gamma}} \tau_{on}^{\frac{1-\sigma}{1+\sigma\gamma}} \left(\sum_{i \in \mathcal{J}_{od}} \varphi_i^{\frac{\sigma-1}{1+\sigma\gamma}} \right) & \text{if } Z_{on} > 1 \\ 0 & \text{otherwise} \end{cases} \quad (13)$$

In this expression, we need to proxy for average firm productivity of firms from origin o that are active in destination n , $\mathbb{E} \left[\left(\sum_{i \in \mathcal{J}_{on}} \varphi_i^{\frac{\sigma-1}{1+\sigma\gamma}} \right) | Z_{on} > 1 \right]$. This term is correlated with observable trade barriers due to selection of high-type firms into destinations with higher trade costs and thus not including it in the regression would lead to biased estimates of our coefficients of interest.

We proxy for (the log of) this term in a non-parametric way with a polynomial in $\log \hat{Z}_{od} = \Phi^{-1}(\hat{\rho}_{od})$. Finally, we also need to proxy for the unobserved variable trade cost $\mathbb{E}(\epsilon_{on}^\tau | Z_{on} > 1) \neq 0$ which is also correlated with observed trade flows. For this, we use the inverse Mills ratio $\lambda_{on} = \phi(\log \hat{Z}_{on}) / \Phi(\log \hat{Z}_{on})$, which corresponds to the standard Heckman selection correction.

Putting everything together, we obtain the following sector-level gravity equation for positive export flows:

$$\begin{aligned} \log \tilde{r}_{on} = & \zeta_o + \zeta_n + \beta \frac{1-\sigma}{1+\sigma\gamma} X_{on} \\ & + \omega_1 \log \hat{Z}_{on} + \omega_2 \log \hat{Z}_{on}^2 + \omega_3 \log \hat{Z}_{on}^3 + \omega_4 \lambda_{on} + \eta_{on}, \end{aligned} \quad (14)$$

where

$$\log \tilde{r}_{on} \equiv \log r_{on} + \frac{\hat{\sigma} - 1}{1 + \hat{\sigma}\hat{\gamma}} s_{on} HHI_{on}$$

is the value of the export flows from o to n , purged from oligopolistic market power effects.

We can thus consistently estimate the intensive-margin trade elasticity $\beta \frac{1-\sigma}{1+\sigma\gamma}$.

5 Empirical Implementation

In this section, we show how to implement our methods for firm- and sector-level gravity estimations empirically. We first discuss our datasets. We then present basic descriptive statistics on our data and the GMM estimates for σ and γ . Finally, we run firm and sector-level regressions with and without oligopoly correction terms and investigate if and under which circumstances ignoring oligopolistic behavior can lead to quantitatively important coefficient bias.

5.1 Data Sources

As discussed, 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 dataset, we observe all the products and destinations to which a firm exports, as well as the quantity and value of the underlying flow. Both datasets record export data at the 8-digit level but we aggregate this information up to the 6-digit level of the Harmonised System (HS) which is the lowest level at which the two national classifications are comparable with each other. Because we observe both values and quantities, we can compute unit values which are a commonly used proxy for prices in the trade literature.

A final challenge for our firm-level analysis is to obtain information on market shares at a level of disaggregation that is sufficient to capture meaningful strategic interaction between firms. To our knowledge, the only suitable database here is Eurostat's PRODCOM database which allows computation of absorption at a level at, or close to, HS 6-digit.¹⁴ Together with the information on the value of destination-product-level exports by individual French and Chinese firms, this allows the computation of market shares at the HS-6digit "plus" level, where "plus" means that some products have to be aggregated further to make the classifications of the trade and production data consistent with each other. The downside of

¹⁴Absorption is defined as domestic production + imports - exports and thus is the theoretical equivalent to E_n 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 6digit plus", that is comparable over time and also 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 1-to-1 mapping between them. See their paper for an in-depth discussion of the procedure.

using PRODCOM is that absorption data is only available for approximately thirty European countries. After combining our data sources, we end up with information on export values and quantities as well as market shares for 32 European destination markets, approximately 1,800 products and 250,000 exporters for the period 2000-2010.¹⁵

For our sector-level gravity regressions, we require product-level data on the value of bilateral exports, absorption data for the computation of market shares and exporter-destination-product-specific HHIs. To make the estimation sample consistent with our firm-level regressions, we aggregate our firm-level data at the 6-digit HS level and use the firm level data to compute exporter HHIs.¹⁶ Finally, we source information on bilateral distance from CEPIL.

5.2 Descriptive Statistics

The key determinants of our oligopoly correction term are firm-level market shares as well as estimates for demand elasticities (σ) and returns to scale (γ). The first line of Table 1 presents information on the firm-level market shares for the French and Chinese exporters in our firm-level dataset. The average market share across the approximately 14 million firm-destination-product-year combinations in our data is small at 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 necessarily 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. The remaining columns of Table 1 focuses on such firms. The second line begins by showing descriptive statistics for the top exporters (i.e., the French and Chinese firms with the largest total export values for a given 6-digit product and year). The average top-exporter market share is around 6%, substantially larger than the average exporter’s market share. Moreover, at the 95th percentile the top

¹⁵Possibly because of measurement issues in PRODCOM, we occasionally observe cases where absorption is smaller than a firm’s exports to a given market, resulting in market shares larger than one; in such cases, we winsorize market shares to 0.95.

¹⁶Researchers may not generally have firm-level datasets for several countries available. As an alternative, we have used PRODCOM for absorption data and combined these 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 destination-specific Herfindahl indices 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 this data is still at the product level rather than the level of individual exporters, it is 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 datasets were similar to the ones based on aggregated firm-level information. (Results are available upon request).

firm enjoys a market share of almost 30%. In the third column we show the market shares for the sample of the top 3 exporters (i.e., the largest three French and Chinese firms in terms of total export values for a given 6-digit product and year). The mean market share in this sample is around 3.9% and at the 95th percentile it equals around 18%. In the final column, we present the cumulative market shares of the top-3 exporters. For them, the average cumulative market share is equal to 7.30% and at the 95th percentile, they have a total market share of 33.4%. Thus, there is a small set of exporters with large market shares in most markets.

Table 1: Summary Statistics for French and Chinese Firm-Level Market Shares

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

Note: annual Chinese and French firm-level market share data for period 2000-2010.

The picture that aggregate exports are very concentrated across a few firms is confirmed by Table 2, which presents summary statistics on exporter HHIs and sectoral aggregate market shares of Chinese and French firms for our HS 6-digit product-level sample. The mean exporter HHI is 0.55 and at the 90th percentile a single firm accounts for the total market share of each country. Moreover, the mean aggregate market share of China and France in each destination is around 9% and at the 90th percentile the market share reaches 24%. Thus, in many cases, Chinese and French firms have substantial market power in individual markets. The third column of Table 2 provides summary statistics for the sector-level markup correction term (computed for $\sigma = 5$ and $\gamma = 0$). It has a mean of 0.18 and reaches from 0 at the 10th percentile to 0.4 at the 90th percentile.

Table 3 shows descriptive statistics for our estimates of σ and γ . As discussed, we constrain coefficient estimates to be identical within 2-digit HS codes 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 5.39 and a median of 3.74. Reassuringly, these

Table 2: Summary Statistics for Sector-Level Market Shares and Exporter HHIs

	Exporter Herfindahl	Destination Market Share	Markup Correction Term
Mean	0.55	9%	0.18
5th pctl	0.08	0.01%	0.0004
10th pctl	0.13	0.06%	0.001
Median	0.50	2%	0.03
90th pctl	1	24%	0.41
95th pctl	1	42%	0.82

Note: data for 6-digit HS sectors. Sample 2000-2010. Markup correction term computed for $\sigma = 5$, $\gamma = 0$.

numbers are very similar to estimates at comparable levels of aggregation estimated in the literature (e.g., Broda and Weinstein, 2006).

Table 3: Price Elasticities and Returns-to-Scale Estimates – Cournot Competition

	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
HS 2-digit products	78	78

Note: 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 products.

5.3 Firm-Level Gravity Estimation Results

We now turn to the estimation of our firm-level gravity equations with and without correction for oligopoly bias. In all firm-level regressions, we consider the top-3 exporters of any given 6-digit product as potential exporters of that product to any given destination and fill in the zero export flows if they do not export the product to a destination. As a first step, we pool across all firms in our data and estimate equations 6 and 7 via PPML using a full set of firm-product-year and destination-product-year fixed effects. We aim at identifying the intensive-margin trade elasticity given by $\beta \frac{1-\sigma}{1+\sigma\gamma}$. As our main regressor of interest, we include bilateral distance. Note that because our destination countries are all in Europe, there is insufficient variation to include other commonly used indicators such as dummies for

common language, or policy-related variables, such as membership in a free trade agreement and bilateral tariffs.¹⁷

Table 4: Firm-Level Gravity Estimates, $\sigma = 5$ and $\gamma = 0$.

Regressor	PPML w/ corr	PPML w/o corr	OLS w/ corr	OLS w/o corr
log distance	-1.518*** (0.220)	-0.874*** (0.021)	-0.275*** (0.015)	-0.232*** (0.014)
$\hat{\beta}_{distance}$	0.375	0.218	0.069	0.058
Observations	11,955,786	11,955,786	708,392	708,386
R-squared			0.05	0.06
Firm-year FE	YES	YES	YES	YES
Product-dest.-year FE	YES	YES	YES	YES

Note: Firm-level data, pooled across sectors. Results for top 3 exporters. Cournot model with $\sigma = 5$ and $\gamma = 0$. Standard errors in brackets, clustered at the destination-year level.

Table 5: Firm-Level Gravity Estimates, $\sigma = 5.39$ and $\gamma = 0$.

Regressor	PPML w/ corr	PPML w/o corr	OLS w/ corr	OLS w/o corr
log distance	-1.522*** (0.221)	-0.874*** (0.0210)	-0.279*** (0.015)	-0.232*** (0.014)
$\hat{\beta}_{distance}$	0.354	0.199	0.064	0.053
Observations	11,955,786	11,955,786	708,392	708,386
R-squared			0.05	0.06
Firm-year FE	YES	YES	YES	YES
Product-dest.-year FE	YES	YES	YES	YES

Note: Firm-level data, pooled across sectors. Results for top 3 exporters. Cournot model with mean of estimated σ and $\gamma = 0$. Standard errors in brackets, clustered at the destination-year level.

Table 4 presents the results using $\sigma = 5$ and $\gamma = 0$ (CRS) for constructing the markup correction term. Column (1) presents the PPML estimate for the specification including the markup correction term, while column (2) presents the PPML estimate without the markup correction. As becomes immediately clear, without the markup correction term, the distance coefficients is heavily downward biased in absolute magnitude. While the point estimate in column (1) is -1.52, it is -0.87 in column (2). This confirms our theoretical insight that the distance coefficient suffers a substantial attenuation bias because firms systematically

¹⁷The destination countries in our sample were either already EU member states or had implemented free trade agreements 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 before 2010 and EU external tariffs for imports from China only had industry variation. Thus, there is no variation in the FTA dummy or in tariffs that is not absorbed by our firm-product-year dummies. Likewise, there is insufficient variation to include an indicator common language which is identified only for the pair France-Belgium and France-Switzerland.

Table 6: Firm-Level Gravity Estimates, $\sigma = 5.39$ and $\gamma = 0.34$.

Regressor	PPML w/ corr	PPML w/o corr	OLS w/ corr	OLS w/o corr
log distance	-1.201*** (0.118)	-0.874*** (0.0210)	-0.248*** (0.0142)	-0.231*** (0.0136)
$\hat{\beta}_{distance}$	0.792	0.577	0.160	0.149
Observations	11,955,786	11,955,786	708,392	708,386
R-squared			0.05	0.06
Firm-year FE	YES	YES	YES	YES
Product-dest.-year FE	YES	YES	YES	YES

Note: Firm-level data, pooled across sectors. Results for top 3 exporters. Cournot model with mean of estimated σ and γ . Standard errors in brackets, clustered at the destination-year level.

reduce their markups in markets where they face higher variable trade costs and thus have lower market shares. As a consequence, firm-level export values decrease by less than they would have decreased under constant markups. The point estimate on distance corresponds to $(\sigma - 1)\hat{\beta}_{distance}$, where $\hat{\beta}_{distance}$ is the fundamental trade cost elasticity of distance. Given values of σ equal to 5 and γ equal to zero, the implied values for this coefficient are 0.354 in column (1) and 0.199 in column (2), implying a downward bias in coefficient magnitude of around 44 percent. Note that while distance is not a policy variable, a very similar attenuation bias would arise for any iceberg-type variable, such as ad-valorem tariffs or transport costs.

In columns (3) and (4) we report estimates for the log-linear versions of equations (7) and (6) with OLS estimated on the positive trade flows. In both cases, the distance coefficient is heavily down-ward biased due to the heteroscedasticity bias.

In Table 5 we report results for setting σ equal to its mean estimated value ($\sigma = 5.39$) when computing the markup correction term, while keeping γ equal to zero. Results remain very similar to the previous specification, with the exception of the estimate in column (3), where the estimate of the distance coefficient in the log-linear OLS specification with the markup correction term now turns significantly negative. Finally, in Table 6, we report results using the mean estimated σ and γ . The distance coefficient is now a bit smaller in absolute magnitude (-1.2), but still substantially larger than without the markup correction term, implying a coefficient bias of around 30%. The implied fundamental trade cost elasticity of distance is larger in this case and corresponds to 0.79 in the specification including the markup correction term, compared to 0.58 without the correction. The reason is that we estimate significant decreasing returns to scale for the average sector ($\gamma = 0.34$), implying that marginal costs increase in the exported quantity. Thus, given the actually observed trade value, markups decrease more for more distant markets than with constant marginal costs.

Correspondingly, for a given export value, trade costs must be more sensitive to distance if we hold markups constant.

So far, we have reported pooled estimates and constrained the coefficients to be the same across products. We now turn to an estimation product by product, since oligopolistic market structure is plausibly more important for some products compared to others and thus the markup bias may also vary across products. We illustrate the importance of oligopoly bias at the product level by estimating firm-level gravity equations separately for each of the 78 HS 2-digit product in our data, the level of variation of σ and γ (pooling observations across 6-digit products within a given HS 2 product). Table 7 reports summary statistics for the median coefficient estimates across products with and without the correction for oligopoly bias for all three cases ($\sigma = 5, \gamma = 0$, σ estimated, $\gamma = 0$, (σ, γ) estimated). In all specifications, the median point estimate on 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 markup correction the median distance coefficient is only -0.51. The absolute percentage bias in the coefficient, defined as $abs(\frac{\beta_{distance,w/corr} - \beta_{distance,wo/corr}}{\beta_{distance,w/corr}})$, varies across sectors. For a product at the 10th percentile, this bias is around 10%, but it increases to 160% for a product at the 90th percentile. Thus, for some products the oligopoly bias in the distance coefficient is much larger than the pooled estimates would suggest.

Table 7: Firm-level Gravity Estimates by 2-digit Product

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

Note: Firm-level data. Coefficients by 2-digit HS product for top 3 exporters. Cournot model.

5.4 Sector-Level Gravity Estimation Results

While firm-level trade data are increasingly becoming available to researchers, in many cases gravity estimations are still based on more aggregate types of data, such as sector-level trade flows. Of course, the issue of coefficient bias due to oligopolistic behavior does not go away

at this level as aggregate exports are simply the sum of underlying firm-level exports. In this section, we use the HHI-based correction term proposed in Section 4 to investigate the quantitative importance of oligopoly bias for sector-level regressions. Like at the firm level, we want to identify the intensive-margin trade elasticity $\beta \frac{1-\sigma}{1+\sigma\gamma}$ and thus we expect to obtain similar coefficient estimates. (Observe that the aggregate trade elasticity is not constant in our model.¹⁸)

We first present results for the pooled regressions. We balance the dataset by adding all zero trade flows at the 6-digit product-destination-year level. For computational reasons, we focus on the year 2010. (Similar results are obtained for other years). First, we run regressions with and without the markup correction term but without correcting for selection into exporting. As we have shown in Section 4, controlling for selection is in theory necessary to correctly identify of the intensive-margin elasticity. We thus regress bilateral sector-level exports on bilateral distance as well as product-origin and product-destination fixed effects.

$$\log \tilde{r}_{on} = \zeta_o + \xi_n + \beta \frac{1-\sigma}{1+\sigma\gamma} X_{on} + \eta_{on}$$

Table 8 reports the results for $\sigma = 5$ and $\gamma = 0$. In columns (1) and (2), we apply the PPML estimator, which allows us to include the zeros in bilateral trade flows, while in columns (3) and (4) we apply the OLS estimator on positive trade flows. In principle, we expect to recover the same elasticity as at the firm level. However, the PPML coefficient estimates on (log) distance in columns (1) and (2) are much smaller in absolute magnitude: -0.26 without the markup correction term and -0.19 with the correction term. Thus, the non-random nature of the zeros leads to a heavy attenuation bias in the PPML estimates. The OLS results presented in columns (3) and (4) estimated on the selected sample of positive trade flows instead look more reasonable: the coefficient on distance is -1.12 without and -1.26 when including the markup correction term, which is somewhat smaller than the firm-level estimates. Thus, there is plausibly a selection bias in the OLS estimates, which combines several opposing effects: at the extensive margin, we see fewer firms exporting to more distant markets and thus exports should fall more with distance; at the selection margin, the surviving firms are positively selected and thus exports should fall less with distance. Moreover, also at the sector level, a significant markup bias arises, even though the bias in coefficient magnitudes from not including the markup correction term seems smaller than at

¹⁸In the absence of Pareto-distributed productivity even with constant markups the aggregate trade elasticity varies along the firm-size distribution and does not usually have a closed-form solution (Bas, Mayer, and Thoenig, 2017). In addition, in our model markups for any given product vary at the firm-destination level.

the firm level (around 10%).

Table 8: Sector-level Gravity Estimates without Controlling for Selection

Regressor	PPML w/o corr	PPML w/ corr	OLS w/o corr	OLS w/ corr
log distance	-0.263 (0.168)	-0.186 (0.147)	-1.128*** (0.195)	-1.260*** (0.216)
Observations	107,064	107,064	66,563	66,563
R-squared			0.314	0.285
Product-origin FE	YES	YES	YES	YES
Product-dest. FE	YES	YES	YES	YES

Note: Sector-level data. Cournot model with $\sigma = 5$ and $\gamma = 0$. Standard errors clustered at destination level

We now apply the HMR methodology in order to correct for selection into exporting in addition to the markup bias and to obtain correct estimates of the intensive-margin trade elasticity. To apply the HMR method, we first estimate the propensity to export (extensive margin) using a Probit estimator. We include 2-digit-product-origin and 2-digit-product-destination fixed effects.¹⁹ To proxy for fixed market entry costs between origin o and destination n , we also follow HMR: we add dummies for the business startup time and the startup cost being above the sample median for both the origin and the destination country. By assumption, these variables only impact on the fixed export cost but not on the iceberg-type trade costs and thus exclusively affect firms' entry decision but not their quantity choice conditional on entry. The source for these variables is the Worldbank's Doing Business Database. We report results for the first-stage Probit regression in Table 9. As expected, the dummies for high business startup cost and long business startup time are significantly negatively associated with the probability to export at the sector level. Moreover, distance also affects the sectoral export propensity significantly negatively.

We next turn to the second-stage HMR results, which refer to the intensive-margin export decisions at the sector level. Remember that the second-stage regression is specified as

$$\begin{aligned} \log \tilde{r}_{onz} = & \zeta_{oz} + \xi_{nz} + \beta \frac{1 - \sigma}{1 + \sigma\gamma} X_{on} \\ & + \omega_1 \log \hat{Z}_{onz} + \omega_2 \log \hat{Z}_{onz}^2 + \omega_3 \log \hat{Z}_{onz}^3 + \omega_4 \lambda_{onz} + \eta_{onz} \end{aligned}$$

¹⁹Using 6-digit product-origin and product-destination fixed effects is not computationally feasible with the Probit estimator. However, results using a linear probability model indicate hardly any changes in the point estimates when adding these more detailed fixed effects compared to 2-digit product-origin and product-destination fixed effects.

Table 9: Sector-level Gravity Estimates – Export Propensity

Regressor	Export > 0
log distance	-0.420*
high startup cost	-1.340***
long startup time	-2.102***
Observations	107,064
Product-origin FE	YES
Product-dest. FE	YES

Note: Sector-level data. HMR first-stage Probit regression of propensity to exporting. Standard errors clustered at destination level in parentheses

Here, $\tilde{r}_{onz}$ corresponds to the markup-corrected export revenue of origin country o in destination n in sector z . X_{on} is log distance and $\hat{Z}_{onz}$ is the predicted export-profit-to-fixed-cost ratio for the highest-type firm. We include a second or third-order polynomial in this variable to proxy for the exporting firms' type mix in each origin-destination-sector combination. Finally, λ_{onz} is the inverse Mills ratio, which controls for unobserved variable trade costs. As usually, we first report results for this regression when the correction term is computed with $\sigma = 5$ and $\gamma = 0$ (CRS) in Table 10. In columns (1) and (2), we report results without and with the markup correction term, but only including the inverse Mills ratio (columns labeled *Heck*), in columns (3)-(6) we instead implement the full HMR procedure. Columns (3) and (4) include a quadratic function of $\hat{Z}_{onz}$, while columns (5)-(6) include a third-order polynomial. In all three cases, the coefficient on distance in the specification including the markup correction term is significantly larger in magnitude compared to the one without markup correction. In our preferred specifications (columns (4)-(6)), the coefficient on log distance is around -1.284, compared to -1.15 without correction, corresponding to around 10% downward bias in absolute magnitude without correction. The corresponding point estimates of the fundamental trade cost elasticity to distance are 0.32 vs. 0.29. Finally, note that the inverse Mills ratio and the polynomial terms in $\hat{Z}_{onz}$ have the expected signs in columns (4)-(6) and are mostly significant. Lower unobserved trade barriers (a higher inverse Mills ratio) and higher average types both increase the value of sectoral exports to a given destination.

In Table 11, we repeat the same specifications, using the mean estimated σ of 5.39 and $\gamma = 0$ to compute the markup correction term. Not surprisingly, our results are hardly affected by this modification. Finally, in Table 12, we use both the mean estimates of σ and γ to correct for oligopolistic markups. While the point estimates on log distance hardly change, the estimated fundamental distance coefficient does become significantly larger in this

case. The reason is that we estimate significantly decreasing returns to scale for the average sector ($\gamma = 0.34$, implying that marginal costs increase in the exported quantity. Thus, given the actually observed trade value, markups decrease more for more distant markets than with constant marginal costs. Correspondingly, for a given export value, trade costs must be more sensitive to distance if we hold markups constant.

Table 10: Sector-level Gravity Estimates – intensive margin, $\sigma = 5$, $\gamma = 0$

Regressor	Heck w/o	Heck w/	HMR ² w/o	HMR ² w/	HMR ³ w/o	HMR ³ w/
log distance	-1.189*** (0.198)	-1.327*** (0.220)	-1.151*** (0.190)	-1.284*** (0.209)	-1.150*** (0.193)	-1.284*** (0.212)
inv mills	-0.121 (0.165)	-0.130 (0.182)	0.678*** (0.169)	0.799*** (0.192)	0.639** (0.309)	0.805** (0.344)
log $\hat{Z}$			0.907*** (0.265)	1.067*** (0.293)	0.736 (1.305)	1.097 (1.396)
log $\hat{Z}^2$			-0.103* (0.0565)	-0.125** (0.0621)	-0.0297 (0.534)	-0.137 (0.567)
log $\hat{Z}^3$					-0.0102 (0.0693)	0.00176 (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	0.302	0.275	0.304	0.277	0.304	0.277
Origin-product FE	YES	YES	YES	YES	YES	YES
Destination-product FE	YES	YES	YES	YES	YES	YES

Note: Sector-level data. Cournot model with $\sigma = 5$ and $\gamma = 0$. Standard errors clustered at destination level in parentheses.

Finally, we report regression results for estimating the sector-level regressions separately by 2-digit sector (pooling over 6-digit products within a given 2-digit product), which is the level of variation of σ and γ . In Table 13 we report the median point estimates for the distance coefficients and the implied fundamental trade cost elasticity of distance with and without including the markup correction. In addition, we report the absolute percentage bias in the coefficient point estimates for a sector at the 10th percentile, at the median and at the 90th percentile of the bias. Again we find that without oligopoly correction there is a downward bias in the distant coefficients. While the bias is relatively small in the median sector, it increases to around 40% at the 90th percentile. Moreover, the size of the coefficient bias increases systematically with the exporter HHI index. When running the bias on country-pair fixed effects and the exporter HHI, we find a very strong positive correlation:

$$pct.bias_s = \alpha_{od} + 0.794^{***} \times HHI_{ods}$$

Thus, in sectors with highly concentrated exports, sector-level gravity estimates are heav-

Table 11: Sector-level Gravity Estimates – intensive margin, $\sigma = 5.39$, $\gamma = 0$

Regressor	Heck w/o	Heck w/	HMR ² w/o	HMR ² w/	HMR ³ w/o	HMR ³ w/
log distance	-1.189*** (0.198)	-1.341*** (0.222)	-1.151*** (0.190)	-1.297*** (0.211)	-1.150*** (0.193)	-1.297*** (0.214)
inv mills	-0.121 (0.165)	-0.131 (0.184)	0.678*** (0.169)	0.810*** (0.194)	0.639** (0.309)	0.822** (0.348)
log $\hat{Z}$			0.907*** (0.265)	1.083*** (0.295)	0.736 (1.305)	1.132 (1.405)
log $\hat{Z}^2$			-0.103* (0.0565)	-0.127** (0.0626)	-0.0297 (0.534)	-0.148 (0.570)
log $\hat{Z}^3$					-0.0102 (0.0693)	0.00293 (0.0736)
$\hat{\beta}_{distance}$	0.277	0.313	0.268	0.302	0.268	0.302
Observations	60,662	60,662	60,662	60,662	60,662	60,662
R-squared	0.302	0.272	0.304	0.274	0.304	0.274
Origin-product FE	YES	YES	YES	YES	YES	YES
Destination-product FE	YES	YES	YES	YES	YES	YES

Note: Sector-level data. Cournot model with mean of estimated σ and $\gamma = 0$. Standard errors clustered at destination level in parentheses.

Table 12: Sector-level Gravity Estimates – intensive margin, $\sigma = 5.39$, $\gamma = 0.34$

Regressor	Heck w/o	Heck w/	HMR ² w/o	HMR ² w/	HMR ³ w/o	HMR ³ w/
log distance	-1.189*** (0.198)	-1.243*** (0.206)	-1.151*** (0.190)	-1.202*** (0.197)	-1.150*** (0.193)	-1.202*** (0.200)
inv mills	-0.121 (0.165)	-0.124 (0.171)	0.678*** (0.169)	0.725*** (0.178)	0.639** (0.309)	0.703** (0.322)
log $\hat{Z}$			0.907*** (0.265)	0.969*** (0.275)	0.736 (1.305)	0.876 (1.338)
log $\hat{Z}^2$			-0.103* (0.0565)	-0.112* (0.0586)	-0.0297 (0.534)	-0.0714 (0.546)
log $\hat{Z}^3$					-0.0102 (0.0693)	-0.00556 (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	0.302	0.292	0.304	0.294	0.304	0.294
Origin-product FE	YES	YES	YES	YES	YES	YES
Destination-product FE	YES	YES	YES	YES	YES	YES

Note: Sector-level data. Cournot model with mean of of estimated σ and γ . Standard errors clustered at destination level in parentheses.

ily biased.

Table 13: Sector-level estimates by 2-digit product.

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

6 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 in which the estimation challenges (due to oligopolistic behavior and selection into export markets) discussed in Sections 2.1 and 4.1 are present. We then apply our firm- and industry-level estimators to that data-set and confirm that our oligopoly correction terms significantly improve the accuracy of our estimates.

6.1 Setup

The oligopoly model is as described in Section 2, with $\lambda = 1$ (Bertrand-Nash conduct). In the following, we focus on a sector z and drop the sector index to ease notation. We now put more structure on the distribution of cost and quality shocks, as well as 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 a normal distribution with mean zero and variance v^2 and θ^2 , respectively. Moreover, the iceberg-type trade cost τ_{in} is set equal to 1 if firm i is based in country n and otherwise to $\tau_{on} \equiv \mathcal{T} \times (\text{dist}(o, n))^\beta$, where $o \neq n$ denotes the country in which firm i is located, $\text{dist}(o, n)$ is the distance between countries o and n , and $\mathcal{T}$ 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 .²⁰

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_{od}^o and ϕ_{od}^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_{od}^o is observable to the econometrician whereas ϕ_{od}^u is not, so that ϕ_{od}^o can be used as an excluded first-stage variable when applying the HMR methodology. 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 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, $(c_{in}\tau_{on})^{\frac{1-\sigma}{1+\sigma\gamma}}/f_{on}$ (for every origin country o and firm $i \in \mathcal{J}_o$), 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.

6.2 Calibration

We choose parameter values to generate a Monte Carlo data-set broadly similar to the firm-level data-set used in Section 5. We use the same set of countries as in the empirical implementation and take the bilateral distance matrix $\text{dist}(o, d)$ directly from the data. Market size in country n , $\alpha_n E_n$ in Section 2, is set equal to (something proportional to) country- n

²⁰Thus, using the notation of Section 2, we are setting

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

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 the assumption of 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.

GDP in the data. We assume that, for every country o , $|\mathcal{J}_o|$, the number of firms based in o , is proportional to that country's GDP, with the proportionality coefficient chosen so that the total number of firms is 220. 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 $\beta = 0.35$, which is our baseline empirical estimate of the distance coefficient (see Table 4).

We still require values for the following four parameters: F , the intercept of the fixed-cost function; $\mathcal{T}$, the intercept of the trade-cost function; ν , 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 Chinese and French firm-level data):

1. The fraction of zeros in all potential (firm-destination-product-year) export relationships (0.92);
2. the mean (by origin-destination-product-year) aggregate combined market share of Chinese and French firms (0.139);
3. the median (by origin-product-year) 90/10 ratio of firm-level total exports (451);
4. the median (by origin-destination-product-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 moment 1, $\mathcal{T}$ to have a strong and negative effect on moment 2, ν to have a strong and positive effect on moment 3, and θ to have a strong and positive effect on moment 4. In practice, we adjust the vector of parameters $(F, \mathcal{T}, \nu, \theta)$ to minimize the sum of the squared Davis-Haltiwanger deviations between theoretical and empirical moments.²¹

We approximate the theoretical moments using Monte Carlo integration. For each parameter vector, we perform 10 Monte Carlo runs. For each run, we randomly draw vectors and matrices of firm baseline costs (ε_i^c) , firm-destination cost shocks (ε_{in}^c) , and fixed-cost shocks (ϕ_{od}^o) and (ϕ_{od}^u) . For each destination within a run, we then compute the equilibrium

²¹The 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.

of the entry game using the behavioral assumption mentioned in the previous subsection, and the oligopoly equilibrium using a variant of Nocke and Schutz (2018b)’s nested fixed-point algorithm. Having done that for all ten runs, we compute arithmetic averages (or medians) across runs to obtain Monte Carlo approximations to our theoretical moments.

Our calibration algorithm converges to $F = 1.34 \times 10^{-9}$ (times total world expenditures in the sector, which we normalized to unity), $\mathcal{T} = 0.144$, $v = 0.254$, and $\theta = 1.13$. At that parameter vector, we obtain nearly perfect matches for moments 3 and 4 (456 and 221, respectively, vs. 451 and 220 in the data), and we slightly under-predict the fraction of zeros in the firm-level export matrix (80.4% vs. 92% in the data) and the combined market share of Chinese and French firms (12.3% vs 13.9% in the data).

6.3 Data Generation and Results

Now that the parameters have been calibrated, we can 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 baseline costs (ε_i^c), firm-destination cost shocks (ε_{in}^c), and fixed-cost shocks (ϕ_{od}^o) and (ϕ_{od}^u), and can thus be thought of as a different sector 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, as well bilateral distance and observable fixed-cost shocks. To make the data-set comparable to the one used in our empirical application, we only keep observations for firms based in the countries corresponding to China and France in the Monte Carlo experiment. We thus obtain a firm-level data-set, which we also aggregate up to construct an industry-level data-set. We can then run our firm- and industry-level regressions on those Monte Carlo data-sets to evaluate the performance of our oligopoly correction terms.

The results of firm-level regressions can be found in Table 14. A first observation is that all OLS estimates are strongly biased towards zero, consistent with the heteroskedasticity bias discussed in Section 2.1 and with the empirical results in Table 4. Focusing now on PPML estimates, we see that specifications that do not use our correction term also tend to underestimate the absolute value of the distance coefficient due to the omitted variable bias. The PPML specifications with the oligopoly correction term deliver estimates that are very close to the true distance coefficient (-1.4).

The results of industry-level regressions are in Table 15. In all specifications, our oligopoly correction term improves the accuracy of the distance-coefficient estimate. Specifications

Table 14: Monte Carlo: Firm-Level Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	OLS	PPML	PPML	OLS	OLS	PPML	PPML
VARIABLES	all	all	all	all	top3 exp	top3 exp	top3 exp	top3 exp
	no corr	corr	no corr	corr	no corr	corr	no corr	corr
ldist	-0.576*** (0.0777)	-0.588*** (0.0781)	-0.966*** (0.0801)	-1.446*** (0.291)	-0.624** (0.261)	-0.645** (0.276)	-1.064*** (0.146)	-1.632*** (0.451)
Observations	116,483	116,483	427,505	427,505	8,981	8,981	21,405	21,405
R-squared	0.326	0.322			0.455	0.446		
Firm-year FE	YES	YES	YES	YES	YES	YES	YES	YES
Destination-year FE	YES	YES	YES	YES	YES	YES	YES	YES

Note: True distance coefficient is -1.4 .

that explicitly account for selection (columns Heckman and HMR), when combined with our correction term, deliver estimates that are very close to the true distance coefficient (-1.4).

Table 15: Monte Carlo: Industry-Level Results

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	OLS	OLS	PPML	PPML	Heckman	Heckman	HMR	HMR
VARIABLES	all	all	all	all	all	all	all	all
	no corr	corr	no corr	corr	no corr	corr	no corr	corr
ldist	-1.165*** (0.0632)	-1.300*** (0.0720)	-0.966*** (0.0801)	-1.192*** (0.123)	-1.202*** (0.0748)	-1.341*** (0.0864)	-1.233*** (0.0754)	-1.378*** (0.0852)
$\log \hat{Z}$							-15.96 (25.41)	-23.24 (29.61)
$(\log \hat{Z})^2$							11.21 (15.27)	15.64 (18.00)
$(\log \hat{Z})^3$							-2.890 (3.175)	-3.869 (3.785)
inv. Mills					0.770 (0.758)	0.845 (0.878)	-1.006 (4.095)	-2.068 (4.636)
Observations	7,628	7,628	8,432	8,432	7,628	7,628	7,628	7,628
R-squared	0.703	0.668			0.703	0.668	0.703	0.668
Firm-year FE	YES	YES	YES	YES	YES	YES	YES	YES
Destination-year FE	YES	YES	YES	YES	YES	YES	YES	YES

Note: True distance coefficient is -1.4 .

7 Conclusions

In this paper, we have evaluated the consequences of oligopolistic behavior for the estimation of gravity equations for trade flows. We showed that with oligopolistic competition, firm-level gravity equations based on a standard CES demand framework need to be augmented by markup terms that are functions of firms' market shares. At the aggregate level, the additional term takes the form of the exporting country's market share in the destination country

multiplied by an exporter-destination-specific Herfindahl-Hirschman index. We showed how to construct appropriate correction terms for both cases that can be used to avoid problems of omitted variable bias. Using combined French and Chinese firm-level export data as well as a sample of product-level imports by European countries, we showed that correcting for oligopolistic behavior can lead to substantial changes in the coefficients on standard gravity regressors.

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 as well as its limits; (b) show that the equilibrium condition (10) has a unique solution; (c) show that, at $\lambda = 1$, the first-order conditions of profit maximization are sufficient for global optimality, so that the profile of quantities $(q_j^*(1))_{j \in \mathcal{J}}$ does constitute a Nash equilibrium of the Cournot game. We do so below.

(a) As $1 + \sigma\gamma > 0$, the right-hand of equation (9) is strictly increasing in s_i , whereas the left-hand side is non-increasing in s_i . It follows that equation (9) 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 (10) is strictly decreasing in H , and has limits 0 and $|\mathcal{J}|/\lambda$ as H tends to ∞ and 0, respectively. It follows that equation (10) has a unique solution, $H^*(\lambda)$.

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

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

where we have dropped the destination subscript for ease of notation. 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 . $\square$

A.2 Proof of Proposition 2

Proof. To apply Taylor's theorem, we require the value of $s_e^{*'}(0)$. This requires computing the partial derivatives of $S(\cdot, \cdot)$ at $\lambda = 0$ as well as $H^{*'}(0)$. Differentiating equation (9) with respect to s_i , λ , and $t_i \equiv T_i/H$ at $\lambda = 0$ yields

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

It follows that²²

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

Next, we differentiate equation (10) with respect to λ and H :

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

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

$$\sum_{j \in \mathcal{J}} \left[-\frac{\sigma(1 + \gamma)}{1 + \sigma\gamma} s_j^{*'}(0) \frac{dH}{H} - \frac{\sigma - 1}{1 + \sigma\gamma} (s_j^{*'}(0))^2 d\lambda \right] = 0.$$

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

$$\frac{H^{*'}(0)}{H^*(0)} = -\frac{\sigma - 1}{\sigma(1 + \gamma)} \text{HHI}(0).$$

²²Notation: $\partial_k S$ is the partial derivative of S with respect to its k th argument.

We can now compute $s_i^{*'}(0)$:

$$\begin{aligned}
s_i^{*'}(0) &= \frac{\partial}{\partial \lambda} S \left(\frac{T_i}{H^*(\lambda)}, \lambda \right) \Big|_{\lambda=0} \\
&= -\frac{T_i}{H^*(0)} \partial_1 S \left(\frac{T_i}{H^*(0)}, 0 \right) \frac{H^{*'}(0)}{H^*(0)} + \partial_2 S \left(\frac{T_i}{H^*(0)}, 0 \right) \\
&= \frac{\sigma - 1}{1 + \sigma\gamma} [s_i^*(0) \text{HHI}(0) - (s_i^*(0))^2].
\end{aligned}$$

It follows that

$$\begin{aligned}
\frac{s_e^{*'}(0)}{s_e^*(0)} &= \frac{\sigma - 1}{1 + \sigma\gamma} \frac{1}{s_e^*(0)} \sum_{j \in \mathcal{E}} [s_j^*(0) \text{HHI}(0) - (s_j^*(0))^2] \\
&= \frac{\sigma - 1}{1 + \sigma\gamma} \left[\text{HHI}(0) - s_e^*(0) \sum_{j \in \mathcal{E}} \left(\frac{s_j^*(0)}{s_e^*(0)} \right)^2 \right] \\
&= \frac{\sigma - 1}{1 + \sigma\gamma} [\text{HHI}(0) - s_e^*(0) \text{HHI}_e(0)].
\end{aligned}$$

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

$$\begin{aligned}
\log s_e^*(\lambda) &= \log s_e^*(0) + \frac{d}{d\lambda} \log s_e^*(\lambda) \Big|_{\lambda=0} \lambda + o(\lambda) \\
&= \log s_e^*(0) + \frac{\sigma - 1}{1 + \sigma\gamma} [\text{HHI}(0) - s_e^*(0) \text{HHI}_e(0)] \lambda + o(\lambda) \\
&= \log s_e^*(0) + \frac{\sigma - 1}{1 + \sigma\gamma} [\text{HHI}(\lambda) - s_e^*(\lambda) \text{HHI}_e(\lambda)] \lambda + o(\lambda),
\end{aligned}$$

where the last line follows from the fact that $\text{HHI}(\lambda) - \text{HHI}(0)$ and $s_e^*(\lambda) \text{HHI}_e(\lambda) - s_e^*(0) \text{HHI}_e(0)$ are at most first order. $\square$

B Price Competition

B.1 Theoretical Results

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

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

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

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

$$\begin{aligned} 0 = \frac{\partial \pi_{in}}{\partial p_{in}} &= a_{in} p_{in}^{-\sigma} P_n^{\sigma-1} \alpha_n E_n + (p_{in} - C'_{in}(q_{in})) \left[-\frac{\sigma}{p_{in}} + \frac{\sigma-1}{P_n} \lambda \frac{\partial P_n}{\partial p_{in}} \right] \alpha_n E_n a_{in} p_{in}^{-\sigma} P_n^{\sigma-1} \\ &= q_{in} \left(1 - \frac{p_{in} - C'_{in}(q_{in})}{p_{in}} [\sigma - \lambda(\sigma-1)s_{in}] \right), \end{aligned} \quad (15)$$

where

$$s_{in} \equiv \frac{a_{in} p_{in}^{1-\sigma}}{\sum_{j \in \mathcal{J}} a_{jn} p_{jn}^{1-\sigma}} \quad (16)$$

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

Equation (15) pins down firm i 's optimal markup under price competition:

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

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 sector-level results. As in Section 4, 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

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

Plugging these definitions into equation (15), making use of equation (16), and rearranging, we obtain:

$$\left(1 - s_i^{\frac{1+\sigma\gamma}{\sigma-1}} \left(\frac{H}{T_i} \right)^{\frac{1+\gamma}{\sigma-1}} \right) (\sigma - \lambda(\sigma-1)s_i) = 1. \quad (17)$$

Note that the left-hand side of equation (17) 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 (17) 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:

$$\sum_{j \in \mathcal{J}} S\left(\frac{T_j}{H}, \lambda\right) = 1. \quad (18)$$

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 (18). 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 (17). From equation (16), firm i 's equilibrium price is given by*

$$p_i^*(\lambda) = \left(\frac{s_i^*(\lambda) H^*(\lambda)}{a_i} \right)^{\frac{1}{1-\sigma}}.$$

Proof. All that is left to do is check that first-order conditions are sufficient for optimality when $\lambda = 1$. Combining equations (15) and (17) yields:

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

where

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

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 $\tilde{p}_i > 0$ such that $\chi(p_i) > 0$ if $p_i > \tilde{p}_i$ and $\chi(p_i) < 0$ if $p_i < \tilde{p}_i$. Hence, π_i is strictly increasing on the interval $(0, \tilde{p}_i)$, and firm i 's first-order condition holds nowhere on that interval. The fact that $\lim_{p_i \rightarrow \infty} \chi(p_i) = 1$ and $\lim_{p_i \rightarrow \infty} \phi(p_i) = \sigma$ and the

monotonicity properties of χ and ϕ on $(\tilde{p}_i, \infty)$ imply the existence of a unique $\hat{p}_i$ at which firm i 's first-order condition holds. Moreover, π_i is strictly increasing on $(\tilde{p}_i, \hat{p}_i)$ and strictly decreasing on $(\hat{p}_i, \infty)$. First-order conditions are therefore sufficient for optimality. $\square$

Having characterized the equilibrium in a given destination, we now adapt the first-order approach to sector-level gravity to the case of price competition. As in Section 4, 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 4.

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*

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

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 (17) as

$$s_i = t_i^{\frac{1+\gamma}{1+\sigma\gamma}} \left(1 - \frac{1}{\sigma - \lambda(\sigma - 1)s_i} \right)^{\frac{\sigma-1}{1+\sigma\gamma}}. \quad (19)$$

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 $H^{*'}(0)$, we differentiate equation (18):

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

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:

$$\frac{H^{*'}(0)}{H^*(0)} = -\frac{\sigma - 1}{\sigma(1 + \gamma)} \text{HHI}(0).$$

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

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

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

$$s_e^{*'}(0) = \frac{\sigma - 1}{\sigma(1 + \sigma\gamma)} [\text{HHI}(0) - s_e^*(0) \text{HHI}_e(0)].$$

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

Proposition B motivates the following approximation:

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

As in Section 4, this approximation can then be used to derive the sector-level gravity regression

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

where

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

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 .

B.2 Empirical Results

Table 16 presents results for our estimates of σ and γ using the estimation procedure from Section 3 but replacing the Cournot markup formula with its Bertrand equivalent. This only leads to minor changes in coefficient estimates.

Table 16: 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
HS 2-digit products	78	78

Note: 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 products.

Tables 17-19 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 20-22 show results for the pooled sector-level regressions. Again, the distance coefficient becomes larger in absolute magnitude when including the markup 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 17: Firm-level Gravity Estimates – Bertrand competition, $\sigma = 5$, $\gamma = 0$.

Regressor	PPML w/ corr	PPML w/o corr	OLS w/ corr	OLS w/o corr
log distance	-1.418*** (0.190)	-0.874*** (0.021)	-0.246*** (0.014)	-0.232*** (0.014)
$\hat{\beta}_{distance}$	0.355	0.219		-
Observations	11,955,786	11,955,786	708,386	708,386
R-squared			0.05	0.06
Firm-year FE	YES	YES	YES	YES
Product-dest.-year FE	YES	YES	YES	YES

Note: Firm-level data, pooled across sectors. Results for top 3 exporters. Bertrand model with $\sigma = 5$ and $\gamma = 0$. Standard errors in brackets, clustered at the destination-year level.

Table 18: Firm-level Gravity Estimates – Bertrand competition, $\sigma = 4.96$, $\gamma = 0$.

Regressor	PPML w/ corr	PPML w/o corr	OLS w/ corr	OLS w/o corr
log distance	-1.415*** (0.189)	-0.874*** (0.0210)	-0.246*** (0.013)	-0.232*** (0.013)
$\hat{\beta}_{distance}$	0.357	0.221		
Observations	11,955,786	11,955,786	708,392	708,386
R-squared			0.06	0.06
Firm-year FE	YES	YES	YES	YES
Product-dest.-year FE	YES	YES	YES	YES

Note: Firm-level data, pooled across sectors. Results for top 3 exporters. Bertrand model with mean of estimated σ and $\gamma = 0$. Standard errors in brackets, clustered at the destination-year level.

Table 19: Firm-level Gravity Estimates – Bertrand competition, $\sigma = 4.96$, $\gamma = 0.31$.

Regressor	PPML w/ corr	PPML w/o corr	OLS w/ corr	OLS w/o corr
log distance	-0.957*** (0.040)	-0.874*** (0.021)	-0.237*** (0.014)	-0.232*** (0.014)
$\hat{\beta}_{distance}$	0.613	0.560		
Observations	11,955,786	11,955,786	708,392	708,386
R-squared			0.06	0.06
Firm-year FE	YES	YES	YES	YES
Product-dest.-year FE	YES	YES	YES	YES

Note: Firm-level data, pooled across sectors. Results for top 3 exporters. Bertrand model with mean of estimated σ and γ . Standard errors in brackets, clustered at the destination-year level.

Table 20: Sector-level Gravity Estimates – Bertrand competition, $\sigma = 5$, $\gamma = 0$

Regressor	Heck w/o	Heck w/	HMR ² w/o	HMR ² w/	HMR ³ w/o	HMR ³ w/
log distance	-1.189*** (0.198)	-1.217*** (0.202)	-1.151*** (0.190)	-1.177*** (0.194)	-1.150*** (0.193)	-1.177*** (0.197)
inv mills	-0.121 (0.165)	-0.123 (0.168)	0.678*** (0.169)	0.702*** (0.174)	0.639** (0.309)	0.672** (0.315)
log $\hat{Z}$			0.907*** (0.265)	0.939*** (0.270)	0.736 (1.305)	0.809 (1.322)
log $\hat{Z}^2$			-0.103* (0.0565)	-0.108* (0.0575)	-0.0297 (0.534)	-0.0513 (0.540)
log $\hat{Z}^3$					-0.0102 (0.0693)	-0.00780 (0.0700)
$\hat{\beta}_{distance}$						
Observations	60,662	60,662	60,662	60,662	60,662	60,662
R-squared	0.302	0.297	0.304	0.299	0.304	0.299
Origin-product FE	YES	YES	YES	YES	YES	YES
Destination-product FE	YES	YES	YES	YES	YES	YES

Note: Sector-level data. Bertrand model with $\sigma = 5$ and $\gamma = 0$. Standard errors clustered at destination level in parentheses.

Table 21: Sector-level Gravity Estimates – Bertrand Competition, $\sigma = 4.96$, $\gamma = 0$

Regressor	Heck w/o	Heck w/	HMR ² w/o	HMR ² w/	HMR ³ w/o	HMR ³ w/
log distance	-1.189*** (0.198)	-1.217*** (0.202)	-1.151*** (0.190)	-1.177*** (0.194)	-1.150*** (0.193)	-1.177*** (0.197)
inv mills	-0.121 (0.165)	-0.123 (0.168)	0.678*** (0.169)	0.702*** (0.173)	0.639** (0.309)	0.672** (0.315)
log $\hat{Z}$			0.907*** (0.265)	0.939*** (0.270)	0.736 (1.305)	0.808 (1.322)
log $\hat{Z}^2$			-0.103* (0.0565)	-0.108* (0.0575)	-0.0297 (0.534)	-0.0512 (0.540)
log $\hat{Z}^3$					-0.0102 (0.0693)	-0.00780 (0.0700)
$\hat{\beta}_{distance}$						
Observations	60,662	60,662	60,662	60,662	60,662	60,662
R-squared	0.302	0.297	0.304	0.299	0.304	0.299
Origin-product FE	YES	YES	YES	YES	YES	YES
Destination-product FE	YES	YES	YES	YES	YES	YES

Note: Sector-level data. Bertrand model with mean of estimated σ and $\gamma = 0$. Standard errors clustered at destination level in parentheses.

Table 22: Sector-level Gravity Estimates – Bertrand competition, $\sigma = 4.96$, $\gamma = 0.31$

Regressor	Heck w/o	Heck w/	HMR ² w/o	HMR ² w/	HMR ³ w/o	HMR ³ w/
log distance	-1.189*** (0.198)	-1.200*** (0.199)	-1.151*** (0.190)	-1.161*** (0.192)	-1.150*** (0.193)	-1.161*** (0.194)
inv mills	-0.121 (0.165)	-0.122 (0.166)	0.678*** (0.169)	0.687*** (0.171)	0.639** (0.309)	0.652** (0.311)
log $\hat{Z}$			0.907*** (0.265)	0.919*** (0.267)	0.736 (1.305)	0.765 (1.311)
log $\hat{Z}^2$			-0.103* (0.0565)	-0.105* (0.0569)	-0.0297 (0.534)	-0.0382 (0.536)
log $\hat{Z}^3$					-0.0102 (0.0693)	-0.00925 (0.0696)
$\hat{\beta}_{distance}$						
Observations	60,662	60,662	60,662	60,662	60,662	60,662
R-squared	0.302	0.300	0.304	0.302	0.304	0.302
Origin-product FE	YES	YES	YES	YES	YES	YES
Destination-product FE	YES	YES	YES	YES	YES	YES

Note: Sector-level data. Bertrand model with mean of of estimated σ and γ . Standard errors clustered at destination level in parentheses.

C Data Appendix

Table 23: List of export destinations included in the firm-level and product-level data

Austria	Latvia
Belgium	Lithuania
Bulgaria	Luxembourg
Croatia	Malta
Cyprus	Netherlands
Czech Rep.	Norway
Denmark	Poland
Estonia	Portugal
Finland	Romania
France	Serbia
Germany	Slovakia
Greece	Slovenia
Hungary	Spain
Iceland	Sweden
Ireland	Turkey
Italy	UK

References

- 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, K., A. COSTINOT, AND A. RODRIGUEZ-CLARE (2012): “Competition, New Trade Models, Same Old Gains?,” *American Economic Review*, 105(10), 3183–3221.
- ATKESON, A., AND A. BURSTEIN (2008): “Pricing-to-Market, Trade Costs, and International Relative Prices,” *American Economic Review*, 98(5), 1998–2031.
- 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.
- 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.
- 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. (1994): “New Product Varieties and the Measurement of International Prices,” *American Economic Review*, 84(1), 157–177.
- GAUBERT, C., AND O. ITSKHOKI (forthcoming): “Granular Comparative Advantage,” *Journal of Political Economy*.
- HEAD, K., AND T. MAYER (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.
- 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.
- REDDING, S., AND A. J. VENABLES (2004): “Economic geography and international inequality,” *Journal of International Economics*, 62, 53–83.
- ROSE, A. K. (2000): “One Money, One Market: Estimating the Effect of Common Currencies on Trade,” *Economic Policy: A European Forum*, 30, 7–33.
- 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.