PPML, Gravity, and Heterogeneous Effects*

Vladimir Tyazhelnikov†
The University of Sydney

Xinbei Zhou‡
University of Oxford

February 2020

Abstract

The gravity equation is the most popular empirical tool among trade economists. Two of the most common approaches to estimating it are the ordinary least squares (OLS) and Poisson Pseudo-Maximum Likelihood (PPML), with PPML often being preferred to OLS because it does not lead to a bias if the error term of the regression is heteroskedastic. We show theoretically, and document in a series of Monte-Carlo simulations that when the trade elasticity is not constant between country pairs, OLS and PPML estimates of the gravity equation have different interpretations: OLS estimates are the average elasticity and PPML are the elasticity of the average. Furthermore, we employ international trade data and show that more than 100% of the difference between PPML and OLS estimates of distance elasticity is explained by the difference in the interpretation of the coefficients. The bias of OLS estimates associated with the error term heteroskedasticity accounts for 8% of the difference between the estimates and has the sign opposite to what was previously found in the literature.

JEL Classification: F10, F14, C13, C21, C50

Keywords: elasticity, hetoregeneity, heteroskedasticity, gravity model, misspecification

*We thank Scott French, Ye Lu, Rachael Meager, Mark Melatos, and John Romalis for early suggestions.
†vladimir.tyazhelnikov@sydney.edu.au
‡xinbei.zhou@economics.ox.ac.uk. The majority of the project was carried out in Sydney.

1 Introduction

In this paper, we focus on the behaviour of techniques commonly used for the estimation of constant elasticity models in the presence of unobserved heterogeneity. Wooldridge (2005) shows that in the case heterogeneity is independent of the covariates, the Ordinary Least Squares (OLS) estimator, which does not account for heterogeneity, generates estimates that can be interpreted as an Average Partial Effect (APE). It follows then that the coefficient from a log-log OLS regression in case of heterogeneous effects can be interpreted as an average elasticity.

While average elasticity is a meaningful object, depending on the question, researchers might be interested in the elasticity of the average,[1] the effect of one percent change in one variable on percentage change of average value of another variable. In this paper, we use the gravity equation, one of the most widely used empirical tools in the field of international trade, to illustrate this difference. We show that the estimates obtained from Poisson Pseudo Maximum Likelihood (PPML), the most common econometric technique used to estimate the gravity equation, can be interpreted as the elasticity of the average.

This finding has two implications: first, PPML deserves wider recognition outside of the field of international trade. Second, all the difference between PPML and OLS coefficients in the gravity equation has been wrongly attributed to the fact that unlike PPML, OLS estimates are biased when there is heteroskedasticity in an error term, a point that Santos Silva and Tenreyro (2006) (further SST) raised.

In this paper, we obtain evidence of the presence of heterogeneous effects across country pairs from a gravity dataset. We document that there is heterogeneity in the effect of distance on trade volumes. Furthermore, we show, both theoretically and with Monte Carlo simulations, that the presence of these heterogeneous effects causes bias in the estimation of

[1] Consider the simple numerical example below. In a hypothetical country, there are only two individuals, one rich and one poor. The rich one has an income of \$1,000,000, while the poor one has only \$1,000. They face heterogeneous tax rates: 10% and 1%, for the rich one and poor one, respectively. The average of tax rates is calculated as 5.5%, but the total tax revenue is \$100, 000 + \$10 = \$100, 010. The percentage of total tax revenue in terms of total income is $\frac{100{,}010}{1{,}001{,}000} = 9.991\%$, which can be interpreted as the "tax rate on average". The latter is the true effect of tax rates on the whole economy. The reason why the simple average of tax rates is biased from the true effect is that it neglects the different weights the two individuals have in the economy and includes both of them into the computation with equal weights. To accurately calculate the tax rate on average, the heterogeneous tax rates for the two individuals should be assigned with weights according to the individuals' share of income of total income of the whole country: $10\% * \frac{1000000}{1001000} + 1\% * \frac{1000}{1001000} = 9.991\%$. Similarly, countries of different sizes would have various degrees of impact on global trade volumes and should enter into the gravity equation computation with different weights. Otherwise, the use of OLS estimator, the average elasticity, for the effects of a uniform decrease in trade costs, for example, the effect of a worldwide decrease in the gasoline price, on global trade volumes, leads to biased results.

the elasticity of the average by the OLS.[2] We introduce a Weighted Least Squares (WLS) estimator and show that both PPML and WLS estimators are robust to the presence of heterogeneous effects; at the same time, unlike PPML and similarly to OLS, WLS is not robust to heteroskedasticity of an error term.

Comparison of OLS, PPML, and WLS estimates allows us to decompose the bias caused by using the OLS estimator with log-linearized gravity equation into two different sources: heteroskedasticity in the error term and heterogeneity across country pairs, using PPML estimate as a benchmark. We document that the share of bias in estimating the distance elasticity caused by heterogeneity is much larger than that caused by heteroskedasticity; in particular the share of bias caused by heteroskedasticity accounts for approximately 8% of the total difference between OLS and PPML estimates and has the opposite sign to the heterogeneity bias.

While the topic of heterogeneous country-pair trade elasticity is rarely discussed, there exist a few empirical studies examining the distance elasticity across countries. Fratianni and Kang (2005) first use the log form of gravity equation to estimate the distance elasticity of the full sample, which yields a significant estimate of -1.17. Then they test whether the homogeneity assumption for different groups of countries holds with two tests: one for OECD countries and non-OECD countries and the other for Christian and Islamic countries; both tests reject the null hypothesis of distance homogeneity at the 1% level. They find that the distance elasticity is much smaller for OECD members in absolute values than for non-members. Magerman et al. (2015) summarize multiple previous studies of the distance and border effects in international trade. These studies perform sensitivity tests of the effects for various countries, regions, and periods with different methods; Magerman et al. (2015) documents the presence of heterogeneous effects in distance and trade elasticity across different country pairs.

Discussion of mechanisms behind the heterogeneity of trade and distance elasticity are beyond the scope of this paper, but Fieler (2011), Novy (2013), and Bas et al (2017) provide a micro-foundation of the gravity equation with country-pair heterogeneity. Fieler (2011) first confirms that low income countries trade less than rich countries, both with each other and with the rest of the world: in 2000, transactions to and from the 12 Western European countries accounted for 45% of global trade while the 57 African countries accounted for only 4.2%. This fact provides evidence that large countries are overrepresented in global trade flows and, therefore, should be given more weight in calculating the elasticity of the average. The author relaxes two of the assumptions of trade models that generate the gravity equation:

[2]In this paper, we focus on the estimation of the elasticity of the average, from this point of view OLS estimates are biased. Alternatively, we can say that OLS and PPML estimates have different interpretation.

homothetic preferences and the production in low and high income countries differs only in quantitative, not qualitative, aspect. With this new model, the author is able to explain the prevalence of large trade flows among the rich countries and small trade flows among poor countries. Novy (2013) shows that the gravity equation based on translog demand system, unlike standard CES gravity model, generates trade elasticity heterogeneous among country-pairs. Bas et al. (2017) make related point but focus on the supply-side mechanism of the heterogeneity. They introduce an extension of the Melitz (2003) model with the productivity distribution of heterogeneous firms following log-normal instead of commonly used Pareto distribution. They show that this assumption leads to bilateral-specific aggregate trade elasticity for each country-pair.

Finally, our paper fits into a bigger strand of literature devoted to the estimation of the gravity equation (see Head and Mayer (2013) and Yotov et al. (2016) for a detailed review of this literature), and in particular the research on aggregation issues of trade data. Costinot and Rodríguez-Claire (2014), Kehoe et al. (2017), and French (2019) address issues caused by industrial heterogeneous effects. Coughlin and Novy (2019) analyze the consequences of spatial aggregation of trade data. Larch et al. (2019) show that OLS and PPML estimates diverge when there is a large number of small countries in the data.

The rest of the paper is organized as follows: we discuss the methodology and theoretical background in Section 2. In Section 3, we present the results of Monte Carlo simulations. In Section 4, we employ the gravity equation to prove the existence of country-pair heterogeneous effects and decompose the difference between OLS and PPML estimates. In Section 5, we conclude.

2 Methodology

2.1 Heterogeneity

In this paper, we focus on the case of unobserved heterogeneity, which is independent of other covariates. The reason is that if the variation in observed independent variables drives the heterogeneity, this issue can be addressed by the inclusion of the interaction term between the variable of interest and the variable that causes heterogeneity. If the unobserved variable is correlated with other covariates, a suitable proxy variable can address this issue, while the rest of the procedure would remain the same.

We rely on the approach by Wooldridge (2005) to handling the unobserved heterogeneity and incorporate it into a constant elasticity model. A general formulation is

$$y = \beta_0 x^{\beta} x^{q\gamma} v,$$

where y and x are dependent and control variables correspondingly, q is unobserved variable, and v is an error term that satisfies unit conditional mean assumption: $E(v|\mathbf{x}, q) = 1$.

The log-linearized version of this expression is then:

$$\log y = \beta_0 + \beta \log x + \gamma q \log x + \log v$$

The partial effect of $\log x$ on $E(\log y|\mathbf{x}, q)$ is then

$$\frac{\partial E(\log y| \log x, q)}{\partial \log x} = (\beta + \gamma q).$$

As the unobserved term q differs from observation to observation, this partial effect is non-constant, and we interpret it as the observation-specific elasticity. The formulation above can be interpreted as a random coefficient model with the observation-specific elasticity $\theta_i \equiv \beta + \gamma q$; to simplify the notation from now on, we will rely on this interpretation of the unobserved heterogeneity model.

Wooldridge (2005) shows that in case q is independent of x, $E(q) = \mu$, OLS estimates the Average Partial Effect (APE) of $\log x$ on $E(\log y|\mathbf{x}, q)$, which is equal to $\beta + \gamma\mu$. This estimate can be interpreted as the average elasticity as it simply averages out partial effects. Because of Jensen's inequality, however, average elasticity and elasticity of the average are not equal. If there are N observations with values of y_i and corresponding elasticities θ_i, the average elasticity is simply $\frac{\sum_{i=1}^{N} \theta_i}{N}$ and the elasticity of the average can be computed as a weighted average of elasticities with weights equal to the share of y_i in $\sum_{i=1}^{N} y_i$: $\frac{\sum_{i=1}^{N} \theta_i y_i}{\sum_{i=1}^{N} y_i}$. It follows that to calculate the elasticity of the average, we do not need to know the values of unobserved variable q, the weights are proportional to the observed values of y, and thus, an estimator that applies these weights can be interpreted as the elasticity of the average.

2.2 The Gravity Equation and Heteroskedastic Errors

In this section, we use the "naive form" of gravity equation in Head and Mayer (2014) to be consistent with the Constant-Elasticity Models in SST and restrain the trade elasticity to be distance elasticity:

$$Y_{ij} = \frac{GDP_i^{\beta_1} * GDP_j^{\beta_2}}{Distance_{ij}^{\theta}} \varepsilon_{ij} \tag{1}$$

where Y_{ij} is the trade flow between country i and j, $Distance_{ij}$ stands for bilateral distance and ε_{ij} is the error term.

Taking logs on both sides of the equation yields:

$$\log Y_{ij} = \beta_1 \log gdp_i + \beta_2 \log gdp_j + \theta \log distance_{ij} + \log \varepsilon_{ij} \tag{2}$$

where $\theta < 0$ is the trade elasticity.[3]

To obtain consitent estimates of the coefficients in Equation 1 using the log-linearized form Equation 2 by OLS, it is necessary for $\mathbb{E}[\log \varepsilon_{ij}|\mathbf{X}]$ to be constant. However, according to SST, ε_{ij} is generally heteroskedastic, taking simple OLS regression of $\log Y_{ij}$ on $\mathbf{X}$ leads to inconsistent estimates of θ. Moreover, when there is heteroskedasticity in the error term, log transformation of gravity equation potentially leads to the violation of the exogeneity assumption and biases the OLS estimator. From Equation 2:

$$\hat{\log} \varepsilon_{ij} = \log Y_{ij} - \hat{\log} Y_{ij}$$

where $\hat{\log} Y_{ij} = \mathbf{X}\hat{\beta}$ is the predicted log of bilateral trade volumes. Even though the exogeneity assumption $Cov(\varepsilon_{ij}, \mathbf{X}) = 0$ is satisfied in Equation 2, with heteroskedasticity in ε_{ij}, $Cov(\log \varepsilon_{ij}, \mathbf{X})$ does not necessarily equal to 0. Thus, the OLS estimators is biased.

2.3 Country-pair Heterogeneity

To focus on estimating the aggregate elasticity that reflects the effect of a uniform change in the trade costs on bilateral trade volumes, i.e., the elasticity the of average, we incorporate heterogeneous country-pair trade elasticities into Equation 2:

$$\log Y_{ij} = \beta_1 \log gdp_i + \beta_2 \log gdp_j + \theta_{ij} \log distance_{ij} + \log \varepsilon_{ij} \tag{3}$$

In case the error term is homoskedastic, the OLS estimator $\hat{\theta}$ is a consistent estimator for APE: $\overline{\theta}^{APE} = E(\theta_i j)$ given that the assumption of exogeneity with heterogeneity is satisfied: $E(\theta_{ij} \mid \log distance_{ij}) = E(\theta_{ij})$. This estimator was defined previously as the average trade elasticity. However, using the OLS to estimate the elasticity of the average leads to a bias with heterogeneity because $\overline{\theta}^{APE}$ assigns equal weight to every individual θ_{ij}, while different trading country pairs have a differential impact on the world's aggregate trade flows.

A natural solution candidate to the weighting problem is the WLS estimator. The share

[3] The multiplicative error term ε_{ij} used in our model is derived from the "true" additive error term $\eta_{ij} = Y_{ij} - \hat{Y_{ij}}$ where $\varepsilon_{ij} = 1 + \frac{\eta_{ij}}{\exp(\mathbf{X}*\beta)}$ and $\mathbb{E}[\varepsilon_{ij}|\mathbf{X}] = 1$.

of the bilateral trade flows between country i and j over the global trade flows, denoted by $\frac{Y_{ij}}{\mathbf{Y}}$, can be used as the weights. The formal justification of this method follows the solutions to the endogenous stratified sampling issue proposed by Hausman and Wise (1981), where they use the Gary Income Maintenance Experiment as an example to show the extent of selection bias due to endogenous stratified sampling and demonstrate that the bias can be corrected by both maximum likelihood estimator (MLE) and WLS estimator.

In the context of this study, the trade elasticity can be interpreted as the effect of trade liberalization on the representative dollar of world trade flows. If every country-pair is treated as a stratum, the probability of a representative dollar falls in any stratum is $\frac{1}{M}$, where M is the number of country-pairs, and the share of the bilateral trade flows between a country-pair is $\frac{Y_{ij}}{Y}$. The dollars from a country-pair are underrepresented when $\frac{Y_{ij}}{\mathbf{Y}} > \frac{1}{M}$. The relative probability of a particular dollar falling in a stratum is given by $\frac{1}{M} / \frac{Y_{ij}}{\mathbf{Y}} = \frac{Y}{M * Y_{ij}}$. Therefore, with normalization, the weight used in this study is $\frac{Y_{ij}}{\mathbf{Y}}$.

The moment conditions for OLS and WLS are listed below:

$$\sum \mathbf{X}(\log Y_{ij} - \hat{\log} Y_{ij}) = 0$$

$$\sum \mathbf{X}(\log Y_{ij} - \hat{\log} Y_{ij}) Y_{ij} = 0$$

From these conditions one can see that OLS involves log deviations of Y_{ij} from its predicted value. Since percent deviations are approximately equal to log deviations, multiplying the log deviations by actual trade volumes leads to a result close to the level deviations, which is implied by the WLS first-order condition.

We also use the PPML estimator in this study following the method discussed by SST; the moment condition for PPML is:

$$\sum \mathbf{X}(Y_{ij} - \hat{Y_{ij}}) = 0$$

where $\hat{Y_{ij}}$ is the predicted bilateral trade volume. This expression indicates that PPML also involves the level deviations of Y_{ij} from its predicted value, and hence, similarly to WLS, it addresses the heterogeneity issue.

3 Simulations

In this section, we simulate bilateral trade flows under assumptions of two types of countries and heterogeneous trade elasticities across country pairs. The simulation results indicate that the use of OLS estimator under such conditions yields biased estimates for the parameter of interest (the elasticity of the average) and that the application of the PPML and WLS estimators corrects this bias.

3.1 Data Generating Process

We incorporate the heterogeneous effects into SST's specifications of heteroskedastic errors in the data generating process (DGP):

$$Y_{ij} = \exp(\beta_1 \log gdp_i + \beta_2 \log gdp_j + \theta_{ij} \log distance_{ij}) * \varepsilon_{ij} \tag{4}$$

where ε_{ij} is a log normal random variable in with mean 1 and variance σ^2.

In this simple economy, there are two types of countries, large and small, with $GDP_L = 100$ and $GDP_S = 10$, respectively. The total number of countries in the economy is $N = 100$. According to the findings in previous literature, larger countries tend to have a trade elasticity that is smaller in absolute value (Fratianni and Kang, 2006). Therefore, under heterogeneous effects, we assume that the trade elasticity between two large countries is $\theta_L = -0.5$ and the elasticity for the country pairs involving small countries is $\theta_S = -1$.[4] The distance between a country-pair is randomly generated from the uniform distribution $U(2,3)$, and the effects of GDP on bilateral trade flows are assumed to be 1, i.e., $\beta_1 = \beta_2 = 1$.

Following SST, we consider the following two specifications of σ^2 to assess the performance of different estimators under different patterns of heteroskedasticity.

Case 1: $\sigma^2 = \exp(\log Y_{ij})^{-2}$; $V[Y_{ij}|\mathbf{X}] = 1$.

Case 2: $\sigma^2 = 1$; $V[Y_{ij}|\mathbf{X}] = \exp(\mathbf{X}\beta)^2$.

The DGP starts with when all countries are small, i.e., $N_L = 0$, and increases the number of large countries in the economy by one after each iteration until $N_L = 100$. For a total of 101 cases, we run 100 simulations with random errors and distances every time. This process follows the approach of Head and Mayer (2014) that illustrates the robustness features of PML estimators.

3.2 Estimates

Before analyzing the main estimated results from the above specifications with both heterogeneity in country-pairs and heteroskedasticity in the error term, we present a summary table below (Table 1) for a benchmark estimation of both OLS and PPML estimators of the trade elasticity. The benchmark case follows the above Case 1 heteroskedastic error,

[4] We chose country sizes as a source of unobserved heterogeneity to make the interpretation of our Monte-Carlo simulations more intuitive; we are agnostic on the true source of heterogeneity, and country sizes can be replaced with any other country characteristic that has a potential impact on distance or trade elasticity.

but without country-pair heterogeneity, i.e., all the countries considered are small. In this specification, the multiplicative gravity equation with additive error term [5] is homoskedastic, but it becomes heteroskedastic after log-linearization. We also include the estimates with Case 2 type error term but without country-pair heterogeneity.

We express the bias of using different estimators as the percentage difference between the estimates and the elasticity of the average as follows:

$$bias_k = 1 - \frac{\hat{\theta}_k}{\delta_k} \tag{5}$$

where k is the number of large countries in the economy, $\hat{\theta}_k$ is the estimated result when there are k large countries in the economy and δ_k is the elasticity of the average.

Table 1: Benchmark Estimates and Biases

	Case 1		Case 2	
	Estimate	95% Confidence Interval	Estimate	95% Confidence Interval
OLS	-1.000191	(-1.000339; -1.000043)	-1.010592	(-1.024148; -0.997036)
	Bias		Bias	
	-0.010193	(-0.010342 ; -0.010044)	-0.020698	(-0.034389; -0.007006)
PPML	Estimate		Estimate	
	-1.000115	(-1.000259 ; -0.999971)	-1.021753	(-1.039463 ; -1.004043)
	Bias		Bias	
	-0.010115	(-0.010261 ; -0.009970)	-0.031970	(-0.049857 ; -0.014083)
WLS	Estimate		Estimate	
	-1.000033	(-1.000178 ; -0.999887)	-1.033151	(-1.061874 ; -1.004428)
	Bias		Bias	
	-0.010033	(-0.010180 ; -0.009887)	-0.043482	(-0.007492 ; -0.014472)

This table represents the case with no country-pair heterogeneity and all countries are small. For cases 1 and 2 heteroskedasticity follow the corresponding specifications in Section 3.1. Biases calculated using equation 5.

As shown in Table 1, PPML gives an estimate closer to -1, the universal trade elasticity, than the OLS estimator under Case 1 with smaller standard errors. The reverse is true under Case 2. These results are consistent with SST's findings under their DGP. On the other hand, when there is country-pair heterogeneity, the OLS estimate will be drastically different from the PPML estimate, as shown in Table 2, where half of the countries are large. We also include WLS estimates to show the similarity between the estimated results using WLS and PPML, which is consistent with the discussion in Section 3.

[5] $Y_{ij} = \exp(\log Y_{ij}) + \eta_{ij}$, where $\varepsilon_{ij} = 1 + \frac{\eta_{ij}}{\exp(\log Y_{ij}}$, and $\mathbb{E}[\varepsilon_{ij}|\mathbf{X}] = 1$.

Table 2: Example Estimates and Biases under Heterogeneity

	Case 1		Case 2	
	Estimate	95% Confidence Interval	Estimate	95% Confidence Interval
OLS	-0.877038	(-0.879166; -0.874910)	-0.887398	(-0.901115; -0.873680)
	Bias		Bias	
	-0.659387	(-0.664443 ; -0.655330)	-0.680140	(-0.706099; -0.654182)
	Estimate		Estimate	
PPML	-0.560260	(-0.560517 ; -0.560002)	-0.577062	(-0.609140 ; -0.544984)
	Bias		Bias	
	-0.060031	(-0.060524 ; -0.059536)	-0.092612	(-0.153373 ; -0.031851)
	Estimate		Estimate	
WLS	-0.559261	(-0.559526 ; -0.558997)	-0.585113	(-0.636850 ; -0.533377)
	Bias		Bias	
	-0.058142	(-0.058652 ; -0.057632)	-0.107890	(-0.205885 ; -0.009895)

This table represents the heterogeneous case with the number of large and small countries both equal to 50. For cases 1 and 2 heteroskedasticity follow the specifications in Section 3.1. Biases calculated using equation (5).

The following graphs display the estimates of OLS, WLS, and PPML estimators of θ in models with both heterogeneous country-pairs and the above mentioned heteroskedastic errors, where Case 1 is graphed in Figure 1a and Case 2 in Figure 1b. The WLS estimator is estimated with the share of Y_{ij} of total global trade volumes as weight. The PPML estimator is estimated using Y_{ij} as a dependent variable and log of GDPs and log of bilateral distance as independent variables. We also calculate the elasticity of the average as the percentage difference in the level trade volumes before and after an 1% increase in bilateral distance, denoted as δ.

It is clear from both of the graphs that the OLS estimator has distinct behavior from the other two estimators and also the elasticity of the average δ. In general, PPML estimator and WLS estimator outperform OLS, with both smaller deviations from δ and similar trend as the number of large countries k increases. The differences between the OLS estimates and the other two estimates are significant when there is more heterogeneity in the economy, i.e., the number of large and small countries are close to each other, while in case of little heterogeneity in the economy, the estimates are not drastically different. These two points indicate that in our simulations, country-pair heterogeneity generates larger share of bias than heteroskedastic errors, which confirms our theory that PPML and WLS are capable of correcting the heterogeneous effects better than OLS.

Figure 1c and Figure 1d represent biases for Cases 1 and 2 correspondingly; comparing the OLS estimates across Figure 1c and Figure 1d, we can see that while the patterns of

Figure 1: Estimates and Biases of OLS, WLS, and PPML

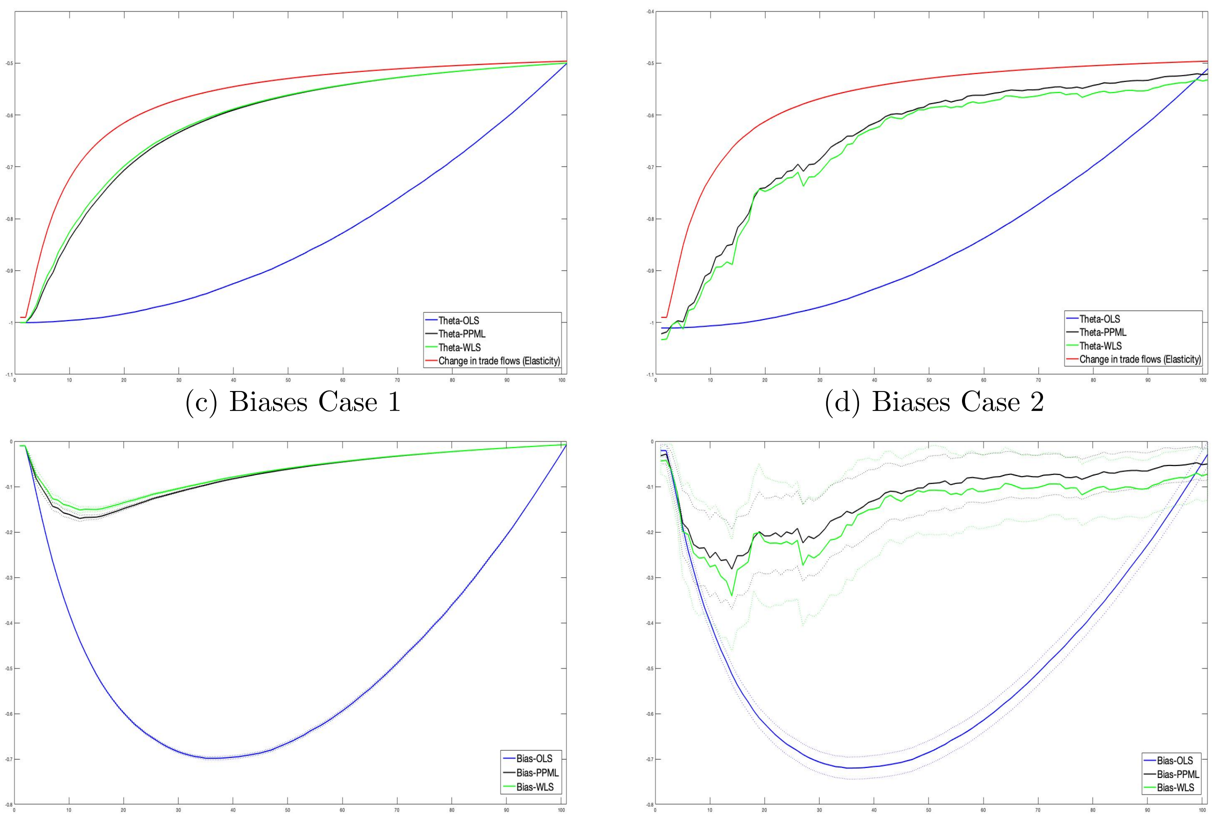

(a) Estimates Case 1

(b) Estimates Case 2

(c) Biases Case 1

(d) Biases Case 2

Case 1 and 2 heteroskedastic errors followed the same specification as in Section 3.1.
Trade elasticity calculated as the change in level trade volumes with an 1% increase in bilateral distance.
Biases calculated using Eq.(5).
95% Confidence Interval of estimated biases included as the dashed lines.

heteroskedasticity in the error term are different, the effect of this difference is trivial on both the shape of the graph and the scale of the bias. These findings suggest that the major determinant of the bias is not heteroskedasticity in the error terms as emphasized by SST but the heterogeneity in country-pairs. In Figure 1c, the heteroskedastic bias is negative as the WLS estimates are smaller, while it is positive with PPML estimates being larger in Figure 1d (notice that the PPML estimates have a narrower 95% confidence interval in this case).

The results in Table 1 and 2 are also present in these graphs. In Figure 1a, when the number of large countries $k = 0$ (no large countries) and $k = 100$ (no small countries), i.e. no heterogeneity, the PPML estimates are the closest to the elasticity of the average δ. On the other hand, OLS estimator is the best performer without heterogeneity in Case 2 (Figure 1b): the OLS estimated elasticity ($\hat{\theta_{OLS}}$) is the closest to δ.

PPML and WLS estimators generate very close estimates in both scenarios, and they both outperform the OLS estimator. Since both estimators are capable of correcting the heterogeneous effects across country pairs, the differences between the two is the result of heteroskedasticity in the error term discussed in Section 3. The sign and scale of the bias, however remain undetermined as they vary with different patterns of heteroskedasticity.

The dashed lines show the narrow 95% confidence interval of the biases under each case, indicating the differences across estimators are significant.

4 The Gravity Equation

In this section, we apply different estimation techniques to the gravity data. We document not only the existence of country-pair heterogeneity but also decompose the bias of OLS estimates by heterogeneity and heteroskedasticity channels.

To make our results comparable with ones in SST, we use the same data and replicate their results for the cases of OLS and PPML. Following SST's specification, we include distance and indicator variables for remoteness, common language, colonial heritage and preferential agreement in the model and replicate the Table 5 from SST with the Anderson-van Wincoop Gravity Equation controlling for multilateral resistance terms with importer and exporter fixed effects. Table 3 reports the estimates of the model from all three estimators and corresponding standard errors.

First, we present the evidence obtained from the gravity dataset that there are country-pair heterogeneous effects. Naturally, we cannot establish a link between the unobserved variable and distance elasticity; trade volumes, on the other hand, are observed. If the OLS estimates of the distance elasticity do not vary much depending on how large trade flows are, then there will not be a significant difference between the average elasticity and the elasticity of the average. The reason is that world trade shares serve as weights while computing the elasticity of the average.

After sorting the bilateral trade flows into ascending order, we run OLS regression multiple times, dropping the 10 smallest trade flows in the sample at each iteration. The estimated $\hat{\theta}$ and corresponding 95% confidence interval are saved and plotted in Figure 2a. This is equivalent to running regressions on different subsamples containing various levels of trade flows, and as the number of iteration increases, the weight of large trade volumes increases as well. The graph demonstrates an unambiguously positive relationship between estimated $\hat{\theta}$ and larger average trade flows in the subsamples. Compared to the results generated from the same procedure without sorting the trade flows in Figure 2b, it is clear that a positive relationship between the estimated $\hat{\theta}$ and larger trade volumes in the subsamples provides

Table 3: The Gravity Equation

	OLS	WLS	PPML (Without ZTF)	PPML (With ZTF)
Distance	-1.347***	-0.722***	-0.770***	-0.750***
	(0.031)	(0.049)	(0.042)	(0.041)
Common border	0.174	0.395***	0.352***	0.369***
	(0.130)	(0.080)	(0.090)	(0.091)
Common language	0.406***	0.419***	0.418***	0.383***
	(0.068)	(0.076)	(0.094)	(0.093)
Colonial ties	0.666***	-0.047	0.038	0.079
	(0.070)	(0.103)	(0.134)	(0.134)
FTA	0.310***	0.424***	0.374***	0.376***
	(0.098)	(0.080)	(0.077)	(0.077)
Exporter FE	Y	Y	Y	Y
Importer FE	Y	Y	Y	Y
R^2	0.75	0.93	0.93	0.93
# Observations	9,613	9,613	9,613	18,630

OLS, WLS, and PPML columns represent the estimates of the main regression estimated by the ordinary least squares, weighted least squares and poisson pseudo-maximum likelihood correspondingly. Zero trade flows are excluded from all specifications except PPML (With ZTF), which includes zero trade flows. Standard errors are in the parentheses. *, **, *** indicate significance at the 10%, 5%, and 1% levels. Details in the main text.

evidence of country-pair heterogeneous effects.[6] Notice, that we included in our regressions exporter and importer fixed effects and standard controls for country pairs, so the results on Figure 2a and 2b are driven by unobserved heterogeneity.

As shown in the previous section, WLS and PPML estimators are robust to country-pair heterogeneous effects in estimating the elasticity of the average in the simulated economy. Now we show that these two methods lead to similar results with the same gravity dataset. We still drop the 10 smallest trade values on each iteration but now apply WLS and PPML estimators to the regression equation. For WLS, the weights are calculated as the share of each observations bilateral trade flows on the global total trade flows.

Compared to Figure 2a, the estimated $\hat{\theta}$ from WLS and PPML shown in Figure 3 do not depend on the choice of subsamples except for the case when the sample size becomes very small. The overall constant estimates show that both estimators are robust to the country-pair heterogeneous effects. The two estimators exhibit a similar pattern, indicating that both of them are capable of addressing the heterogeneity in country-pair elasticity problem,

[6]Notice that the standard errors for the estimates with smaller subsamples in Figure 2b are much larger. It happens because in small subsamples the number of countries increases relative to the number of observations and consequently more fixed effects are included leading to higher standard errors of $\hat{\theta}$.

Figure 2: Evidence of Heterogeneity: OLS Estimates

(a) Increasing Trade Volumes in Subsamples

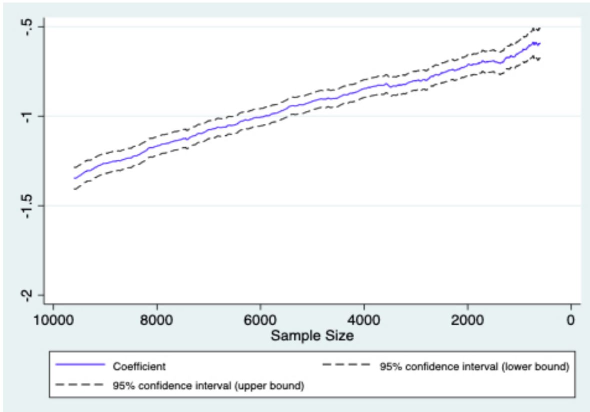

(b) Random Subsamples

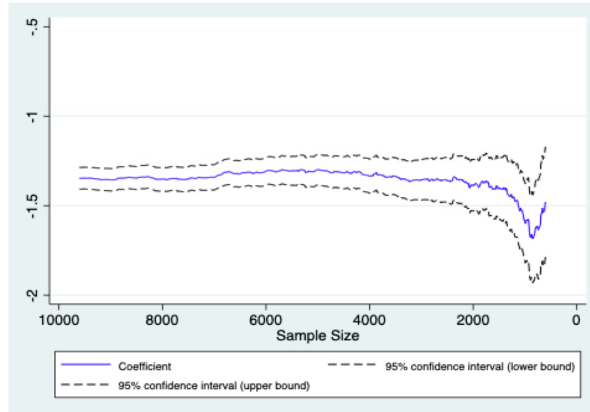

This figure represents OLS estimates of the distance elasticity on different subsamples. On panel (a) at each iteration we drop 10 observations with the smallest corresponding volume of trade. On panel (b) we drop 10 random observations. There are 9613 observations in the full sample. The last iteration has 623 observations.

consistent with our findings in Section 3.

Moreover, SST highlights the robustness of PPML estimators to the bias caused by heteroskedasticity in the error term of the multiplicative gravity equation, and as discussed in Section 3, WLS and OLS both use the log form of the gravity equation to estimate the elasticity. Therefore, WLS and OLS suffer from the same problem that if there is heteroskedasticity in the error term of the original gravity equation. The exogeneity assumption of the transformed regression equation is violated. As a result, the use of the PPML estimator is ultimately the desired way to estimate the ”elasticity of the average” since it is robust to both biases caused by heteroskedasticity and heterogeneity.

With the PPML estimate as the benchmark, we can perform a bias decomposition to show the proportion of bias caused by heterogeneity and heteroskedasticity, respectively. The total bias is calculated as the difference between the OLS estimates and PPML estimates. The bias caused by heterogeneity is calculated as the difference between the OLS estimates and WLS estimates, and the bias caused by heteroskedasticity is the difference between the WLS estimates and PPML estimates. Table 4 summarizes the estimated $\hat{\theta}$, the corresponding standard errors, and the bias decomposition results. It is clear that the bias caused by heteroskedasticity only takes up a small share of the total bias, while heterogeneity is the major driver of the bias.[7] It may seem counter-intuitive that the share of the bias caused by

[7]French (2019) performs an alternative bias decomposition between OLS and PPML, focusing on the aggregation properties of both estimators. In this paper we are agnostic about the source of the heterogeneity elasticity, but if it is industrial heterogeneity, our findings are close to ones in French (2019).

Figure 3: Comparison Across OLS, WLS and PPML Estimators

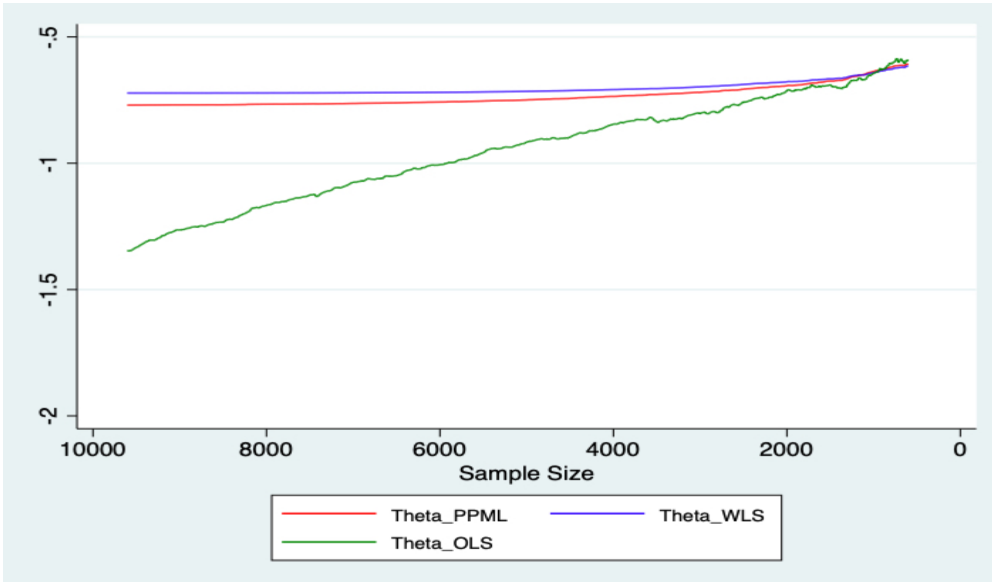

This figure shows OLS, WLS, and PPML estimates on different subsamples. At each iteration we drop 10 observations with the smallest corresponding volume of trade. Full sample has 9613 observations. The last iteration has 623 observations.
Data has been sorted into ascending order.

heterogeneity is greater than 1, but this is the result of having a positive heteroskedasticity bias. As discussed before, the sign and scale of the bias vary with different patterns of the heteroskedasticity. This finding echoes the conclusion of SST that the PPML estimator is so far the best option for estimating the trade elasticity in the gravity equation, and it provides a fresh perspective to appraise the capability of the PPML estimator.

5 Conclusion

In this paper, we argue that the interpretation of PPML regression, the most popular way to estimate gravity equations, is different from ones obtained by the previously dominant OLS regression. We show that the former can be interpreted as the elasticity of the average and the latter as the average elasticity.

We use Monte Carlo simulations to show that when distance elasticity is systematically heterogeneous between country pairs, OLS cannot be used to estimate the elasticity of the average, while WLS and PPML are appropriate methods.

We employ trade data and find evidence of the existence of heterogeneous effects across different country-pairs. Furthermore, comparison of OLS, WLS, and PPML allows us to decompose the difference between OLS and PPML estimates to two channels: extensively

Table 4: Bias Decomposition

	OLS	WLS	PPML (Without ZTF)	PPML (With ZTF)
Estimates	-1.347	-0.722	-0.770	-0.750
	(0.031)	(0.049)	(0.042)	(0.041)
Heterogeneous Bias	-0.625	0	0	-
Percentage	108.32%	0	0	-
Heteroskedastic Bias	0.048	0.048	0	-
Percentage	8.32%	100.00%	0	-
Total Bias	-0.577	0.048	0	-

Total bias is computed as the difference between OLS and PPML (without ZTF) estimates. Heteroskedastic bias is computed as the difference beween WLS and PPML (without ZTF). Heterogeneous bias is computed as the difference between WLS and PPML (without ZTF). PPML (with ZTF) is provided for reference. Details are in main text.

studied previously bias associated with the heteroskedasticity of an error term and different interpretation of two estimators, which we call bias in estimating of the elasticity of the average.

We find that while the bias associated with the heteroskedasticity still exists, it is approximately 7 times smaller compared to the previous findings and has the opposite sign. It gives the choice between PPML and OLS another perspective: when choosing between these two estimation techniques, we recommend making a decision based on the desired interpretation of the results or providing the results obtained with both methods. As the presence of heteroskedasticity bias may still be a problem, analysis can be accompanied by the WLS estimates that would allow a quantification of the size of this bias.

While in this paper we focused on properties of OLS and PPML in estimating the gravity equation, our findings are much more general: any standard log-log regression estimated by PPML will have a different interpretation, which, depending on the research question, may be preferable to the interpretation of a standard OLS regression.

References

Anderson, J.E. 1979, "A Theoretical Foundation for the Gravity Equation", *The American Economic Review*, vol. 69, no. 1, pp. 106-116.

Anderson, J.E. & van Wincoop, E. 2003, "Gravity with Gravitas: A Solution to the Border Puzzle", *The American Economic Review*, vol. 93, no. 1, pp. 170-192.

Anderson, J.E. & Yotov, Y.V. 2010, "The Changing Incidence of Geography", *The American Economic Review*, vol. 100, no. 5, pp. 2157-2186.

Arkolakis, C., Costinot, A. & Rodríguez-Clare, A. 2012, "New Trade Models, Same Old Gains?", *The American Economic Review*, vol. 102, no. 1, pp. 94-130.

Bas, M., Mayer, T. & Thoenig, M. 2017, "From micro to macro: Demand, supply, and heterogeneity in the trade elasticity", *Journal of International Economics*, vol. 108, pp. 1-19.

Bosquet, C., & Boulhol, H. 2009, "Gravity, log of gravity and the" distance puzzle".

Bosquet, C., & Boulhol, H. 2015, "What is really puzzling about the distance puzzle", *Review of World Economics*, vol. 151, no. 1, pp. 1-21.

Brooks, W. J., & Pujolas, P. S. 2018, "Gains from Trade with Variable Trade Elasticities". https://www3.nd.edu/ wbrooks/NLG_May2018.pdf

Chen, N. & Novy, D. 2011, "Gravity, trade integration, and heterogeneity across industries", *Journal of International Economics*, vol. 85, no. 2, pp. 206-221.

Costinot, A. & Rodríguez-Clare, A. 2014, "Trade Theory with Numbers: Quantifying the Consequences of Globalization", *Handbook of International Economics*, vol. 4, pp. 197-261.

Fieler, A.C. 2011, "Nonhomotheticity and Bilateral Trade: Evidence and a Quantitative Explanation", *Econometrica*, vol. 79, no. 4, pp. 1069-1101.

Fratianni, M. & Kang, H. 2006, "Heterogeneous distanceelasticities in trade gravity models", *Economics Letters*, vol. 90, no. 1, pp. 68-71.

French, S., 2019. Comparative Advantage and Biased Gravity. Mimeo

Hausman, J.A., & Wise, D.A., (1981) 'Stratification on Endogenous Variables and Estimation: The Gary Income Maintenance Experiment', in Manski, C.F. & McFadden, D. (eds.) *Structural analysis of discrete data with econometric applications*, Cambridge: MIT Press, pp.365-391.

Head, K. & Mayer, T. 2013, "Gravity Equations: Workhorse, toolkit, and cookbook", *Handbook of International Economics*, vol. 4, pp. 131-195.

Helpman, E., Melitz, M. & Rubinstein, Y. 2008, "Estimating Trade Flows: Trading Partners and Trading Volumes", *The Quarterly Journal of Economics*, vol. 123, no. 2, pp. 441-487.

Imbs, J. & Mejean, I. 2017;2016;, "Trade Elasticities", *Review of International Economics*, vol. 25, no. 2, pp. 383-402.

Kehoe, T.J., Pujols, P.S. & Rossbach, J. 2017, "Quantitative Trade Models: Developments and Challenges", *Annual Review of Economics*, vol. 9, pp. 295-325.

Larch, M., Wanner, J., Yotov, Y.V. and Zylkin, T., 2019. Currency unions and trade: A PPML reassessment with highdimensional fixed effects. Oxford Bulletin of Economics and Statistics, 81(3), pp.487-510.

Magerman, G., Studnicka, Z. & Van Hove, J. 2016, "Distance and border effects in international trade: A comparison of estimation methods", *Economics*, vol. 10, no. 18, pp. 1.

Martin, W. & Pham, C.S. 2015, *Estimating the Gravity Model When Zero Trade Flows are Frequent and Economically Determined*, World Bank, Washington, DC.

Melitz, M.J. 2003, "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity", *Econometrica*, vol. 71, no. 6, pp. 1695-1725.

McCallum, J. 1995, "National Borders Matter: Canada-U.S. Regional Trade Patterns", *The American Economic Review*, vol. 85, no. 3, pp. 615-623.

Novy, D., 2013. International trade without CES: Estimating translog gravity. Journal of International Economics, 89(2), pp.271-282.

Novy, D. and Coughlin, C.C., 2016. Estimating Border Effects: The Impact of Spatial Aggregation. FRB St. Louis Working Paper, (2016-6).

Ossa, R. 2015, "Why trade matters after all", *Journal of International Economics*, vol. 97, no. 2, pp. 266-277.

Santos Silva, J. M. C. & Tenreyro, S. 2006, "The Log of Gravity", *The Review of Economics and Statistics*, vol. 88, no. 4, pp. 641-658.

Tinbergen, J, 1962. "Shaping the World Economy; Suggestions for an International Economic Policy", Books (Jan Tinbergen). Twentieth Century Fund, New York.

Wooldridge, J.M., 2005. Unobserved heterogeneity and estimation of average partial effects. Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg, pp.27-55.

Yotov, Y.V., Piermartini, R., Monteiro, J.A. and Larch, M., 2016. An advanced guide to trade policy analysis: The structural gravity model. Geneva: World Trade Organization.