Estimating the gravity model without gravity using panel data*

Joakim Westerlund† and Fredrik Wilhelmsson‡

February 14, 2006

Abstract

This paper examines the effects of zero trade on the estimation of the gravity model using both simulated and real data. We begin by showing that the usual log-linear estimation method can result in highly deceptive inference when some observations are zero. As an alternative approach, we suggest using the Poisson fixed effects estimator. This approach eliminates the problems of zero trade and is shown to perform well in small samples.

JEL Classification: F10; F15; C15; C23.

Keywords: Gravity Model of Trade; Poisson Regression Model; Panel Data; Monte Carlo Simulation.

1 Introduction

The gravity model of trade has been widely used to estimate the impact of various policy issues, including preferential trade agreements, currency unions, and border effects. The model has a long tradition in social sciences where it has been used to model, for example, migration. In economics, the model has become very popular due to its success in explaining trade flows among countries. Some critique for the lack of theoretical underpinnings has emerged but much progress has been made and now the gravity model rests on a solid theoretical foundation. Instead, the focus has shifted towards the estimation techniques used.

*Westerlund gratefully acknowledges financial support from the Jan Wallander and Tom Hedelius Foundation, research grant number P2005-0117:1. Wilhelmsson gratefully acknowledges financial support from Stiftelsen för främjande av ekonomisk forskning vid Lunds universitet and Sparbanksstiftelsen Färs & Frosta.

†Department of Economics, Lund University, P. O. Box 7082, S-220 07 Lund, Sweden. Telephone: +46 46 222 4970; Fax: +46 46 222 4118; E-mail address: joakim.westerlund@nek.lu.se.

‡Department of Economics, Lund University. Telephone: +46 46 222 4670; Fax: +46 46 222 4113; E-mail address: fredrik.wilhelmsson@nek.lu.se.

The gravity model has traditionally been estimated using cross-sectional data. However, this has been shown to generate biased results since heterogeneity among the countries is typically not controlled for in an appropriate way. To mitigate this problem, researchers have turned towards panel data, which have the advantage that they permit more general types of heterogeneity. For example, consider estimating the impact of currency unions on trade while controlling for country-pair propensity to trade. For a single cross-section, these controls can only depend on observed country-pair attributes such as common language, and estimates can thus be biased if there is additionally an unobserved component to the country-pair propensity to trade. With panel data, such heterogeneity can be readily controlled for by means of a country-pair fixed effect.

The single most popular approach to estimating the gravity model using panel data is to first make it linear by taking logarithms and then to estimate the resulting log-linear model by the fixed effects ordinary least squares (OLS). However, although simple to implement, this approach is problematic because the log-linearized model is not defined for observations with zero trade. Moreover, even though the proportion of observations with zero trade may vary somewhat depending on, among other things, the size of the sample, it is usually quite significant suggesting that the proper handling of these zeros is potentially very important. Another problem is that the OLS estimates of the log-linearized model may be both biased and inefficient in the presence of heteroskedasticity.

Two common approaches to handle the presence of zero trade include simply discarding the zeros from the sample or to add a constant factor to each observation on the dependent variable. The first strategy is correct as long as the zeros are randomly distributed. However, if the zeros are not random, as is usually the case, then this induces a selection bias. This problem is often ignored in applied work, but could be handled by using sample selection correction. In a recent contribution, Helpman *et al.* (2005) propose a theoretical model rationalizing the zero trade flows and suggest estimating the gravity equation with a correction for the probability of countries to trade. To estimate the model they apply a two-step estimation technique similar to sample selection models. However, in order to implement the new estimator, the researcher needs to find a suitable exclusion restriction for identification of the second stage equation, which can be very difficult. The problem with bias and inefficiency in the presence of heteroskedasticity has been largely ignored by applied researchers.

In this paper, we propose estimating the gravity model directly from its non-linear form by using the fixed effects Poisson maximum likelihood (ML) estimator. Since this removes the need to linearize the model by taking logarithms, the problem with zero trade disappears. Our simulation results suggest that the new estimation method is superior to the conventional approach of applying OLS to the log-linearized model. In particular, it is shown that the conventional approach is likely to result in severe bias and misleading inference even if the fraction of observations with zero trade is very small. On the other hand, the Poisson ML estimator generally performs very well with only small bias and size distortion. Therefore, since the Poisson ML estimator is becoming increasingly available using standard statistical software packages, these results suggest that it should be a valuable tool for econometric analysis of the gravity model. As an empirical illustration, we consider the trade effects of the 1995 European Union (EU) enlargement.

The remainder of this paper is organized as follows. Section 2 briefly outlines the gravity model and the problems of zero trade. Section 3 then presents the Monte Carlo simulations, while Section 4 contains the application. Section 5 concludes.

2 The problem of zero gravity

Let M_{ijt} denote the bilateral trade between countries i and j at time $t = 1, ..., T$, as measured by the imports of country i from country j. For convenience, the total number of country pairs is henceforth denoted by N. A common empirical formulation of the gravity model for bilateral trade includes the GDP levels of the two countries, Y_{it} and Y_{jt} say, as well as D_{ijt}, a dummy variable representing for example some contiguity, common language or free-trade agreement effect. This formulation of the gravity equation can be written algebraically as

$$\mathrm{E}(M_{ijt}|Y_{it}, Y_{jt}) = \exp(\gamma D_{ijt}) Y_{it}^{\beta_1} Y_{jt}^{\beta_2}. \tag{1}$$

Because only a very limited amount of heterogeneity between the country pairs is allowed in the parametrization of the regression function, conventional cross-section estimates of the gravity model are generally biased. With panel data, on the other hand, we can easily permit for such heterogeneity by means of N country-pair specific effects, denoted α_{ij}. These effects may be different depending on the direction of trade and enters (1) multiplicatively in the following fashion

$$\mathrm{E}(M_{ijt}|Y_{it}, Y_{jt}, \alpha_{ij}) = \exp(\alpha_{ij} + \gamma D_{ijt}) Y_{it}^{\beta_1} Y_{jt}^{\beta_2}.$$

This implicitly defines the following regression

$$M_{ijt} = \exp(\alpha_{ij} + \gamma D_{ijt}) Y_{it}^{\beta_1} Y_{jt}^{\beta_2} + e_{ijt},$$

which can be written equivalently as

$$M_{ijt} = \exp(\alpha_{ij} + \gamma D_{ijt}) Y_{it}^{\beta_1} Y_{jt}^{\beta_2} v_{ijt}, \tag{2}$$

where e_{ijt} is a mean zero disturbance that is independent of the regressors and $v_{ijt} = 1 + e_{ijt}/ \exp(\alpha_{ij} + \gamma D_{ijt}) Y_{it}^{\beta_1} Y_{jt}^{\beta_2}$ is a heteroskedastic disturbance term with $\mathrm{E}(v_{ijt}|Y_{it}, Y_{jt}) = 1$. Moreover, since α_{ij} will generally be correlated with the explanatory variables, random effects estimation of (2) will be inconsistent. To circumvent this, it is common to treat α_{ij} as fixed, which is equivalent to allowing each country pair to have its own dummy variable.

Suppose for a moment that M_{ijt} is strictly positive. One of the most common approaches to estimate the regression in (2) is to first make it linear by taking logarithms, which yields

$$\ln(M_{ijt}) = \alpha_{ij} + \gamma D_{ijt} + \beta_1 \ln(Y_{it}) + \beta_2 \ln(Y_{jt}) + \ln(v_{ijt}). \tag{3}$$

Since the model is now linear, it is readily estimable using OLS. However, this is only possible as long as M_{ijt} is nonzero, which is not always the case. Indeed, a common feature of trade data is that the bilateral trade can sometimes be zero. Although this poses no problem when estimating the gravity model based on its multiplicative form in (2), as the logarithm is defined only for positive outcomes, the log-linear regression (3) is no longer admissible. A common solution to this problem is to drop all observations with zero trade and then to estimate (3) based on the resulting truncated sample. However, although this approach certainly eliminates the zeros, it simultaneously induces a bias to the OLS estimator, which is why truncating the sample should be avoided as a matter of practice.

A natural alternative approach in situations such as this, when the model cannot be log-linearized, is to estimate it from its multiplicative form directly. In so doing, note that the fixed effects conditional mean can be written as

$$\mathrm{E}(M_{ijt}|Y_{it}, Y_{jt}) = \exp(\alpha_{ij} + \gamma D_{ijt} + \beta_1 \ln(Y_{it}) + \beta_2 \ln(Y_{jt})), \tag{4}$$

which is known as the exponential regression function. This regression follows naturally from the multiplicative form of (1) and ensures that $\mathrm{E}(M_{ijt}|Y_{it}, Y_{jt})$ is nonnegative, which is very convenient as trade cannot be negative. Thus, the conventional additive regression $\mathrm{E}(M_{ijt}|Y_{it}, Y_{jt}) = \alpha_{ij} + \gamma D_{ijt} + \beta_1 \ln(Y_{it}) + \beta_2 \ln(Y_{jt})$ is likely to be unsatisfactory here as it cannot ensure the nonnegativity of trade.

The estimation of (4) has been studied by Hausman *et al.* (1984), who consider the special case when the data is measured in nonnegative integers. They propose using a version of the conventional Poisson ML estimator, which is modified to account for the fixed effects. In so doing, the authors condition on $\sum_{t=1}^{T} M_{ijt}$, which is shown to eliminate the fixed effects from the log-likelihood function. As noted by the authors, given that the regression in (4) is correctly specified, consistency of the resulting fixed effects Poisson ML estimator follows directly by standard ML theory (see, e.g. Gourieroux *et al.*, 1984).

On the other hand, valid inference requires the correct specification of both the conditional mean and variance, which necessitates that

$$\mathrm{E}(M_{ijt}|Y_{it}, Y_{jt}) = \mathrm{var}(M_{ijt}|Y_{it}, Y_{jt}). \tag{5}$$

However, note that the validity of (4) and (5) does not require the data to be Poisson distributed. In fact, M_{ijt} does not have to be an integer at all. This suggests that we can use the fixed effects Poisson ML to estimate the gravity model. Since this estimator does not require M_{ijt} to be nonzero, it is expected to produce better results than OLS in panels where some trade flows are zero. Moreover, if it is consistency that we are interested in, then (5) does not have to hold either, so the data do not have to be equidispersed. In the next section, we elaborate on this point.

3 Monte Carlo study

In this section, we investigate the small-sample properties of the OLS and ML estimators in the presence of zero observations through Monte Carlo simulations. The data generating process used for this purpose is given as follows

$$M_{ijt} = \exp(\alpha_{ij} + \gamma D_{ijt} + \beta Y_{ijt}) v_{ijt}, \tag{6}$$

where $\alpha_{ij} = \gamma = \beta = 1$ for simplicity. Since Y_{ijt} is usually positive in applied work, we set $Y_{ijt} \sim U(0,1)$. Moreover, if we let $\tau_{ij} \sim U(0,1)$ denote the location of the break, then the dummy variable D_{ijt}, representing for example a preferential trade agreement, is such that $D_{ijt} = 1$ if $t > \tau_{ij}T$ and zero otherwise.

The disturbance v_{ijt} is key in this data generating process. In particular, it is assumed that v_{ijt} is a log-normally distributed variable with mean one and variance σ_{ij}^2. For the variance, we have two scenarios. In the first, $\sigma_{ij}^2 = 1$, which implies that $\mathrm{var}(M_{ijt}|Y_{ijt}) = \exp(\alpha_{ij} + \gamma D_{ijt} + \beta Y_{ijt})^2$, while, in the second, we have $\sigma_{ij}^2 = \exp(\alpha_{ij} + \gamma D_{ijt} + \beta Y_{ijt})^{-1}$ so that $\mathrm{var}(M_{ijt}|Y_{ijt}) = \exp(\alpha_{ij} + \gamma D_{ijt} + \beta Y_{ijt})$. Thus, we expect the OLS estimator to perform relatively well in Case 1, while, since $\mathrm{var}(M_{ijt}|Y_{ijt}) = \mathrm{E}(M_{ijt}|Y_{ijt})$, we expect the Poisson ML estimator to perform relatively well in Case 2.[1] In both cases, we generate data by drawing $1,000$ panels with N cross-sectional and T time series observations.

The results are organized according to the two cases described above. In each case, we want to examine the effect of zero observations in the data. Both the OLS and Poisson ML estimators are considered.[2,3] The former is implemented using both truncated data and $\ln(M_{ijt} + 1)$ as dependent variable. However, note that since $M_{ijt} > 0$ in this data generating process, the log-linear model is no longer inadmissible. Hence, to be able to study the effect of truncating the

[1]Other values of σ_{ij}^2 produced very similar results and are thus not reported.

[2]We also examined the negative binomial ML estimator of Hausman *et al.* (1984), which relaxes the Poisson condition that $\mathrm{var}(M_{ijt}|Y_{ijt}) = \mathrm{E}(M_{ijt}|Y_{ijt})$. However, since this estimator generally performed very unsatisfactory, the results are not reported here but are available from the authors upon request.

[3]The Poisson ML estimator is implemented using the GAUSS optimization library OPTMUM. We use the BFGS gradient algorithm with numerical derivatives. The standard errors of the estimated parameters are computed based on the conventional Hessian method, which generally worked best in the simulations. The truncated OLS is used to start up the estimation.

sample we use a positive truncation threshold parameter, which is such that the fraction of truncated observations is exactly λ. For brevity, we only report the mean bias and the size of a nominal 5% level t-test of the null hypothesis that the parameter of interest is equal to its true value versus the alternative that it is not.

The results reported in Table 1 can be summarized as follows. First, as expected, OLS estimation with $\ln(M_{ijt}+1)$ as the dependent variable generally produces very poor results. In particular, it is seen that the estimators of γ and β both suffer from substantial downwards bias, which do not tend to vanish as the sample size increases. Moreover, the results on the size of the t-tests suggest that inference based on this estimation method is likely to be highly deceptive. In fact, with this method, we always end up rejecting the null hypothesis. Thus, based on these results, we recommend not using OLS estimation based on $\ln(M_{ijt}+1)$.

Second, the results on the truncated OLS estimator are mixed. At one end of the scale, we have Case 1 when there is no truncation, in which the performance, both in terms of bias and size accuracy, is very good. At the other end, we have the $\lambda > 0$ case where Table 1 shows that the performance is poor, and that the problems with bias and size distortion are highly potent, even for a truncation as small as 10%. Apparently, the truncation makes the OLS estimator both downwards biased and unfit for inference. Thus, from an empirical point of view, it seems highly unlikely that the truncated OLS is able to deliver any meaningful results at all.

In addition to the problems associated with truncating the data, Table 1 points to another important shortcoming with the truncated OLS estimator. In particular, it seems as that the heteroskedasticity in Case 2 induces both severe size distortions as well as a sizeable bias that persists even in large panels.

Although this may appear somewhat counterintuitive at first, as pointed out by Santos Silva and Tenreyro (2005), it is actually a direct consequence of the well-known Jensen inequality. To appreciate this, consider the data generating process in (6) where $\mathrm{E}(v_{ijt}|Y_{ijt}) = 1$. The OLS estimates of the parameters in the log-linear model (3) are consistent only if $\mathrm{E}(\ln(v_{ijt})|Y_{ijt}) = 0$. However, although $\ln(\mathrm{E}(v_{ijt}|Y_{ijt})) = 0$, by the Jensen equality, $\mathrm{E}(\ln(v_{ijt})|Y_{ijt}) \neq 0$. Indeed, since $\mathrm{E}(v_{ijt}|Y_{ijt})^2 = 1$ in our case, by using the properties of the log-normal distribution, we have that

$$\mathrm{E}(\ln(v_{ijt})|Y_{ijt}) \quad = \quad \ln\left(\frac{1}{1+\sigma_{ij}^2}\right),$$

which is not equal to zero unless of course σ_{ij}^2 is zero too. As a result, the OLS estimator in (3) will generally be biased.

Third, except possibly for Case 1 when there is no truncation, the results show that the Poisson ML consistently outperforms the other estimators in terms of bias. In fact, by looking at Table 1, it would appear as that the bias is practically nonexisting even for as small panels as $T = 10$ and $N = 500$, which correspond approximately to 10 time series observations on six countries. We also see that the size is very close to the nominal 5% level in Case 2 but that it is distorted in Case 1, which is partly expected since $\mathrm{var}(M_{ijt}|Y_{ijt}) \neq \mathrm{E}(M_{ijt}|Y_{ijt})$ in this case.

One possibility to get rid of the distorted standard errors of the ML estimator is to use the bootstrap. This approach has become very popular in applied work and it will therefore be used in this paper. Some simulations on the resulting bootstrapped t-statistic are reported in Table 2.[4] As expected, we see that the size of the bootstrapped test generally lies much closer to the 5% level than the size of the asymptotic test. Also, the t-statistics appear to be well centered around zero.

In summary, we find that the Poisson ML show smaller bias than the two OLS estimators considered and, at the same time, maintain relatively good size properties in small samples. Since the Poisson ML with bootstrapped standard errors is now readily available through existing software packages such as STATA, it should be considered a feasible alternative to estimation by OLS.

4 An application to the 1995 EU enlargement

We have shown that log-linear OLS estimation of the gravity model yields biased results. In this section, we demonstrate these findings by estimating the trade effects of the adhesion of Austria, Finland and Sweden to the EU in 1995. The sample that we use for this purpose cover the period 1992 to 2002 and consists of import data for EU and other developed countries from all trade partners except oil exporting countries and formerly planned economies in Central and Eastern Europe, as defined in Direction of Trade Statistics (International Monetary Fund, 2005). The GDP and population data comes from World Development Indicators (World Bank, 2005).

The estimated gravity equation can be written as

$$M_{ijt} = \exp(\alpha_{ij} + \lambda_t + \gamma_1 D_{it} + \gamma_2 D_{jt} + \gamma_3 D_{ijt}) Y_{it}^{\beta_1} Y_{jt}^{\beta_2} N_{it}^{\beta_3} N_{jt}^{\beta_4} v_{ijt}, \quad (7)$$

or equivalently in its log-linear form

$$\begin{aligned} \ln(M_{ijt}) &= \alpha_{ij} + \lambda_t + \gamma_1 D_{it} + \gamma_2 D_{jt} + \gamma_3 D_{ijt} + \beta_1 \ln(Y_{it}) + \beta_2 \ln(Y_{jt}) \\ &+ \beta_3 \ln(N_{it}) + \beta_4 \ln(N_{jt}) + \ln(v_{ijt}), \quad (8) \end{aligned}$$

where M_{ijt} denotes the nominal imports of country i from country j, Y_{it} and Y_{jt} denote the real GDP of the two countries, and N_{it} and N_{jt} denote their population. The fixed effects α_{ij} capture all types of unobserved country-pair

[4] The algorithm used is taken from Gonçalves and White (2004), who proposes a moving block bootstrap for use with non-linear models. The bootstrap results are based on 100 replications and a block length of one.

specific heterogeneity that is constant over time, while the time effects λ_t capture all forms of time-varying heterogeneity that is shared among the country pairs.

The dummy variables D_{it}, D_{jt} and D_{ijt} are key in this model. The variable D_{it} equals one if country i is a member of the EU at time t while country j belongs to the rest of the world. The second dummy variable D_{jt} equals one if country j is a member of the EU while i belongs to the rest of the world. Similarly, D_{ijt} equals one if both i and j are members of the EU at time t. In other words, the three dummy variables take the value one for EU imports from the rest of the world, EU exports to the rest of the world and intra-EU trade, respectively.

The rest of the world is defined as all countries in the sample that are not members of the EU at any given time in the sample. This enables us to identify the effect of the enlargement on the trade of new EU members as opposed to the effect of changes in the size of the rest of the world. To appreciate this, note that if the rest of the world also included new members, the dummy variable D_{it} would capture not only the import effect on the new members but also the effect of the change in the composition of the rest of the world, as the imports from the new members to the old ones would no longer be classified as imports from the rest of the world. A similar argument applies to the construction of D_{jt}.

A consequence of this definition of the rest of the world is that, since fixed effects absorb all heterogeneity that is constant over time, the trade effect for countries that have been members of the EU for the whole sample period cannot be identified. Thus, the dummy variables capture only the effect on countries that have changed their EU status at least one time. That is, the dummy variables capture the effect of the Austrian, Finnish and Swedish accession to the EU. Specifically, γ_1 measures the trade diversion or changes in EU imports from the rest of the world. Similarly, γ_2 measures the effect on EU exports to the rest of the world, sometimes called export diversion. Finally, γ_3 measures trade creation, resulting from the increased intra-EU trade following the enlargement.

Economic integration should increase trade between countries integrating. Thus, we expect the trade creation, as measured by γ_3, to be positive. This effect can be separated into pure trade creation, i.e. increased trade due to lower prices on imports from the other countries in the EU, and trade diversion, which implies a shift in imports from more efficient producers in the rest of the world to less efficient producers within the EU. A negative sign on γ_1 would thus indicate trade diversion. Similarly, export diversion occurs if exports to the rest of the world decreases as a result of the integration process, but exports could also increase. The expected sign of γ_2 is therefore ambiguous.

The empirical results are contained in Table 3. It is seen that the enlargement of the EU induced significant trade diversion but no trade creation. This absence of trade creation is, however, not surprising since the new members were part of a free trade area with the EU prior to the membership. When joining the EU, the new members implemented the Common External Tariff, which changed the tariffs on their imports from the rest of the world. Note that the trade diversion effect is rather large in comparison to the trade creation effect. Although counterintuitive at first, one should keep in mind that several countries with preferential access to the EU market, such as those that joined the EU in 2004, have been excluded from our sample, so trade might have been diverted away from suppliers on the world market to suppliers with preferential access to the EU market. Moreover, taken as a fraction of total trade, the diversion effect is probably quite small since the estimation results only capture the effect on imports to Austria, Sweden and Finland and not changes in the total imports of the EU.

Even though the number of zeros is comparatively small in our sample, only 10%, when comparing the results obtained from the various estimators, we see that the difference can be substantial. In particular, for the GDP and population variables, the Poisson ML estimates are typically larger than their OLS counterparts. This finding is well in line with the Monte Carlo evidence suggesting that both OLS estimators are downwards biased. Moreover, while the truncated OLS estimator indicates that changes in the importing countries GDP does not effect imports, the ML estimator gives a more plausible estimate close to unity.

For the dummy variables, the differences are less marked. In particular, although the sign and significance of the estimates do not differ much, the magnitude of the estimates varies quite substantially. The OLS estimator indicates that the trade diversion is twice as large as implied by the ML estimator and, while the OLS estimate of the trade creation effect is slightly negative, it is positive for the ML estimator.

In summary, the results presented in this section highlight the importance of using appropriate estimation techniques to be able to draw correct inference.

5 Conclusions

The gravity model has become a standard tool for evaluating policies affecting trade and it is widely used to assess the effects of preferential trade agreements and currency unions or to calculate trade potential, among other things. It is well known that the gravity model should be estimated by panel data to mitigate the bias due to failure to fully control for country heterogeneity. A very popular way to accomplish this is to first linearize the model by taking logarithms and then to apply the conventional fixed effects OLS estimator.

In this paper, we argue that this approach is likely to be very misleading with severely biased estimates and t-statistics. There are two reasons for this. Firstly, since trade cannot be zero in the log-linearized model, all zeros must either be discarded or replaced by some arbitrary positive value, which induces a sample selection bias. Secondly, the heteroskedasticity inherent in the log-linear formulation of the gravity model can render the OLS estimates both biased and inefficient. By contrast, being based on the gravity model in its original non-linear form, the fixed effects Poisson ML estimator does not suffer from these weaknesses and is therefore expected to yield more accurate results.

Our assertion is verified by means of Monte Carlo simulations and illustrated via an application to the 1995 EU enlargement. The simulations show that the performance of the log-linear approach is likely to be so poor that it may not even be meaningful to interpret the results. On the other hand, the Poisson ML estimator performs well with only a very small bias and good size accuracy in most cases. Still, in some data generating processes, the results show that the estimated standard errors can be downward biased. To alleviate this, we suggest using bootstrapped standard errors. The empirical application points to a significant difference between the estimators with respect to both the main explanatory variables and the trade effects of the 1995 EU enlargement, thus underlining the importance of using the proper estimation technique.

To conclude, we recommend not estimating the gravity model from its log-linear form. Instead, we propose estimating the model directly from its non-linear form using the fixed effects Poisson ML estimator with bootstrapped standard error. Since this estimator can now be implemented using standard statistical software packages such as STATA. That is, our analysis provide a foundation for a simple way of estimating the gravity model avoiding biased results due to country-pairs not trading; hence gives empirical researchers, using the gravity model, a valuable tool.

Table 1: Simulated bias and size for the ML and OLS estimators.

				Mean bias						Size at the 5% level					
λ	Case	N	T	$\widehat{\gamma}^{ols}$	$\widehat{\beta}^{ols}$	$\widetilde{\gamma}^{ols}$	$\widetilde{\beta}^{ols}$	$\widehat{\gamma}^{mle}$	$\widehat{\beta}^{mle}$	$\widehat{\gamma}^{ols}$	$\widehat{\beta}^{ols}$	$\widetilde{\gamma}^{ols}$	$\widetilde{\beta}^{ols}$	$\widehat{\gamma}^{mle}$	$\widehat{\beta}^{mle}$
0	1	500	10	0.1	−0.1	−36.4	−36.9	0.1	−0.1	7.3	5.8	100.0	100.0	24.6	31.8
		1000	10	0.0	0.0	−36.5	−36.8	−0.1	−0.1	6.6	5.8	100.0	100.0	25.6	32.2
		500	20	0.1	0.1	−36.5	−36.7	0.1	0.2	6.2	6.1	100.0	100.0	24.1	33.4
		1000	20	−0.1	0.0	−36.5	−36.8	−0.1	−0.2	5.9	6.2	100.0	100.0	26.5	31.9
	2	500	10	13.8	14.1	−25.8	−26.6	−0.1	0.0	100.0	99.8	100.0	100.0	5.0	6.3
		1000	10	13.8	14.2	−25.8	−26.5	0.1	0.1	100.0	100.0	100.0	100.0	3.7	4.4
		500	20	13.8	14.1	−25.8	−26.6	0.0	0.0	100.0	100.0	100.0	100.0	6.1	4.8
		1000	20	13.9	14.2	−25.8	−26.6	0.1	0.1	100.0	100.0	100.0	100.0	4.9	5.3
0.1	1	500	10	−21.8	−21.8	−36.5	−36.8	−0.1	−0.1	100.0	99.9	100.0	100.0	25.3	29.6
		1000	10	−21.9	−21.8	−36.6	−36.8	−0.1	−0.1	100.0	100.0	100.0	100.0	23.9	29.5
		500	20	−21.9	−21.9	−36.6	−36.8	−0.1	−0.1	100.0	100.0	100.0	100.0	27.1	31.3
		1000	20	−21.8	−21.9	−36.5	−36.8	0.0	0.0	100.0	100.0	100.0	100.0	25.3	32.2
	2	500	10	−9.4	−13.4	−25.7	−26.6	0.1	−0.1	99.8	99.8	100.0	100.0	4.7	4.9
		1000	10	−9.5	−13.5	−25.8	−26.6	0.0	−0.1	100.0	100.0	100.0	100.0	5.5	4.9
		500	20	−9.4	−13.4	−25.8	−26.6	0.0	−0.1	100.0	100.0	100.0	100.0	5.6	4.5
		1000	20	−9.5	−13.3	−25.8	−26.5	0.0	0.1	100.0	100.0	100.0	100.0	5.7	5.7
0.3	1	500	10	−44.8	−40.9	−36.5	−36.8	−0.1	−0.2	100.0	100.0	100.0	100.0	27.2	29.5
		1000	10	−44.7	−41.0	−36.5	−36.8	0.0	−0.2	100.0	100.0	100.0	100.0	26.6	34.3
		500	20	−44.6	−40.9	−36.5	−36.8	0.1	0.1	100.0	100.0	100.0	100.0	28.8	33.2
		1000	20	−44.8	−41.0	−36.6	−36.8	−0.1	−0.1	100.0	100.0	100.0	100.0	26.7	32.9
	2	500	10	−41.4	−29.0	−25.9	−26.6	−0.1	−0.1	100.0	100.0	100.0	100.0	4.0	5.1
		1000	10	−41.4	−28.9	−25.8	−26.6	0.0	0.1	100.0	100.0	100.0	100.0	6.1	3.9
		500	20	−41.4	−29.2	−25.8	−26.7	0.0	−0.1	100.0	100.0	100.0	100.0	5.3	5.5
		1000	20	−41.4	−29.1	−25.8	−26.6	0.0	0.1	100.0	100.0	100.0	100.0	5.0	5.2

Notes: The value λ refers to the fraction of truncated observations, $\widehat{\gamma}^{ols}$ and $\widehat{\beta}^{ols}$ refer to the truncated OLS estimates, $\widetilde{\gamma}^{ols}$ and $\widetilde{\beta}^{ols}$ refer to the OLS estimates with $\ln(M_{ijt}+1)$ as dependent variable, and $\widehat{\gamma}^{mle}$ and $\widehat{\beta}^{mle}$ refer to the Poisson ML estimates. Case 1 refers to the data generating process with $\sigma_{ij}^2 = 1$ while Case 2 refers to the data generating process with $\text{var}(M_{ijt}|Y_{ijt}) = \text{E}(M_{ijt}|Y_{ijt})$. The reported bias results refer to the mean bias times 100.

Table 2: Simulation results for the bootstrapped ML t-statistic.

		Case 1				Case 2			
N	T	$t(\widehat{\gamma})$	$t(\widehat{\beta})$	$t_b(\widehat{\gamma})$	$t_b(\widehat{\beta})$	$t(\widehat{\gamma})$	$t(\widehat{\beta})$	$t_b(\widehat{\gamma})$	$t_b(\widehat{\beta})$
Size on the 5% level									
500	10	23.8	33.4	9.2	10.0	5.8	4.0	10.2	7.4
1000	10	25.2	31.4	10.4	9.6	4.6	5.8	10.0	7.2
500	20	28.4	33.2	7.6	9.6	5.0	6.6	7.8	10.0
1000	20	26.4	39.4	8.4	10.6	6.4	4.8	8.2	6.6
Mean of t-statistic									
500	10	0.0	−0.1	0.0	−0.1	0.1	0.0	0.1	0.1
1000	10	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.0
500	20	0.0	−0.1	0.0	0.0	0.1	0.0	0.1	0.0
1000	20	0.1	0.0	0.1	0.0	0.0	−0.1	0.0	−0.1
Standard deviation of t-statistic									
500	10	1.7	2.1	1.2	1.2	1.0	1.0	1.2	1.1
1000	10	1.8	1.9	1.2	1.2	1.0	1.0	1.2	1.1
500	20	1.8	2.1	1.1	1.2	1.0	1.1	1.1	1.1
1000	20	1.8	2.2	1.1	1.2	1.0	1.0	1.1	1.1

Notes: The values $t(\widehat{\gamma})$ and $t(\widehat{\beta})$ refer to the conventional asymptotic ML t-ratios while $t_b(\widehat{\gamma})$ and $t_b(\widehat{\beta})$ refer to the bootstrapped counterparts. See Table 1 for an explanation of the remaining features of the table.

Table 3: Empirical estimation results.

Estimator	OLS	OLS	Poisson ML
Dependent variable	$\ln(M_{ijt})$	$\ln(M_{ijt}+1)$	M_{ijt}
β_1	-0.091	0.229^{***}	0.931^{***}
	(0.191)	(0.062)	(0.173)
β_2	1.438^{***}	0.820^{***}	1.483^{***}
	(0.084)	(0.039)	(0.110)
β_3	4.055^{***}	1.765^{***}	2.471^{***}
	(0.612)	(0.267)	(0.629)
β_4	-1.275^{***}	-0.979^{***}	-0.580
	(0.190)	(0.074)	(0.357)
γ_1	-0.403^{***}	-0.211^{***}	-0.232^{***}
	(0.046)	(0.016)	(0.074)
γ_2	0.000	0.102^{***}	0.041
	(0.032)	(0.023)	(0.047)
γ_3	-0.002	0.033^{*}	0.035
	(0.025)	(0.018)	(0.034)
No. of country-pairs	2719	2748	2719
No. of observations	32487	35600	35256

Notes: The numbers within the parantheses are the robust OLS standard errors or the bootstrapped Poisson ML standard errors. The superscripts (***), (**) and (*) denote significance at the 1%, 5% and 10% levels, respectively.

References

Hausman, J., B. H. Hall and Z. Griliches (1984). Econometric models for count data with an application to patents-R&D relationship, *Econometrica*, Vol. 52, pp. 909-938.

Helpman, E., M. Melitz and Y. Rubinstein (2005). Trading Partners and Trading Volumes, Unpublished manuscript.

Gonçalves, S., and H. White (2004). Maximum likelihood and the bootstrap for nonlinear dynamic models, *Journal of Econometrics*, Vol. 119 pp. 199-219.

Gourieroux, C., A. Monfort and A. Trognon (1984). Pseudo maximum likelihood methods: Theory, *Econometrica*, Vol. 52, pp. 681-700.

International Monetary Fund (2005). Direction of Trade Statistics, CD-Rom, *International Monetary Fund*, Washington DC.

Santos Silva, J. M. C., and S. Tenereyro (2005). The log of gravity, *Review of Economics and Statistics*, forthcoming.

World Bank (2005). World Development Indicators, CD-Rom, *The World Bank*, Washington DC.