# Gravity Estimation:

# Best Practices and Useful Approaches

Peter R. Herman

ECONOMICS WORKING PAPER SERIES
Working Paper 2023–10–C

U.S. INTERNATIONAL TRADE COMMISSION
500 E Street SW
Washington, DC 20436

October 2023

Gravity Estimation: Best Practices and Useful Approaches
Peter R. Herman
Economics Working Paper 2023–10–C
October 2023

**Abstract**

The gravity model has proven to be an exceptionally successful tool for empirically studying international trade. However, deriving good estimates requires careful consideration of the econometric specification being used. Common applications of the model are prone to a variety of biases and identification challenges that should not be overlooked. This paper describes notable approaches for mitigating these issues and provides guidance on how to produce high quality gravity estimates. It also demonstrates how to implement these techniques in Stata.

Peter R. Herman
Research Division, Office of Economics
U.S. International Trade Commission
peter.herman@usitc.gov

# 1 Introduction

The gravity model has been a workhorse empirical model in international trade for decades. However, during that time frame, the ways in which the model is specified and estimated have changed drastically. Many innovations have addressed pervasive issues like biases that were common in early iterations of the model. Others have introduced useful techniques for improving the identification of parameters of interest. Given these advancements, many specifications in top publications from 20+ years ago are unrecognizable compared to current best practices.

This paper provides a concise survey of the most influential and important evolutions in empirical gravity modeling since its emergence and rise in popularity. It highlights methods that should be employed in order to produce reliable, unbiased estimates. It also notes approaches that practitioners may find generally useful in many common applications.

The paper also aims to provide instruction on how to implement each of these techniques in common statistical programs. While it focuses on Stata, which has some of the best specialized support for gravity analysis, the applied examples should help make the general implementation more concrete and help with adaptation to other software.

In total, six major advances in empirical gravity methods are discussed (Table 1). As a general rule of thumb, practitioners should incorporate approaches 1 and 2—multilateral resistances and Poisson Pseudo Maximum Likelihood estimation (PPML)—in virtually every gravity specification. The other five approaches offer ways to address many of the challenges widely encountered in this type of analysis, such as empirical biases and other identification issues. While these latter techniques may not be effective or appropriate in every case, they are generally useful, widely applicable, and ought to be considered for most analyses.

The remained of the paper is structured as follows. Section 2 provides a brief description of some of the theoretical foundations of the gravity model of trade. Section 3 describes each of the different approaches and how they can be easily implemented in Stata. Section 4 presents a series of estimates derived using multiple different empirical specifications for comparison. Section 5 describes useful sources of data for estimating gravity models. Finally, section 6 concludes.

# 2 Gravity Foundations

The gravity model has rapidly become one of the workhorse models in international trade. Its usefulness and appeal lie in the fact that it is both based on strong theoretical foundations and performs exceptionally well in empirical applications using extensive, readily available data. It is a valuable tool for econometrically

Table 1: Notable empirical approaches in gravity modeling

| Approach | | Discussion |
|---|---|---|
| 1. | Multilateral resistances | Section 3.2 |
| 2. | Poisson Pseudo Maximum Likelihood | Section 3.3 |
| 3. | Country-pair fixed effects | Section 3.4 |
| 4. | Domestic trade | Section 3.5 |
| 5. | Global openness and country-pair trends | Section 3.6 |
| 6. | Three-way bias correction | Section 3.7 |
| 7. | Disaggregated gravity | Section 3.8 |

identifying the factors that drive trade patterns and analyzing counterfactual policy scenarios based on these factors. While this guide is primarily focused on providing a survey of empirical innovations in gravity modeling, it is helpful to begin with a short discussion of a theoretical formulation of the model.

## 2.1 Basic Theory

It has been widely shown that many different models of trade arrive at comparable gravity formulations (Arkolakis et al., 2012). One of the most popular formulations is that of Anderson (1979) and Anderson and van Wincoop (2003). Following their notation, the gravity model is represented by the following system:

$$X_{ijt} = \frac{Y_{it}E_{jt}}{Y_t}\left(\frac{\tau_{ijt}}{\Pi_{it}P_{jt}}\right)^{1-\sigma}, \tag{1}$$

$$\Pi_{it}^{1-\sigma} = \sum_j \left(\frac{\tau_{ijt}}{P_{jt}}\right)^{1-\sigma}\frac{E_{jt}}{Y_t}, \tag{2}$$

$$P_{jt}^{1-\sigma} = \sum_i \left(\frac{\tau_{ijt}}{\Pi_{it}}\right)^{1-\sigma}\frac{Y_{it}}{Y_t}, \tag{3}$$

Equation 1 is the "gravity" equation. Bilateral trade flows ($X_{ijt}$) from exporter $i$ to importer $j$ in period $t$ are modeled as a function of two components: the economic size of the two countries and trade costs. Economic size is expressed as the product of the exporter's total output ($Y_{it}$) and the importers total expenditure ($E_{jt}$), divided by total world output ($Y_t$). Intuitively, trade between $i$ and $j$ generally increases if $i$ increases its output and/or if $j$ increases its expenditures.

The second component reflects trade costs, both between $i$ and $j$ specifically but also for $i$ and $j$ individually. One part of trade costs are the specific bilateral costs between the two countries ($\tau_{ijt}$). The other parts of trade costs reflect the two so-called "multilateral resistance terms" (MRTs), which capture a broader range of trade frictions. The first is the outward MRT ($\Pi_{it}$), which can be though of as an aggregate export cost faced by the exporter, regardless of the specific partner it is trading with. Similarly, the second is the inward MRT ($P_{jt}$), which represents the aggregate trade costs for imports into country $j$. Equations 2 and 3, respectively, define these terms explicitly. In both cases, they are a function of the trade costs and outputs/expenditures of all their trading partners. Importantly, these terms connected trade between $i$ and $j$ to global trade patterns more broadly. For example, suppose a third country $k$ significantly raised its tariffs and reduced its imports with all countries. This might be expected to increase trade between $i$ and $j$ as they divert some of their previous trade with $k$ towards one another. The MRTs capture this type of relationship via $\tau_{ik}$ in equation (2) and $\tau_{kj}$ in equation (3). Finally, the trade cost term is raised to the elasticity of substitution ($\sigma$), which determines how much trade flows respond to changes in trade costs.

Conveniently, the structural gravity model given by equations 1, 2, and 3 admits a convenient empirical formulation. Taking logs of both sides of equation 1 gives a straight-forward log linear model that can be estimated with relative ease:[1]

$$ln(X_{ijt}) = ln(Y_{it}) + ln(E_{jt}) + ln(Y_t) + (1-\sigma)\left[ln(\tau_{ijt}) - ln(\Pi_{it}) - ln(P_{jt})\right]. \tag{4}$$

In adopting the model to an econometric setting, it becomes necessary to replace several of the theoretical terms with observed proxies. This is especially true for the trade costs ($\tau$, $\Pi$, and $P$), which represent aggregate measures that simplify a potentially long list of specific trade costs into a single value in the theoretical model. One of the primary challenges of the empirical model is finding appropriate proxy variables to adequately capture all important types trade costs. This central question has formed the basis of countless papers exploring different trade cost measures. Some of the most common include geographic distance, contiguous borders, colonial relationships, tariffs, preferential trade agreements (PTAs), language, and EU or WTO membership. Additionally, as discussed below, fixed effects have also been used extensively to attempt to broadly capture these influences. Many of the examples that follow include a standard set of trade cost measures that have been widely in the past. Worth noting is that these variables should not necessarily be considered the "correct" set of variables to use in every situation, which will ultimately depend on the underlying research question.

Finally, there have been many excellent surveys and guides written on gravity modeling in recent years. Many of these sources provide additional details on the topics discussed here. They also present alternative formulations of the theoretical model and additional empirical considerations. Of particular note, interested readers should consider the handbooks of Yotov et al. (2016) and Head and Mayer (2014) as a next step.

[1] As discussed below, most modern approaches estimate a nonlinear version of the model ($X_{ijt} = exp\{ln(\tau_{ijt}) + ...\}$), but the basic concept is the same.

Table 2: Description of variables used in examples

| Variable | Description |
|---|---|
| exp, imp | Exporter and importer identifier, respectively |
| year | Year |
| trade | Bilateral trade flow value |
| sect | Sector identifier |
| ln_trade | Log of trade value |
| pta | Preferential trade agreement indicator |
| eu | Indicator for both being European Union members |
| wto | Indicator for both being WTO members |
| colony | Indicator for colonial ties |
| contiguity | Indicator for shared land border |
| language | Indicator for a shared common language |
| ln_distance | Log population weighted geographic distance |
| ln_gdp_exporter, ln_gdp_importer | Log GDP of exporter and importer, respectively |
| foreign | Indicator for international trade |

# 3 Gravity Estimation

This section presents a collection of empirical approaches for estimating gravity models of trade. The approaches begin in section 3.1 with an early "naive" version of the gravity model that is representative of specifications widely used throughout the literature for many years. Importantly, this specification exhibits several problematic issues and has largely been replaced by improved approaches, which are discussed in sections 3.2–3.8.

For each approach, a Stata code example is provided that demonstrates how the approach can be implemented. The code examples can be run using a sample dataset that is available by request from the author. This dataset contains information on trade between 150 countries for the years 2000, 2003, 2006, 2009, 2012, 2015, and 2018. Table 2 describes the variables referenced in the code examples and found in the sample dataset.[2] These data were sourced from the International Trade and Production Database for Estimation (ITPD-E) of Borchert et al. (2021) (trade flows) and the Dynamic Gravity Dataset of Gurevich and Herman (2018) (covariates). Throughout, the examples take advantage of powerful specialized functions for quickly estimating Ordinary Least Squares (OLS) and Poisson Pseudo Maximum Likelihood (PPML) models with extensive fixed effects: *reghdfe* and *ppmlhdfe*.[3]

## 3.1 "Naive" gravity

The origins of gravity as a model of trade stem from the physics concept of gravity. In physics, gravity between two objects is proportional to each object's size and the distance between them. Beginning with

[2]The dataset is broken into three files that contain aggregate international trade flows, aggregate domestic trade flows, and sector-level international and domestic trade flows.

[3]See Correia (2016) and Correia et al. (2020), respectively

the work of Tinbergen (1962), economists found that this logic performed exceptionally well when applied to international trade. Early empirical gravity models regressed trade values against the GDP values of both countries (measures of market size) and a collection of trade cost variables (measures of the "distance" between markets). The following provides a good example of this type of approach:

**"Naive" gravity**

```
reghdfe ln_trade ln_gdp_exporter ln_gdp_importer pta eu wto colony contiguity \\\
    language ln_distance, cluster(i.exp#i.imp)
```

## 3.2 Controlling for multilateral resistances

For many years, people estimated gravity models using GDP and other country-specific variables to control for the country-level determinants of trade. However, a seminal paper by Anderson and van Wincoop (2003) demonstrated that this seemingly innocuous approach could be highly problematic. Drawing on the earlier theoretical work of Anderson (1979), Anderson and van Wincoop highlighted the importance of MRTs in explaining trade. Importantly, they showed that failure to properly account for multilateral resistances in empirical gravity specifications could bias the model estimates via omitted variables.

While there are multiple ways of incorporating multilateral resistances into empirical gravity frameworks, the most common and easiest to implement is to use country-level fixed effects. As suggested by Feenstra (2002), GDP and other country-level variables can be replaced with exporter-year and importer-year fixed effects. Such a substitution mitigates the omitted variable bias problem as the fixed effects can broadly control for any country-level determinants, whether observed or not.

**Multilateral resistance terms**

```
reghdfe ln_trade pta eu wto colony contiguity language ln_distance, \\\
    absorb(i.exp#i.year i.imp#i.year) cluster(i.exp#i.imp)
```

## 3.3 Poisson psuedo maximum likelihood

In the early days of gravity, models were typically estimated log-linearly via OLS. However, Santos Silva and Tenreyro (2006) demonstrated that OLS presents two problems when applied to trade in this way. First, because log-linear OLS specifications are estimated using logged trade values, they are unable to include observations in which two parties did not trade (i.e. "zero" flows) because the log of zero is undefined. Most studies dropped these observations, using only positive trade values for estimation. Doing so, however, dispenses with potentially valuable information about country pairs for which barriers were too high to trade. The second problem was that gravity models using logged trade values also tend to exhibit heteroskedasticity.[4]

Santos Silva and Tenreyro (2006) proposed a straightforward means of overcoming both problems. The nonlinear Poisson psuedo maximum likelihood (PPML) estimator allows for zeros in the dependent variable and mitigates the heteroskadisticity issues.[5] Today, PPML has become the primary estimator for gravity models.

**Poisson psuedo maximum likelihood**

```
ppmlhdfe trade pta eu wto colony contiguity language ln_distance, \\\
    absorb(i.exp#i.year i.imp#i.year) cluster(i.exp#i.imp)
```

### 3.4 Country-pair fixed effects

Collections of trade cost variables like distance, contiguity, and language have long been used to measure bilateral costs. However, this approach of using specific cost proxies has limitations that could result in problematic biases. These problems are highlighted in the work of Baier and Bergstrand (2007). The first problem is that it may be difficult to fully capture many of the influential factors comprising bilateral trade costs, especially when they are not readily observed (e.g. cultural preferences). The second problem is that many bilateral trade costs of interest may pose endogeneity concerns due to omitted variable and selection bias, simultaneity bias, and measurement error bias. Baier and Bergstrand examine these issues in the context of gravity estimates of trade agreements and highlight the potential for omitted variable bias. In particular, they note that trade agreements are not formed randomly. Instead, countries tend to sign agreements based on many of the same factors that determine trade flows, such as market size, geographic proximity, cultural ties, etc.. Similarly, it may be the case that policy makers pursue trade agreements *because* of large existing trade flows. Ultimately, if unexplained heterogeneity in trade flows is correlated with the determinants of trade agreements, endogeneity biases may be a problem.

Baier and Bergstrand propose an approach to mitigate these concerns for panel data regressions that has become ubiquitous in the literature. The approach includes (time-invariant) country-pair fixed effects in the specification. The country-pair fixed effects control for a much wide range of potential bilateral determinants as they absorb all typical time-invariant determinants, such as distance or borders, as well as

[4] The variance of the error terms is correlated with the dependent variable (i.e. bigger trade flows tend to have higher variance errors).

[5] Santos Silva and Tenreyro also found that PPML performs well compared to other alternative nonlinear estimators.

many determinants that are not observed. Meanwhile, time-variant determinants like trade agreements can still be identified. The explanation they give is that the underlying endogeneity is most likely related to time-invariant, cross sectional characteristics than changes in these factors over time. For example, geographic distance likely influences both trade and agreement formation in the long term. However, yearly *changes* in trade values are much less likely to influence the adoption (or abandoning) of a trade agreement at that time.[6]

The inclusion of country-pair fixed effects alongside importer-year and exporter-year fixed effects has become widespread. Specifications with these three sets of terms are often referred to as "three-way" gravity (with "two-way" gravity referring to just importer- exporter-year effects). Country-pair fixed effects are generally appropriate in most cases. The main exception is when the bilateral variable of interest is not time-varying and therefore collinear with country-pair fixed effects. In these cases, specifications tend to once again use an extensive collection of specific bilateral variables out of necessity instead.

Throughout the literature, pair fixed effects have been included in two ways: symmetric or asymmetric. Asymmetric pair fixed effects are constructed such that the direction of the flow matters, often described instead as exporter-importer fixed effects. That is, the United States (exporter) $\times$ Canada (importer) fixed effect is not the same as the Canada (exporter) $\times$ United States (importer) fixed effect. Alternatively, symmetric fixed effects do not reflect direction and uses a single term for all trade between Canada and the United States. Although both approaches can be effective, asymmetric allows for even greater heterogeneity and should be considered for most applications. Generally, asymmetric fixed effects are a better choice if feasible as they allow for greater trade cost heterogeneity.[7]

**Asymmetric country-pair fixed effects**

```
ppmlhdfe trade pta eu wto, absorb(i.exp#i.year i.imp#i.year i.exp#i.imp) \\\
    cluster(i.exp#i.imp)
```

## 3.5 Domestic trade

Gravity models have most frequently been estimated using only international trade data, due largely to data availability. However, this represents a departure from the theoretical model, which describes both international and domestic (intranational) trade flows. In recent years, there has been a renewed interest in including domestic trade flows in empirical gravity models as doing so offers some significant advantages.

[6] It should be added that this approach does not discount the value of a good instrument.

[7] The inclusion of these terms adds a very large number of variables to even modest models and can become prohibitively difficult to estimate without the use of a high performance estimator like *ppmlhdfe*.

As discussed by Yotov (2022), including domestic trade flows not only strengthens the connection between the empirical and theoretical versions of the gravity model, it also allows for the (improved) identification of many determinants of trade. Domestic trade plays an important role in the model because it is typically a country's largest source of trade and faces the lowest trade costs. When included in the empirical model, it allows international trade costs to be identified relative to domestic trade costs, improving the quality of those estimates. Domestic flows also allow for the identification of non-discriminatory international barriers, such as the effects of most favored nation (MFN) tariffs or many non-tariff measures (Heid et al., 2021).

The inclusion of domestic flows has long been hampered by the limited availability of these data. However, several new trade datasets containing domestic flows have become available in recent years, making it much easier to estimate models using domsetic trade. Some useful sources are the ITPD-E (Borchert et al., 2021), TradeProd (Mayer et al., 2023), and Structural Gravity (Larch et al., 2019) databases.[8]

Specifications using domestic trade flows may need to introduce a new term to the model. Specifically, an indicator term specifying whether the flow is international or domestic should be considered. This term can be defined so that it equals 1 if the flow is international and 0 if domestic, in which case the corresponding estimate will reflect the average costs of foreign trade relative to domestic. Alternatively, it could equal 1 if the flow is domestic and 0 otherwise, in which case the estimate captures the average home bias for trade.

**Domestic trade with foreign trade indicator**

```
ppmlhdfe trade foreign pta eu wto colony contiguity language ln_distance, \\\
    absorb(i.exp#i.year i.imp#i.year)
```

Alternatively, the use of country-pair fixed effects will implicitly provide this control.

**Domestic trade with country-pair fixed effects**

```
ppmlhdfe trade pta eu wto, absorb(i.exp#i.year i.imp#i.year i.exp#i.imp) \\\
    cluster(i.exp#i.imp)
```

## 3.6 Global openness and country-pair trends

As discussed in many of the preceding sections, omitted variable biases stemming from unobserved trade determinants have long been a challenge for gravity estimation. The widespread adoption of three-way fixed

[8] These datasets are available at https://www.usitc.gov/data/gravity/itpde.htm (ITPD-E), http://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=5 (TradeProd), and https://www.wto.org/english/res_e/reser_e/structural_gravity_e.htm (Structural Gravity)

effects has helped mitigate these concerns significantly by controlling broadly for time-varying country-level and bilateral factors. The remaining dimension of variation—time-varying bilateral factors—has generally continued to rely on observable measures of trade frictions/facilitators like PTA, WTO, or EU membership. However, there may still be important unobserved time-varying factors that ought to be controlled for in order to further mitigate omitted variable bias concerns.

Bergstrand et al. (2015) study this issue and suggest a few approaches to help mitigate these concerns. The inclusion of country-pair×year fixed effects is the natural next step from country-pair fixed effects and, in principle, would provide those controls. However, country-pair×year fixed effects typically vary at the level of the data ($ijt$) and are, therefore, generally infeasible. Instead, Bergstrand et al. (2015) propose two alternative fixed effect approaches that vary over time and capture aspects of bilateral characteristics.

The first approach uses a border-year term, which is defined as the interaction between an indicator for foreign (or domestic) trade and year fixed effects. This border-year term is intended to capture general trade openness trends, reflecting variation in bilateral *international* barriers on average rather than for specific country-pairs. To illustrate, Bergstrand et al.'s estimates of the border×year terms generally exhibit a declining effect of international borders over time, suggesting that foreign trade costs had been declining relative to domestic during their sample period (1990–2002). The inclusion of these terms can help identify the effects of specific variables, such as PTAs, and mitigate biases stemming from global trends in trade openness. Importantly, however, this term does require domestic trade flows.

**Border-year fixed effects**

```
ppmlhdfe trade pta eu wto c.foreign#i.year, \\\
    absorb(i.exp#i.year i.imp#i.year i.exp#i.imp)
```

The second approach attempts to feasibly mimic the inclusion of country-pair×year by using country-pair linear trends. The approach allows for heterogeneity over time at the country-pair level but imposes linearity in that variation. However, it does allow for differences in the trends by estimating a separate slope for each country pair. As with the previous approach, the inclusion of linear country-pair trends can help separate the effects of a specific variable of interest from other, often unobserved trends occurring between trading partners, thereby helping to mitigate omitted variable biases.

**Country-pair linear trends**

```
ppmlhdfe trade pta eu wto c.foreign#i.year, \\\
    absorb(i.exp#i.year i.imp#i.year i.exp#i.imp#c.year) cluster(i.exp#i.imp)
```

## 3.7 Three-way bias correction

A drawback of the proliferation of fixed effects in gravity specifications is that their extensiveness may present new empirical issues. There may be concerns that the incorporation of fixed effects could introduce a bias of their own (incidental parameter problems), especially if their is a limited number of years of data ("fixed T"). For two-way gravity, Fernández-Val and Weidner (2016) demonstrate that PPML estimates are unbiased as the number of years and countries becomes large, suggesting this is less of a concern for two-way models. For three-way gravity, however, Weidner and Zylkin (2021) find that PPML point estimates are consistent but that they may be asymptotically biased. To address these asymptotic biases, Weidner and Zylkin recommend an analytical bias correction. Fortunately, they have provided a package in Stata to perform such a correction: *ppml_fe_bias*.

**Three-way gravity bias correction**

```
ppmlhdfe trade pta eu wto, \\\
    absorb(i.exp#i.year i.imp#i.year i.exp#i.imp) cluster(i.exp#i.imp) d
predict lambda
matrix beta = e(b)
ppml_fe_bias trade pta eu wto, i(exp) j(imp) t(year) lambda(lambda) beta(beta)
```

## 3.8 Disaggregated gravity

Gravity models have often been run at many different levels of aggregation. A strength of the theoretical framework is that the model is separable at the sector (or industry) level. Empirically, one might want to derive estimates for different levels of aggregation based on the research question.[9]

There are different options for estimating gravity models using disaggregated, sector-level data (see Borchert et al. (2022) for a more extensive discussion). One option is to estimate a series of gravity models for each sector individually. This will produce a collection of individual estimates for each sector.

[9] There may be important technical considerations as well. Empirically, the choice of aggregation may not be entirely innocuous. For example, as highlighted by Redding and Weinstein (2019), if a gravity models holds at one level of aggregation, it cannot simultaneously hold at a different level. Estimates derived across multiple levels of aggregation may unintentionally reflect composition effects.

**Individual estimation of sectors**

```
foreach s in "Agriculture" "Manufacturing" "MiningEnergy" "Services" {
    ppmlhdfe trade pta eu wto if broad_sector=="`s'", \\\
        absorb(i.exp#i.year i.imp#i.year i.exp#i.imp) cluster(i.exp#i.imp)
}
```

Alternatively, disaggregated sectors can be pooled in a single regression. This approach offers the advantage of being able to exploit all of the underlying variation in the disaggregated data. However, it has the (potentially prohibitive) disadvantage of significantly increasing the computational burden. When sectors are pooled together, covariates can be included as a single term, in which case they will produce an estimate averaged across all included sectors. Alternatively, they can be included as multiple terms based on interactions with sector dummies, in which case they will provide sector-specific estimates. In either case, one should consider interacting the model's country-year (and country-pair) fixed effects with sector dummies as well, which better allows for heterogeneity across sectors.

**Pooled estimation of sector-level data**

```
ppmlhdfe trade pta eu wto, absorb(i.exp#i.year#i.sect i.imp#i.year#i.sect \\\
    i.exp#i.imp#i.sect) cluster(i.exp#i.imp)
```

## 4 Comparison of Estimates

As discussed throughout the previous section, the evolution of gravity models has sought to mitigate numerous biases and reconcile empirical challenges. As a result, we might expect the estimates generated using approach discussed above to yield slightly different estimates. Table 3 presents a side-by-side comparison of each of these specifications.

Table 3: Gravity estimates derived using different empirical specifications

| | Naive<br>(1) | MRTs<br>(2) | PPML<br>(3) | Pair FE<br>(4) | Domestic (Border)<br>(5) | Domestic (Pair FE)<br>(6) | Border-year<br>(7) | Pair trend<br>(8) | Bias correct<br>(9) | Pooled sectors<br>(10) |
|---|---|---|---|---|---|---|---|---|---|---|
| pta | 0.441*** | 0.464*** | 0.249*** | 0.0805*** | 0.0787** | 0.201*** | 0.136*** | 0.138*** | 0.228*** | 0.112*** |
| | (0.0378) | (0.0305) | (0.0534) | (0.0289) | (0.0331) | (0.0453) | (0.0371) | (0.0452) | (0.0671) | (0.0377) |
| eu | 1.541*** | -0.532*** | 0.156* | 0.121*** | 1.078*** | 0.322*** | 0.229*** | 0.222*** | 0.377*** | 0.366*** |
| | (0.0574) | (0.0618) | (0.0804) | (0.0395) | (0.0570) | (0.0473) | (0.0335) | (0.0456) | (0.0712) | (0.0573) |
| wto | 0.784*** | 0.662*** | 0.542*** | 0.120** | 1.099*** | -0.241** | -0.324*** | -0.323*** | -0.272* | 0.244*** |
| | (0.0351) | (0.0850) | (0.154) | (0.0544) | (0.0913) | (0.113) | (0.0841) | (0.114) | (0.155) | (0.0593) |
| colony | 0.840*** | 0.390*** | 0.101 | | 0.524*** | | | | | |
| | (0.0884) | (0.0909) | (0.0636) | | (0.0437) | | | | | |
| contiguity | 0.652*** | 0.842*** | 0.263*** | | 0.668*** | | | | | |
| | (0.0936) | (0.0969) | (0.0580) | | (0.0397) | | | | | |
| language | 0.756*** | 0.815*** | 0.258*** | | 0.231*** | | | | | |
| | (0.0333) | (0.0332) | (0.0521) | | (0.0270) | | | | | |
| ln_distance | -1.170*** | -1.397*** | -0.733*** | | -0.389*** | | | | | |
| | (0.0219) | (0.0213) | (0.0309) | | (0.0203) | | | | | |
| ln_gdp_exporter | 1.313*** | | | | | | | | | |
| | (0.00867) | | | | | | | | | |
| ln_gdp_importer | 1.044*** | | | | | | | | | |
| | (0.00920) | | | | | | | | | |
| foreign | | | | | -3.691*** | | | | | |
| | | | | | (0.0644) | | | | | |
| N | 102729 | 129763 | 150544 | 148251 | 151481 | 149188 | 149188 | 149188 | 149188 | 382291 |
| $R^2$ | 0.615 | 0.752 | 0.941 | 0.992 | 0.983 | 0.998 | 0.998 | 0.998 | 0.998 | 0.997 |
| Estimator | OLS | OLS | PPML | PPML | PPML | PPML | PPML | PPML | PPML | PPML |
| Trade | Log | Log | Level | Level | Level | Level | Level | Level | Level | Level |
| Aggregation | Total | Total | Total | Total | Total | Total | Total | Total | Total | Sector |
| Country-year FE | | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Country-pair FE | | | | ✓ | | ✓ | ✓ | ✓ | ✓ | ✓ |
| Domestic trade | | | | | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ |
| Additional controls | | | | | | | Border-year | Pair trend | | Sector FE interactions |
| Section | 3.1 | 3.2 | 3.3 | 3.4 | 3.5 | 3.5 | 3.6 | 3.6 | 3.7 | 3.8 |

Note: This table contains a series of gravity model specifications. In addition to the listed covariates and fixed effects (FE), several specifications included additional dimensions of controls. Column (7) included border-year effects and Column (8) included linear country-pair trends. Columns (1) and (2) report adjusted $R^2$ values, all others report pseudo $R^2$ values. Standard errors in parentheses. * p¡0.10, ** p¡0.05, *** p¡0.01.

The results demonstrate many of the expected features of each of these models. Moving from "Naive" gravity in column 1 to a specification estimated with MRTs and PPML in column (3) significantly changes the estimates for most variables. For the reasons stated above, the estimates in column (3) are likely better than those in column (1) or (2). Similarly, the inclusion of domestic trade and additional controls in columns (4) through (9) further affects these estimates. Looking at the effect of PTAs, these additions further reduce the estimated impact to between about 0.14 and 0.23 depending on the specification. These estimates also demonstrate that the choice of specification can have non-trivial impacts on the magnitudes of the estimates.

The interpretation of the magnitude of these estimates depends on how they were estimated. For PPML regressions, the effect of a change in a variable on bilateral trade can be computed using the following equation (Yotov et al., 2016).

$$\% \text{ change in } X_{ij} = \left(e^{\hat{\beta}_z * \Delta z} - 1\right) * 100 \tag{5}$$

Here, $\Delta z$ represents a change in the variable of interest and $\hat{\beta}_z$ its estimated coefficient. Using the estimates from column (3), a PTA is estimated to increase bilateral trade by $\left(e^{(0.249*1)} - 1\right) * 100 = 28.3\%$. An increase in the distance between two countries from 4000km to 5000km would decrease trade by $\left(e^{-0.733*ln(5000/4000)} - 1\right) * 100 = -15.1\%$. Importantly, this estimate and interpretation is only a partial estimate of the total effect on trade. It fails to capture the effect that the change in the variable has of multilateral resistances. These secondary effects are typically much smaller in magnitude, meaning that the partial effects are generally good approximations of the impacts of changes in trade costs. None the less, the fact that they are only partial effects should not be overlooked.[10]

# 5 Data Resources

Armed with the techniques described in section 3, the only remaining component is the data. This section provides a brief discussion of some prominent sources of bilateral trade and covariate data that have been used extensively. These sources are curated databases that are specially suited for gravity estimation. Table 4 lists a variety of trade data sources as well as information on their coverage in terms of countries, years, industries/products, and domestic trade.

While the covariates used to explain trade differ across specifications and research questions, there are numerous databases containing a variety of traditional variables that have been long shown to significantly explain large portions of bilateral trade costs. Prominent sources for many of these variables are:

[10] Extensions of the gravity framework make it possible to estimate the full "general equilibrium" effects of changes in trade costs. See Herman (2021) for an introduction to these approaches.

Table 4: Sources of curated bilateral trade data

| Name | Citation | Countries | Years | Coverage | Domestic flows |
|---|---|---|---|---|---|
| BACI | Gaulier and Zignago (2010) | 200 | 1995–2021 | 5000 products | |
| ITPD-E | Borchert et al. (2021) | 265 | 1986–2019 | 170 industries | Yes |
| Structural Gravity | Larch et al. (2019) | 186 | 1980–2016 | 1 sector | Yes |
| TRADHIST | Fouquin and Hugot (2016) | 319 | 1827–2014 | Total | |
| TradeProd | Mayer et al. (2023) | 162 | 1966–2018 | 9 industries | Yes |

- **Dynamic Gravity Dataset:** A collection of standard gravity covariates, including country-level macroeconomic indicators, geographic characteristics, preferential trade agreements, and cultural proxies (Gurevich and Herman, 2018).[11]
- **CEPII Gravity Database:** A similar collection of standard gravity covariates (Conte et al., 2022).[12]
- **DESTA**: The DESTA (Design of Trade Agreements Database) is a specialized dataset describing the content of trade agreements since 1945 across multiple dimensions (Dür et al., 2014).[13]
- **Deep Trade Agreements Database**: Like DESTA, this database provides an accounting of the content of trade agreements with an especially high level of granularity (Mattoo et al., 2020).[14]
- **DICL**: The DICL (Domestic and International Languages Database) contains multiple bilateral measures of linguistic similarity between and within countries (Gurevich et al., 2021).[15]

# 6 Conclusion

The gravity model of international trade has proven to be an exceptionally successful tool for empirically studying international—and domestic—trade. However, deriving high quality estimates requires careful consideration of the econometric specification being used. Common applications of the model are prone to a variety of biases and identification challenges that should not be overlooked. Fortunately, there exist readily implemented approaches for mitigating these issues, as detailed throughout this paper. Most notably, gravity specifications should include controls for multilateral resistances and be estimated using Poisson Pseudo Maximum Likelihood. In many cases, additional features such as pair fixed effects, domestic trade flows, trends, and bias correction should be considered as well. Practitioners armed with these methods can be more confident in their analyses and the insights that they produce.

[11] https://www.usitc.gov/data/gravity/dgd.htm
[12] http://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=8.
[13] https://www.designoftradeagreements.org/.
[14] https://datatopics.worldbank.org/dta/overview.html#
[15] https://www.usitc.gov/data/gravity/dicl.htm.

# References


Anderson, J. E. (1979). A theoretical foundation for the gravity equation. *The American Economic Review 69*(1), 106–116. https://www.jstor.org/stable/1802501.

Anderson, J. E. and E. van Wincoop (2003). Gravity with gravitas: A solution to the border problem. *American Economic Review 93*, 170–192. https://doi.org/10.1257/000282803321455214.

Arkolakis, C., A. Costinot, and A. Rodríguez-Clare (2012). New trade models, same old gains? *American Economic Review 102*(1), 94–130. https://doi.org/10.1257/aer.102.1.94.

Baier, S. L. and J. H. Bergstrand (2007). Do trade agreements actually increase members' international trade? *Journal of International Economics 71*(1), 72–95. https://doi.org/10.1016/j.jinteco.2006.02.005.

Bergstrand, J. H., M. Larch, and Y. V. Yotov (2015). Economic integration agreements, border effects, and distance elasticities in the gravity equation. *European Economic Review 78*, 307–327. https://doi.org/10.1016/j.euroecorev.2015.06.003.

Borchert, I., M. Larch, S. Shikher, and Y. V. Yotov (2021). The international trade and production database for estimation (ITPD-E). *International Economics 166*, 140–166. https://doi.org/10.1016/j.inteco.2020.08.001.

Borchert, I., M. Larch, S. Shikher, and Y. V. Yotov (2022). Disaggregated gravity: Benchmark estimates and stylized facts from a new database. *Review of International Economics 30*, 113–136. https://doi.org/10.1111/roie.12555.

Conte, M., P. Cotterlaz, and T. Mayer (2022). The CEPII gravity database. CEPII Working Paper 2022-05. http://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=8.

Correia, S. (2016). Linear models with high-dimensional fixed effects: An efficient and feasible estimator. Working paper. http://scorreia.com/research/hdfe.pdf.

Correia, S., P. Guimarães, and T. Zylkin (2020). Fast Poisson estimation with high-dimensional fixed effects. *The Stata Journal 20*(1), 95–115.

Dür, A., L. Baccini, and M. Elsig (2014). The design of international trade agreements: Introducing a new database. *The Review of International Organizations 9*(3), 353–375. https://doi.org/10.1007/s11558-013-9179-8.

Feenstra, R. C. (2002). Border effects and the gravity equation: Consistent methods for estimation. *Scottish Journal of Political Economy 49*(5), 491–506. https://doi.org/10.1111/1467-9485.00244.

Fernández-Val, I. and M. Weidner (2016). Individual and time effects in nonlinear panel models with large n, t. *Journal of Econometrics 192*(1), 291–312. https://doi.org/10.1016/j.jeconom.2015.12.014.

Fouquin, M. and J. Hugot (2016). Two centuries of bilateral trade and gravity data: 1827-2014. CEPII Working Paper 2016-14. http://www.cepii.fr/CEPII/fr/publications/wp/abstract.asp?NoDoc=9134.

Gaulier, G. and S. Zignago (2010). BACI: International trade database at the product-level. the 1994-2007 version. CEPII Working Paper 2010-23. http://www.cepii.fr/CEPII/fr/publications/wp/abstract.asp?NoDoc=2726.

Gurevich, T. and P. R. Herman (2018, February). The dynamic gravity dataset: 1948-2016. Office of Economics Working Paper 2018-02-A, U.S. International Trade Commission. https://www.usitc.gov/publications/332/working_papers/gurevich_herman_2018_dynamic_gravity_dataset_201802a.pdf.

Gurevich, T., P. R. Herman, F. Toubal, and Y. V. Yotov (2021, January). One nation, one language? domestic language diversity, trade and welfare. Office of Economics Working Paper 2021-01-B, U.S. International Trade Commission. https://www.usitc.gov/publications/332/working_papers/ghty_jan2021_one_nation_one_language_usitc.pdf.

Head, K. and T. Mayer (2014). Gravity equations: Workhorse, toolkit, and cookbook. In G. Gopinath, E. Helpman, and K. Rogoff (Eds.), *Handbook of International Economics*, Volume 4, pp. 131–195. Elsevier. https://doi.org/10.1016/B978-0-444-54314-1.00003-3.

Heid, B., M. Larch, and Y. V. Yotov (2021). Estimating the effects of non-discriminatory trade policies within structural gravity models. *Canadian Journal of Economics 54*(1), 376–409. https://doi.org/10.1111/caje.12493.

Herman, P. R. (2021, April). gegravity: General equilibrium gravity modeling in python. Office of Economics Working Paper 2021-04-B, U.S. International Trade Commission. https://www.usitc.gov/publications/332/working_papers/herman_2021_gegravity_modeling_in_python.pdf.

Larch, M., J.-A. Montiero, R. Piermartini, and Y. Yotov (2019). On the effects of GATT/WTO membership on trade: They are positive and large after all. WTO Staff Working Papers ERSD-2019-09, World Trade Organization (WTO). https://www.wto.org/english/res_e/reser_e/ersd201909_e.htm.

Mattoo, A., N. Rocha, and M. Ruta (2020). *Handbook of Deep Trade Agreements*. Washington, DC: World Bank.

Mayer, T., G. Santoni, and V. Vicard (2023). The CEPII trade and production database. CEPII Working Paper 2023-01. http://www.cepii.fr/CEPII/en/bdd_modele/bdd_modele_item.asp?id=5.

Redding, S. J. and D. E. Weinstein (2019, May). Aggregation and the gravity equation. *AEA Papers and Proceedings 109*, 450–455. https://doi.org/10.1257/pandp.20191006.

Santos Silva, J. M. C. and S. Tenreyro (2006). The log of gravity. *The Review of Economics and Statistics 88*(4), 641–658. https://doi.org/10.1162/rest.88.4.641.

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

Weidner, M. and T. Zylkin (2021). Bias and consistency in three-way gravity models. *Journal of International Economics 132*, 103513. https://doi.org/10.1016/j.jinteco.2021.103513.

Yotov, Y. V. (2022). On the role of domestic trade flows for estimating the gravity model of trade. *Contemporary Economic Policy 40*(3), 526–540. https://doi.org/10.1111/coep.12567.

Yotov, Y. V., R. Piermartini, J.-A. Monteiro, and M. Larch (2016). *An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model (Online Revised Version)*. World Trade Organization and the United Nations Conference on Trade and Development. https://doi.org/10.30875/abc0167e-en.