Gravity for Undergrads

Yoto V. Yotov

DREXEL UNIVERSITY
Center for
Global Policy Analysis
LeBow College of Business

The CGPA working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the CGPA Board. The CGPA does not hold any responsibility for the content and correctness of the papers in this series.

To the Instructor

Thank you for using this chapter for your course! I have devoted significant time and effort to write the chapter, and I hope that it is useful. The chapter, the data and codes for the hands-on analysis, and the practice questions and their solutions are free public goods. However, please consider the following:

- Ask your students to download the chapter, along with the data and codes, directly from this web site https://yotoyotov.com/Gravity_Undergrads.html. This will help me track its impact and determine whether continuing this effort is worth it.
- Do let me know if you have any concerns or detect inconsistencies, typos, and errors with the chapter or with the data and the codes that come with it. Suggestions for improvements to make the chapter more useful are always very welcome at yotov@drexel.edu.
- E-mail me at yotov@drexel.edu if you want to obtain the solutions to the practice problems that are included at the end of the chapter.
- Finally, encourage your students to cite the chapter properly if they use it for research projects.

About the Author

This chapter combines my two greatest professional passions – my love for teaching and my dedication to the gravity model of trade. As such, I consider this chapter one of my most valuable contributions and hope that it is useful to many undergraduate students, master's students, and others who may benefit from an introduction to the gravity model of trade.

I decided to pursue a Ph.D. because of teaching and gravity, and I have devoted my career to the gravity model of trade. I have made many contributions to gravity, including the development of theoretical models, estimation techniques, computation methods, and the construction of gravity data. My gravity work has been published in respected academic journals, and I have used the gravity model to advise and consult many international organizations, governments, and think tanks.

Perhaps most important for the current purposes, I have taught gravity many times, to hundreds of students, researchers, and policy makers from more than 130 countries. I also have supervised gravity research by undergraduate students, which convinced me that the gravity model is indeed accessible and can be useful to undergrads. I view this chapter as a unique opportunity to bring cutting-edge methods and practical policy tools to the undergraduate classroom, and I consider writing it a privilege.

© 2025, Y. V. Yotov. All rights reserved.

Acknowledgments and Disclaimer

I am very grateful to many colleagues and friends who have contributed directly and indirectly to this chapter, including Delina Agnosteva, James Anderson, Scott Baier, Richard Baldwin, Richard Barnett, Jeff Bergstrand, Cosimo Beverelli, Ingo Borchert, Sebastien Bradley, Federico Carril-Caccia, Carsten Eckel, Peter Egger, Gabriel Felbermayr, Lisandra Flach, Javier Florez, Lionel Fontagne, Rebecca Freeman, Gene Grossman, Jean Grossman, James Harrigan, Beno Heid, Inga Heiland, Julian Hinz, Peter Herman, Paul Ko, Ohyun Kwon, Mario Larch, Maia Linask, Jeff Luckstead, Martina Magli, Inma Martinez-Zarzoso, Xenia Matschke, Jose Antonio Monteiro, Arne Nagengast, Sergei Nigai, Kevin O'Rourke, Jordi Paniagua, Blake Peters, Roberta Piermartini, Ray Riezman, Fernando Rios-Avila, Ana Maria Santacreu, Mauricio Sepulveda, Serge Shikher, João Santos Silva, Bob Staiger, Costas Syropoulos, Angelos Theodorakopoulos, Farid Toubal, Joschka Wanner, Erdal Yalcin, Jangsu Yoon, and Tom Zylkin. Very special thanks to Blaize Giangiulio for providing invaluable assistance with all aspects of the chapter. I acknowledge support and encouragement from the Center for Global Policy Analysis at Drexel University. Consistent with most undergraduate textbooks, I have kept the number of references low. However, throughout the chapter, I refer the reader to some key papers and surveys. I take full responsibility for all statements and errors.

Gravity for Undergrads

The gravity equation is the "workhorse" model of international trade and the most popular tool for trade policy analysis. For example, all quantitative analysis of U.S. President Donald Trump's tariffs in early 2025, the effects of BREXIT, and the impact of the trade sanctions on Russia following its invasion of Ukraine, which were all widely covered in the popular media, were obtained from versions of the gravity model. Thus, it is not surprising that the gravity equation is perhaps the only empirical model that has been featured on the front page of the Financial Times (see Figure 1).

Figure 1: Gravity in the Financial Times

FINANCIAL TIMES

TUESDAY 19 APRIL 2016 · WORLD BUSINESS NEWSPAPER · UK £2.70 Channel Islands £3.00; Republic of Ireland €3.00

Treasury's grim forecasts spark fury from Tory Brexit rebels

...untry has brought the deep split on ...rope within the Conservative party to the heart of government.
Justice secretary Michael Gove, the ...binet's leading Leave campaigner, will ...day accuse the pro-EU campaign of ...eating voters "like children" in a furi...s response to Mr Osborne's declara...on that pro-Brexit advocates were ...conomically illiterate".
In the most significant day of the EU ...ferendum campaign so far, Mr ...borne sought to destroy the economic ...guments for Brexit in a 200-page ...easury analysis of the impact of a ...ave vote.
This week is seen in Number 10 as ...ucial, with Barack Obama arriving in ...ondon later in the week with an ...pected message that the US would ...efer the UK to stay in the EU.
The Treasury paper's main scenario ...aimed that an exit would over the ...edium term lead to falling trade and ...t a £36bn annual dent in the public ...nances, equivalent to 8p on the basic ...te of income tax.
Over 15 years the economy w... be 2 per cent smaller than w... have ...en the case, costing each h...hold ...4,300 a year. The paper ...nounced by senior Tory MP Bernard ...nkin as "highly tendentious".

-9.5%

Negotiated bilateral agreement (like Canada)

-6.2%

Membership of EEA (like Norway)

-3.8%

...and the key equation

The Treasury's formula for calculating trade ... countries

$$\ln(T_{ijt}) = \alpha_{ij} + \gamma_t + \alpha_1 \ln(Y_{it} * Y_{jt}) + \alpha_2 \ln(POP_{it} * POP_{jt}) + \alpha_3 \ln(DIST_{ij}) + \alpha_4 COMLANG_{ij} + \alpha_5 COLONY_{ij} + \alpha_6 BORDER_{ij} + \varepsilon_{ijt}$$

Briefing

Big companies look outside for a leader
Big UK companies are breaking with global peers and hiring more chiefs from outside than ever. Last year, 58 per cent came from another group, versus 23 per cent globally.— PAGE 17; LOMBARD, PAGE 22

King quits Villa board after two months
Former governor of the Bank of England Mervyn King has resigned from the board of Premier League club Aston Villa — just two months after joining and two days after the club was relegated.— PAGE 17

Crude dips but recovers after talks fail
Brent was roughly flat after the failure of Doha talks aimed at an output freeze on Sunday but a serious sell-off was avoided. Meanwhile, Riyadh policy has a new voice.— PAGE 5; LEX, PAGE 14; MARKETS, PAGES 30-32

Buenos Aires ends 15-year debt exile
Argentina has secured one of the most anticipated comebacks, ending 15 years in exile with a debt sale that attracted bids of more than $50bn. Early signs put 10-year debt yield at under 8 per cent.— PAGE 10

Israeli soldier charged over killing
An Israeli soldier has been charged with manslaughter over the shooting of an unarmed Palestinian assailant as he lay on the ground, the first case of its kind in more than a decade.— PAGE 6

Datawatch

Counterfeit and pirated goods
% of seized goods' value, 2011-2013

Source: The Financial Times, April 19, 2016. Inspired by Peter Neary's 2019 RES presidential address.

The main reasons for the celebrity status of the gravity equation are that: (i) it has unmatched predictive power; (ii) it has solid theoretical foundations; (iii) it is very intuitive and accessible to diverse audiences, including undergraduate students; and (iv) it is straightforward to implement empirically to study many policy applications. Yet, despite its theoretical foundations, remarkable empirical success, intuitive appeal, and ease of implementation, the gravity equation has not received proper and well-deserved coverage in undergraduate trade textbooks. To fill this gap, the main objectives of this chapter are to:

- Introduce the gravity model of trade to the undergraduate student. Discuss why and how gravity can be beneficial to students. Describe the properties that have made the gravity model so popular and successful. This is achieved in Section 1.
- Provide theoretical motivation for the naive gravity equation. Transform naive gravity into a structural gravity model. Highlight the important implications of gravity theory for gravity applications and estimations. This is achieved in Section 2.
- Translate the theoretical gravity equation into an econometric model. Discuss the main econometric challenges with gravity estimations. Synthesize several intuitive and easy to implement recommendations for gravity estimations. This is achieved in Section 3.
- Offer hands-on analysis with real data and simple codes for practical applications. Provide practice questions, guidance, and additional data, which may be useful for empirical term projects, in econometrics and seminar classes, and for writing an undergraduate thesis. This is achieved in Section 4 and a series of practice questions that are included at the end of the chapter.

1. The Gravity Model. Why Study It? Why So Popular?

In response to inquiries by fellow professors, based on feedback from undergraduate students who have been exposed to the gravity model of trade, and consistent with my own teaching and consulting experience, I firmly believe that gravity can and should be taught to undergraduate students. Box 1 provides a list of reasons why the gravity model should be covered in undergraduate courses and how and why it can benefit undergraduate students.

Box 1. Why Study the Gravity Model & How Can It be Useful to Undergraduates?

- The gravity model of trade is more successful at predicting trade flows and is more widely used in both academic and policy work than any of the standard models that are covered in the undergraduate international trade textbooks.
- The gravity model can be used to complement and reinforce the classic trade theories, especially because the empirical gravity equation enables students to empirically test some of these theories.
- The gravity model has clear practical uses and many applications. It is an empirical model that treats trade costs seriously and realistically, and it can be easily implemented with real data, thus enriching the theory-heavy undergraduate trade courses.
- Because of its intuitive appeal and simple theoretical representation, the gravity model is accessible and easy to understand to wide audiences, including undergraduate students.
- Due to recent econometric and computational developments, the implementation of a cutting-edge empirical gravity model is possible with very basic econometrics knowledge (e.g., using ordinary least squares (OLS) and fixed effects). To facilitate this, the chapter includes real data, econometric codes, hands-on analysis, and practice policy questions.
- The gravity model can benefit students beyond trade. Specifically, the intuition and the empirical tools from this chapter can be applied directly to bilateral migration flows, foreign direct investment (FDI), cross-border patents, etc.
- In addition to being appropriate for an undergraduate international trade course, the gravity model can be useful in econometrics courses, undergraduate seminar courses, and to students who want to write an undergraduate thesis or work on research projects.

The Gravity Model of Trade. The gravity equation of international trade predicts that the bilateral trade flows (X_{ij}) between two countries, i.e., the imports of country j from country i, should be proportional to the product of the sizes of the exporter (Y_i) and the importer (Y_j), and inversely proportional to the bilateral trade frictions (T_{ij}) between the two countries:

$$X_{ij} = \tilde{G}\frac{Y_i Y_j}{T_{ij}^{\theta}}, \tag{1}$$

where $\tilde{G}$ is the trade gravitational constant and θ is the elasticity of trade flows with respect to trade frictions, which captures the responsiveness of trade flows to changes in the trade frictions. Intuitively, Equation (1) implies that the larger and the closer two countries are, the more they will trade with each other. The simplicity and intuitive appeal of the gravity model of trade is one of its most attractive features.

Why Is Gravity So Popular? There are four main reasons why the gravity equation of trade is so popular among economists and policy practitioners. First, as discussed earlier, through its analogy with Newtonian gravity, the gravity model of international trade is very intuitive.

1. The gravity equation is very intuitive.

The remarkable resemblance between the trade gravity equation and Newton's law of universal gravitation is captured in Box 2, which reveals that trade (the gravitational force) between two countries (objects) is directly proportional to the product of their sizes (masses) and inversely proportional to the trade frictions (the square of distance) between them. Put simply, the larger and the closer two countries are, the more they will trade with each other.

Box 2. Gravity in Physics vs. Gravity in Trade

Gravity in Physics

$$F_{ij} = G\frac{M_i M_j}{D_{ij}^2}$$

where:

F_{ij} is the gravitational force between objects i and j;

G is the gravitational constant in physics;

M_i & M_j are the masses of objects i and j;

D_{ij} is the distance between i and j;

2 is the elasticity of the gravitational force with respect to distance.

Gravity in Trade

$$X_{ij} = \tilde{G}\frac{Y_i Y_j}{T_{ij}^{\theta}}$$

where:

X_{ij} is the value of trade flows between countries i and j;

$\tilde{G}$ is the gravitational constant in trade;

Y_i & Y_j are the economic sizes of countries i and j;

T_{ij} denotes the trade costs/frictions between i and j;

$\theta > 0$ is the elasticity of trade flows with respect to trade costs.

What makes this analogy, and the striking similarity between the gravity equations in trade vs. physics, even more impressive is that the trade gravity equation can be derived from solid microeconomic theories.

The second main reason why the gravity equation is so popular is that it works. The empirical gravity model predicts bilateral trade flows remarkably well.

2. The gravity equation has tremendous predictive power.

The gravity equation consistently delivers a very strong fit between the actual and predicted trade data and plausible (in terms of magnitudes and signs) estimates on a number of independent variables known as the '*standard gravity variables*', including the logarithm of bilateral distance, and indicator variables for contiguous borders, common official language, colonial relationships, and regional trade agreements (RTAs). For example, bilateral trade decreases with distance but increases when RTAs are formed.

Figure 2 visualizes the great performance of the gravity model by plotting the aggregate bilateral trade flows that are predicted from gravity against the corresponding actual trade flows for the 100 largest exporters in the world over the period 1990-2023. The predictions in Figure 2 are from a modern econometric gravity model, which will be implemented in Section 4 of this chapter, and they show that the empirical performance of the modern gravity model is indeed unprecedented (e.g., the correlation between actual and predicted trade is 0.99). However, as will also be demonstrated in Section 4 and in the practice questions that are included at the end of the chapter, even the naive gravity model with only three independent variables predicts bilateral trade flows quite successfully.

The predictions in Figure 2 are based on data that varies over the period 1990-2023. Importantly, the gravity equation works remarkably well with data for a single year (i.e., cross-section data) and with data pooled over years (i.e., panel data). Moreover, the gravity equation works very well with aggregate and with disaggregated data across any level of aggregation, e.g., product, industry, sector, etc.

Figure 2: The Gravity Model Works

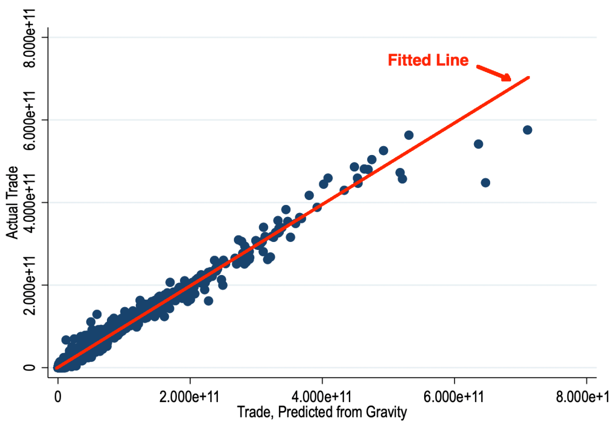

Source: The author. From the analysis in Section 4.

The third main reason for the celebrity status of the gravity equation is that it is very flexible.

3. The gravity equation is very flexible.

The gravity model has been used to explain trade flows and to quantify the impact of their determinants in hundreds of academic papers, and it is the go-to model for trade policy analysis, too. Box 3 includes a detailed (but still far from exhaustive) list of *'traditional'* vs. *'more exotic'* determinants of trade flows. It is safe to conclude that, to study the impact of any determinant of trade flows or other economic outcomes via trade flows, one inevitably should resort to some version of the gravity model.

Box 3. Applications of the Gravity Equation

'Traditional' trade determinants. Distance, Contiguity, GDP, Population, Development, Preferential Trade Agreements, Tariffs, Tariff Wars, Export Subsidies, Geography, Non-tariff Measures, World Trade Organization Membership, Customs Unions, the European Union, Common Currency and Currency Unions, OECD Membership, IMF Membership, Foreign Direct Investment, Immigration, Cultural Ties, Colonial Relationships, Common and Shared Language, ... etc.

'More Exotic' trade determinants. Institutional Quality, Foreign Aid, Trust, Reputation for People, Reputation for Products, Exchange Rates, Covid, Brexit, Export Promotion, Patents, Technical Barriers to Trade, Sanitary and Phytosanitary Standards, Corporate Income Taxes, Value Added Taxes, Mega Sporting Events (Olympic Games and World Cup), Embargoes and Sanctions (e.g., on Russia), Conflict and Wars, Piracy, Ice Cap Melting, The Closures of the Suez Canal, Trump's Tariff Wars ... etc.

For specific references to most of the applications above, I refer the interested reader to Yotov (2024).

In addition to its ability to accommodate any determinant of trade, four other flexibility dimensions of the gravity model arise. First, gravity applies at any level of aggregation, i.e., for products, industries, sectors, etc. Second, despite less developed theoretical foundations as compared to trade, gravity has been applied successfully to study other bilateral flows, including FDI, migration, international technology transfers, etc. Third, as demonstrated later in this chapter, it is very easy to implement a modern empirical gravity model with standard statistical software. Finally, the gravity model can be flexibly nested within broader models, including investment in physical capital, environmental models, labor markets, etc.

The fourth main reason for the success of the gravity model is that it can be derived based on solid microeconomic theories.

4. The gravity equation has solid (and many) theoretical foundations.

One of the most remarkable features of the gravity equation of trade is that, under relatively standard assumptions, it can be obtained from many alternative microeconomic foundations. As visualized in Figure 3, the gravity equation can be derived from the classic trade theories that are routinely covered in undergraduate textbooks, including the Ricardian model, the Heckscher-Ohlin model, the Monopolistic Competition model, and the Specific Factors model. In addition, the same gravity equation emerges from sectoral trade models, trade models with dynamics, trade models with input-output linkages, and trade models with heterogeneous firms.

Figure 3: Some Theoretical Foundations of the Gravity Model of Trade

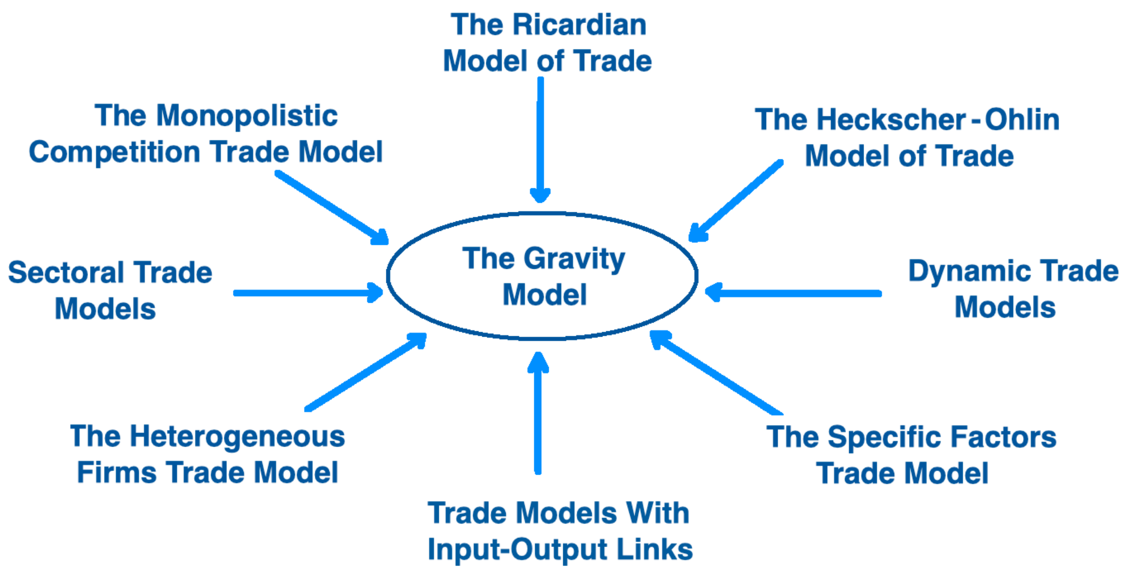

Source: The author. Adapted from Yotov et al. (2016).

The trade profession took a long time to realize that so many different trade theories converge to the same simple and intuitive gravity equation (see Box 4). However, in retrospect, perhaps it should not be surprising that the drivers of trade can be summarized in three intuitive terms: forces on the exporter side, forces on the importer side, and bilateral forces. The same intuition applies to any socioeconomic bilateral relationship. Accordingly, the empirical methods and hands-on analysis that will be presented in Sections 3 and 4 should apply more broadly, e.g., to FDI, migration, technology flows, etc.

The theoretical foundations of the gravity model have many important implications for empirical analysis with the gravity equation. For example, closer adherence to theory leads to improved overall predictive power of the gravity model and to more accurate estimates of the effects of trade policies. On a related note, accounting for some of the key theoretical features of gravity has led to solutions to several prominent puzzles in the trade literature, e.g., the counterintuitive finding that the effects of distance on trade have not fallen over time or that the gravity model cannot properly account for bilateral trade imbalances, among others. Gravity theory has motivated some of the most important recommendations for estimating gravity models, e.g., to account for multilateral trade costs, which are discussed in detail in the next section. Finally, gravity theory allows for translating the effects of trade to other economic outcomes, e.g., labor market effects, environmental impact, investment in physical capital, economic growth, etc.

Before continuing, we note that, similar to many other great inventions, the gravity model did not become a celebrity instantaneously. Instead, as summarized in Box 4, it took many years for the gravity model to win the hearts of the trade economists and to become the workhorse model of trade.

Box 4. The Story of Gravity

In 2025, the gravity model rules the field of international trade! However, its road to fame was not an easy one. The first applications of gravity to economics were a-theoretical and by analogy with physics. Many credit Tinbergen (1962) as the first application of gravity to economics. This may be true, but only from a trade perspective, as the 1969 Nobel Laureate was (arguably) the first to apply gravity to international trade flows. Long before Tinbergen, Ravenstein (1885) applied the gravity equation to migration flows. Some of Tinbergen's students continued his gravity work in the 1960s and 1970s, however, most economists of this period did not consider gravity to be a serious trade model due to lack of theoretical foundations.

While there has been some debate and discussion about who was the first to apply gravity to economics, trade economists seem to be in unanimous agreement that the first theoretical foundation of the gravity equation of trade, as we know it today, belongs to Anderson (1979). In fact, it is truly remarkable that, subject to some 'cosmetic' theoretical improvements, Anderson's 1979 gravity model is perfectly consistent with all modern empirical gravity analysis.

Despite its intuitive appeal, good empirical performance, and already solid theoretical foundation, the gravity model struggled to win the hearts of mainstream trade economists during the 1980s, and even 1990s. Several influential surveys from the *Handbook of International Economics* – the compass for trade research – were not kind to gravity. They questioned its theoretical heritage and predicted that it would have no effect on the subject of international trade. In retrospect, Anderson (2011) characterized the experience of the gravity model during this period as *'an intellectual orphan'*.

Until the early 2000s, trade economists continued to ignore gravity due to its lackluster reputation. During this time, it was used primarily for policy analysis. However, several major advances in the early 2000s led to the golden age of 'Structural Gravity' (2002-2012). First, and most important, the seminal theoretical contributions of Eaton and Kortum (2002) and Anderson and van Wincoop (2003) left no doubt about the fact that the gravity model had very solid theoretical foundations. Second, empirical contributions by influential economists, e.g., Frankel and Romer (1999) and Rose (2000), broke the stigma of using gravity for serious work and thus started the resurrection of gravity as a respectable empirical tool. Third, this period witnessed the construction of large and high-quality trade data and better econometric methods and capabilities (Baldwin and Taglioni, 2006).

During its golden age (2002-2012), gravity established itself as the workhorse model in trade and appeared in hundreds of publications, including applications, theoretical developments, contributions to estimation, and new datasets. Most gravity papers during this period were empirical applications aiming to estimate the impact of various policies and determinants of bilateral trade flows (e.g., free trade agreements (FTAs), membership to the World Trade Organization (WTO), distance, colonial relationships, etc.). While many still applied gravity only intuitively, closer adherence to theory led to better understanding of the drivers of the growth in trade flows and more plausible gravity estimates.

The 'Golden Age of Gravity' also witnessed significant developments on the theory front, e.g., sectoral gravity, gravity with heterogeneous firms, and dynamic gravity. The increased interest in gravity theory and applications was facilitated by significant improvements in computing power, and was accompanied by new contributions on the estimation front, e.g., the use of exporter, importer, and country-pair fixed effects, which will be discussed in the next section, and the introduction of the Poisson Pseudo Maximum Likelihood (PPML) estimator by Santos Silva and Tenreyro (2006), which later firmly established itself as the best gravity estimator.

In a seminal paper, Arkolakis et al. (2012) cemented the structural gravity model's hegemony in trade by demonstrating that different micro-theoretical foundations converge to exactly the same gravity equation. Along with many new applications and theoretical developments, the gravity model got its revenge with prominent coverage in the 2014 edition of the *Handbook of International Economics*, the same publication where it was dismissed in the 1980s and 1990s. The gravity model was also featured in the 2018 *Handbook of International Trade and Transportation* and in dedicated books for trade policy analysis and the impact of globalization.

The empirical gravity equation remained the go-to model for new applications, while others revisited existing results with new, better methods to measure trade costs and to quantify the impact of various determinants of trade flows (e.g., currency unions, piracy, common language, exchange rates, economic sanctions, etc.). Better data (in terms of country, sector, and time coverage) allowed for exploring the heterogeneous effects of many trade policies (e.g., RTAs, currency unions, sanctions, etc.) across various dimensions. It was also established that gravity works quite well for services, mining, and agricultural trade, and at any level of aggregation, i.e., from the product level to the aggregate level. It became quite clear that, in order to study the impact of any determinant of trade flows, one must rely on some version of the gravity model. Moreover, capitalizing on the advances in the trade literature, scholars and policy makers adapted the trade gravity model to study other bilateral flows, e.g., migration, FDI, cross-border patents, etc.

This period also witnessed major estimation, data, and theory contributions. On the estimation front, it was established that the PPML estimator was perfectly consistent with gravity theory, which in turn allowed for performing comprehensive trade policy analysis in standard statistical software and without the need of any custom programming. There were also important advances on the theory front (e.g., gravity with input-output linkages and gravity with bilateral dynamics). In combination with its remarkable empirical performance and new econometric and computational tools, the advances in gravity theory led to a series of contributions that linked trade to various economic outcomes, e.g., technology diffusion, unemployment, carbon emissions, etc. A new generation of gravity databases, covering explanatory gravity variables as well as international and domestic trade flows at various levels of aggregation, were built to support the theoretical advances and the new application needs.

The unprecedented power and success of the gravity model manifested itself in 2025, in response to the frequent tariff changes (e.g., the steel tariffs on Canada were imposed in the morning on March 9 and lifted in the afternoon on the same day) of the U.S. President Trump. Thanks to the great advances in the gravity literature, we witnessed something truly remarkable – we could analyze the full (partial and general equilibrium) impact of trade policies in real time.

This summary of the evolution of the gravity model is adapted from Yotov (2024), and I refer the interested reader to this paper for a more detailed discussion and references.

2. Gravity with Gravitas: The Structural Gravity Model

This section has three main learning objectives. First, it will introduce the theoretical/structural gravity model of trade. Second, it will compare the naive and structural gravity equations while highlighting the key theory-driven differences between the two. Third, it will draw several important implications for modeling trade flows and estimating empirical gravity equations in a theory-consistent manner.

2.1. All Roads Lead to . . . Structural Gravity

As discussed earlier, one of the most remarkable features of the gravity model of trade is that the same gravity equation is representative of and can be derived from many alternative theoretical microeconomic foundations, some of which are standard topics in undergraduate international trade textbooks (see Figure 3). Following the evolution of the theoretical gravity literature, we will introduce the contemporary structural gravity model in two steps. First, consistent with the classic theories of trade, we present the cross-section structural gravity equation:[1]

$$X_{ij} = \frac{Y_i E_j}{Y}\left(\frac{t_{ij}}{T_j T_i}\right)^{-\theta}. \qquad (2)$$

Three theory-driven features distinguish the structural gravity Equation (2) from the naive gravity equation shown earlier in Equation (1). First, the structural term behind the gravitational constant $\tilde{G}$ is the value of world output Y. The implication is that bilateral trade flows are proportional to the product of the economic sizes of the two trading partners as a share of the corresponding world output.

Second, the theoretical gravity equation distinguishes between the value of output (Y_i), as the size measure on the exporter side, and expenditure (E_j), as the measure of size on the importer side. Intuitively, the relevant size measure for the exporter should be its production capacity, while the relevant measure of the size of the importer should be its consumption capacity. The importance of distinguishing between output and expenditure as the appropriate measures of size in the gravity model is even more pronounced at the disaggregated level, where, e.g., due to specialization, the sectoral trade imbalances (i.e., the gap between national exports and imports at the sectoral level) may be very large.

Third, the trade cost term (T_{ij}) from Equation (1) is decomposed into three structural components in Equation (2). Specifically, t_{ij} denotes any bilateral trade frictions that directly affect the trade flows between two countries, e.g., bilateral distance, tariffs, trade agreements, sanctions, etc.. T_i and T_j are multilateral trade costs on the exporter and on the importer side, respectively, which capture the fact that trade between two countries will not only depend on their sizes and the direct bilateral trade frictions between them, but also on how costly it is for these countries to trade with other countries. Hence, the label 'multilateral' trade costs.

Figure 4: Multilateral Trade Costs

Source: The author. Inspired by Krugman (1995).

To highlight the intuition behind the multilateral trade costs, Figure 4 teleports the United States and Canada on Mars. Even though the sizes (Y_i and E_j) and the direct bilateral trade frictions (t_{ij}, e.g., distance) between the United States and Canada are unchanged, the two countries would trade more with each other if they were on Mars than if they were in their current location on Earth. Why? Because on Mars, they would be more isolated from the rest of the world. Hence, trade between

[1] This equation is consistent with the most influential gravity papers of Eaton and Kortum (2002) and Anderson and van Wincoop (2003), and I refer the interested reader to Yotov et al. (2016) for alternative derivations of the gravity model with the same notation.

the U.S. and Canada does not only depend on the trade costs between these two countries, but also their trade costs with all other countries. This is actually one of the key differences between naive gravity and structural gravity: Naive gravity (wrongly) assumes that only direct bilateral trade costs matter, whereas structural gravity also takes into account the multilateral trade costs.

Two final theory-motivated adjustments deliver the contemporary structural gravity model shown below:

$$X_{ij,t}^k = \frac{Y_{i,t}^k E_{j,t}^k}{Y_t^k} \left(\frac{t_{ij,t}^k}{T_{j,t}^k T_{i,t}^k} \right)^{-\theta^k}, \qquad \forall\, i, j, t, k. \tag{3}$$

First, the gravity model can be derived at any level of aggregation, e.g., for products, industries, sectors, and at the aggregate level. This is reflected by the superscript k in Equation (3), which can refer to a particular product, industry, or sector k. The second adjustment is the subscript t, which is motivated by the fact that, just like trade flows, both the size variables and the trade cost terms in Equation (3) vary over time. The time subscript is also motivated by dynamic trade theories, which imply that trade and trade liberalization may lead to accumulation of factors of production, e.g., physical capital such as the Maquiladoras – manufacturing plants located along the U.S.-Mexico border. Note that without the superscript k and the subscript t, Equation (3) is identical to Equation (2). Finally, each of the underlying gravity theories implies that gravity applies to both international, $i \neq j$, as well as domestic trade, $i = j$.

2.2. Theory Implications for Gravity Applications & Estimations

The theoretical foundations of the gravity equation have very important implications for the empirical success of the gravity model and for the proper specification of the econometric gravity equation. I discuss the importance of each of the main theoretical implications in turn.

Multilateral trade costs. The multilateral trade costs have two implications for quantifying the effects of trade policies. First, if the multilateral trade costs are not properly controlled for in the empirical gravity model, then the gravity equation will under-predict trade between countries that are more isolated from the rest of the world and over-predict trade for countries that are surrounded by many other trade partners, e.g., the European states. Second, the multilateral trade costs capture the ease with which countries may divert trade to other countries. This is extremely important for quantifying the effects of many contemporary policies. For example, one of the main reasons for the ineffectiveness of the sanctions that were imposed on Russia due to its invasion of Ukraine is that Russia was able to divert its trade to non-sanctioning states, e.g., China, India, and Turkey. Such effects are captured by the multilateral trade costs. Another recent example are the 2025 U.S. tariffs on Canada. These tariffs would be particularly harmful to Canada because Canada is relatively isolated from the rest of the world.

Disaggregated Gravity. Theory implies that gravity holds at any level of aggregation – from the product to the aggregate level. The main implication of this theoretical property is that the empirical gravity model is very flexible and, depending on the question of interest, the analysis may focus on a particular product, industry, sector or, more broadly, on goods vs. services. The ability to analyze trade flows separately for individual industries or sectors is important because trade costs (e.g., transportation costs) vary across sectors, and many trade policies (e.g., tariffs) are imposed and implemented at a disaggregated level. Even when the policies are implemented at the aggregate level (e.g., complete trade embargoes), their effects could differ substantially across products, industries, and sectors. Thus, it is often desirable to use a disaggregated gravity model. Gravity theory offers clear support and guidance for disaggregated analyses.

Time-varying Gravity. The time dimension of the gravity equation has several empirical implications. First, adding more years to the gravity data leads to more precise gravity analysis. Second, as demonstrated in the next section, the use of panel (i.e., time-varying) data allows for flexible, comprehensive, and easy modeling of all time-invariant bilateral trade costs (e.g., distance) in the gravity model. Third, using panel data allows for capturing the evolution and adjustment of bilateral trade costs and the effects of various policies over time. For example, the effects of RTAs are not instantaneous, and capturing the adjustment of trade flows over time in response to the formation of RTAs may be useful for policy purposes. Moreover, the effects of the RTAs from the 1990s could be very different from the impact of modern RTAs, and such differences cannot be captured without time-varying data.

Output vs. Expenditure. The difference between the value of output and expenditure, which is implied by gravity theory, has several empirical implications. The first implication is that, if trade costs are symmetric (e.g., the distance between two countries), not allowing for the differences between the value of output and expenditure will lead to poor performance of the gravity model, because it will always predict symmetric bilateral trade flows. Second, distinguishing between output vs. expenditure is even more important at the disaggregated level, where the differences between national production and consumption are much more pronounced. Third, even with aggregate data, GDP is not the best proxy for economic size in the gravity model, not only because theory implies that one should account for trade imbalances, but also because GDP is measured as value added, while trade flows are measured on a gross basis, which causes an inconsistency. Fortunately, as demonstrated in the next section, there are very easy econometric techniques to account for the size variables in the empirical gravity model without the need to worry about such measurement challenges and additional data requirements.

Domestic Trade. Trade theory implies that gravity also applies to domestic sales. Moreover, using domestic trade has significant benefits for gravity estimations, e.g., it allows for identification of possible trade-diversion effects of trade agreements to and from countries that are not part of the agreements, the extraterritorial effects of sanctions, the effects of non-discriminatory trade policies. However, mainly due to data constraints, most gravity estimations are still performed exclusively with international data, and the use of domestic trade flows in gravity analysis is not standard yet. Therefore, consistent with most of the existing gravity literature, we will focus only on international trade flows for the current purposes.

3. From Naive to Modern Gravity Estimations

Much of the great success of the gravity model is due to its unprecedented predictive power and flexibility to accommodate many applications. The objective of this section is to translate the naive empirical gravity equation into a modern econometric gravity model, which can be used to estimate the effects of many determinants of trade flows. The analysis assumes/requires very basic knowledge of econometrics (e.g., using ordinary least squares (OLS) and fixed effects). Thus, the material in this section should be accessible to any undergraduate student who has taken an econometrics course. In addition, this section could be beneficial for students who will be taking an econometrics or an economics seminar course or who plan to write an undergraduate thesis or do an independent research project.

Capitalizing on the theoretical foundations of the gravity model and the knowledge that we have accumulated thus far, we will develop the econometric gravity specification in five steps, which follow the chronological evolution of the estimations of the gravity equation as depicted in Figure 5.[2]

Figure 5: The Evolution of Gravity Estimations

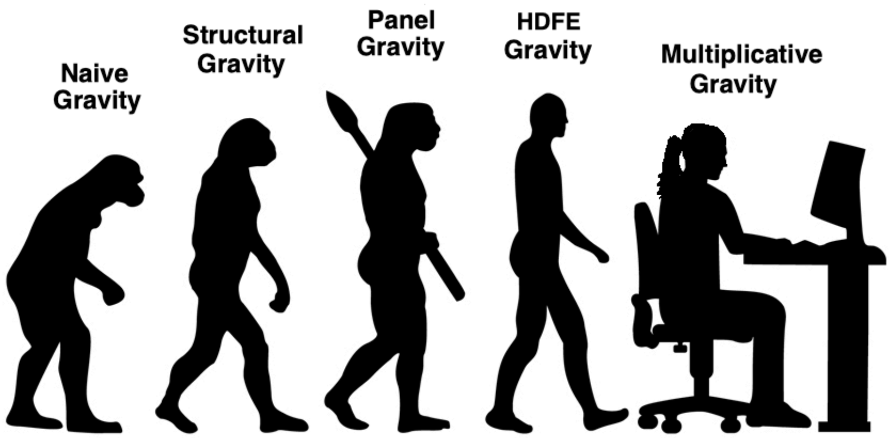

Source: The author.

[2] I refer the reader to Larch et al. (2025) for a detailed motivation and discussion of the recommendations for gravity estimations.

Naive Gravity. Following the approach from the previous section, where we translated the naive gravity equation into a structural gravity model, we begin by specifying a naive estimating gravity equation. To this end, we proceed in three simple steps. First, we log-linearize Equation (1) to obtain:

$$\ln(X_{ij}) = \ln(\tilde{G}) - \theta \ln(T_{ij}) + \ln(Y_i) + \ln(Y_j). \tag{4}$$

Second, we translate Equation (4) into an econometric model by using proxies for the independent variables. Specifically, we use GDP to proxy for the size of the exporter (GDP_i) and of the importer (GDP_j), and we use bilateral distance ($DIST_{ij}$) and regional trade agreements (RTA_{ij}) to proxy for bilateral trade costs. Distance and RTAs are the two most widely used proxies for trade costs in the gravity literature.[3] The use of distance and RTAs as representative proxies for trade costs is also helpful from a pedagogical perspective because one of them (distance) is a continuous variable and the other one (RTA) is an indicator/dummy variable, which takes only values of zeroes and ones. Specifically, the RTA variable is equal to one if two countries have an RTA in force in a given year, and it is equal to zero otherwise.

Finally, we add a constant term (β_0), which corresponds to the gravity constant ($\tilde{G}$), and an error term (ϵ_{ij}), which completes our first econometric gravity specification:

$$\ln(TRADE_{ij}) = \beta_0 + \beta_1 RTA_{ij} + \beta_2 \ln(DIST_{ij}) + \beta_3 \ln(GDP_i) + \beta_4 \ln(GDP_j) + \epsilon_{ij}. \tag{5}$$

Equation (5) is the most popular version of the econometric gravity model. It can be estimated with Ordinary Least Squares (OLS), and the resulting estimates can be interpreted as follows. The estimates of the coefficients on the continuous variables (i.e., distance and GDP) are elasticities. Thus, for example, if the estimate of the coefficient on distance is $\hat{\beta}_2 = -0.8$, this means that, *ceteris paribus*, a 1% increase in the distance between two countries would lead to an 0.8% decrease in the bilateral trade between them. To interpret the impact of any indicator variable in the gravity model (e.g., RTAs), one can apply the following formula $[exp(\beta_1) - 1] \times 100$. Thus, for example, if the gravity estimate on RTAs is $\hat{\beta}_1 = 0.5$, this implies that the RTAs that entered into force during the period of investigation have led to about 65 percent increase ($[exp(0.5) - 1] \times 100 = 64.87$) in trade between the RTA members, *ceteris paribus*.

Structural Gravity. Next, we improve the econometric specification by accommodating the theoretical implications from Equation (2). Specifically, we need to take into account two theory considerations. First, we need to capture the value of output and expenditure. Second, we need to account for the multilateral trade costs. In principle, it is possible to attempt to handle both of these challenges directly, i.e., by using observable data. However, a much easier, theory-consistent, and fully comprehensive econometric approach to account for both (i) output vs. expenditure and (ii) the multilateral trade costs challenges is to use exporter and importer fixed effects, as follows:

$$\ln(TRADE_{ij}) = \beta_0 + \beta_1 RTA_{ij} + \beta_2 \ln(DIST_{ij}) + \pi_i + \psi_j + \epsilon_{ij}. \tag{6}$$

Here, π_i is a set of exporter fixed effects, i.e., dummy variables that take a value of one for each exporter in the data, and ψ_j is a set of importer fixed effects, i.e., dummy variables that take a value of one for each importer in the data. Thus, for example, if the data includes 100 countries and each of them appears as an importer and as an exporter, there will be 100 importer fixed effects and 100 exporter fixed effects.[4] The main advantages of the fixed effects are that: (i) They will fully control for any (observable and unobservable) exporter- and importer-specific characteristics, including size and the multilateral trade costs. Thus, one does not have to worry about omitted variables and collecting data across these dimensions; and (ii) They are very easy to construct and implement empirically with standard statistical software. We demonstrate this in the next section. The disadvantage of these fixed effects is that if they are used, then one cannot identify the effects of any exporter- or importer-specific variables, because these variables will be perfectly collinear with and absorbed by the fixed effects. This does not affect our objectives to estimate the effects of bilateral trade costs and policies (e.g., distance and RTAs).

[3] In the practice questions at the end of the chapter, we also introduce other trade cost proxies and policy variables.

[4] Technically, one of the fixed effects on the exporter side and on the importer side will be dropped due to perfect collinearity. However, this will have no implications for the estimates of the independent variables of interest to us.

Panel Gravity. The next step is simple: it is to introduce a time dimension to our econometric model. The implication is that the error term and all variables (except for distance, which does not vary over time), now have a time subscript:

$$\ln(TRADE_{ij,t}) = \beta_0 + \beta_1 RTA_{ij,t} + \beta_2 \ln(DIST_{ij}) + \pi_{i,t} + \psi_{j,t} + \epsilon_{ij,t}. \quad (7)$$

As discussed in the previous section, adding a time dimension to the econometric analysis is consistent with the theory, and having more data should improve the performance of our econometric model. Moreover, the time dimension enables us to study the evolution of the effects of RTAs before and after their entry into force as well as possible differences in the effects of RTAs that have been formed in different periods of time, e.g., the RTAs of the 1990s vs. the RTAs of the 2000s.

HDFE Gravity. The use of time-varying data additionally allows for the inclusion of country-pair fixed effects (γ_{ij}) – hence, the name of this specification as gravity with High-Dimensional Fixed Effects (HDFE). The country-pair fixed effects are dummy variables that take a value of one for every pair in the data, and they are equal to zero otherwise. The HDFE gravity model becomes:

$$\ln(TRADE_{ij,t}) = \beta_0 + \beta_1 RTA_{ij,t} + \gamma_{ij} + \pi_{i,t} + \psi_{j,t} + \epsilon_{ij,t}. \quad (8)$$

Similar to the exporter and the importer fixed effects, the country-pair fixed effects will fully control for all time-invariant bilateral trade costs, e.g., distance, contiguous borders, etc. The downside of using country-pair fixed effects is that they will not allow us to obtain estimates of the effects of distance or any other time-invariant bilateral variables. Thus, specification (8) no longer includes $\ln(DIST_{ij})$ because it is absorbed by the country-pair fixed effects. However, as long as the objective is to obtain estimates of the effects of time-varying policies, e.g., tariffs, RTAs, sanctions, etc., the use of country-pair fixed effects is strongly recommended for gravity estimations as their inclusion will control for many potentially omitted variables, limit the need for extensive data collection, and significantly improve the predictive power of the gravity model. This will be highlighted and reinforced by the hands-on analysis in the next section.

Multiplicative Gravity. The last step, which leads to the preferred econometric model, is to address the challenge that the log-linear gravity specification drops all bilateral trade flows that are equal to zero (because the logarithm of zero is undefined). The simple adjustment to the econometric gravity equation is to exponentiate both sides, which leads to:

$$TRADE_{ij,t} = \exp\left[\beta_0 + \beta_1 RTA_{ij,t} + \gamma_{ij} + \pi_{i,t} + \psi_{j,t}\right] \times \varepsilon_{ij,t}. \quad (9)$$

Two more features of the multiplicative econometric model make it attractive. First, as demonstrated in the next section, the multiplicative gravity model can be estimated easily and quickly with fast built-in commands in standard statistical packages. Second, the interpretation of the estimates of the gravity coefficients remains the same as in the log-linear specification, which is estimated with OLS.

Equation (9) represents the modern econometric gravity model that is used to obtain estimates of the effects of various policies by academic researchers and policy analysts. Based on Equation (9), Box 5 summarizes the links between the theoretical gravity implications and their empirical implementation.

Box 5. SUMMARY: From Theory to Estimation

Theoretical Implication	Empirical Implementation
Control for the value of output vs. expenditure	Use exporter(-time) and importer(-time) fixed effects
Control for multilateral trade costs	Use exporter(-time) and importer(-time) fixed effects
Account for time-invariant bilateral trade costs	Use country-pair fixed effects (or time-invariant variables)
Account for time-varying bilateral trade costs	Use time-varying (panel) data and policy variables
Account for zero trade flows	Use the multiplicative PPML estimator

Before we continue, I remind the reader that, consistent with theory, the gravity equation applies at any level of aggregation. Thus, Equation (9) can be estimated for individual products, per industry, per sector, or with aggregate data. Importantly, Equation (9) is very easy to implement and estimate with standard statistical software. We demonstrate this with a simple yet state-of-the-art example in the next section.

4. Hands-on Gravity: The Effects of Distance, RTAs, & the EU

This section includes a hands-on analysis that sequentially implements the estimating equations from the previous section. The only difference is that, in addition to estimating the effects of distance and RTAs, we also will isolate the impact of the European Union (EU) from all other RTAs. To this end, we will introduce to our specification a new indicator variable – $EU_{ij,t}$ – that takes a value of one if two countries, i and j, are members of the EU in year t, and it is equal to zero otherwise. Thus, the RTA variable includes all other trade agreements, but it does not include the EU, which is accounted for separately. There are four reasons for this adjustment. First, the impact of EU membership on trade has been of significant interest to both academics and policy makers alike. Second, the impact of the EU is expected to be different/stronger than the effects of other RTAs. Third, this is an example of how to isolate the effects of different trade agreements.[5] Fourth, from a methods perspective, the EU is an example of a cluster of many strongly-integrated countries, which, consistent with the discussion from the theory section, means that the multilateral trade costs may have significant implications for estimating the EU effects.

The empirical analysis in this section will be performed in Stata, and the section includes all commands that are needed to obtain the results, along with some motivation and interpretation of the findings. A single “do” file, which combines all commands from this section, along with the dataset that is used for the estimations, is available at https://yotoyotov.com/Gravity_Undergrads.html. Importantly, the estimating commands that will be introduced in this section are easy to implement in other standard (and free!) statistical packages. Thus, they should be accessible to any advanced undergraduate student. To further facilitate accessibility, the web site includes the data and code used in this section in R too.

The Gravity Data. To accompany this chapter, I compiled the *“Gravity for Undergraduates”* (GU) database, which draws from the most recent (as of March, 2025) trade, policy, and gravity data that are available to academics and policy makers. Thus, the data for the empirical analysis in this section and in the practice questions at the end of the chapter is appropriate for policy analysis and research projects. The GU dataset includes aggregate trade and gravity variables for the 100 largest exporters in the world, covering 98.9% of world exports, 97.7% of world imports, and 98.3% of world GDP, 1990-2023. The list below includes the variables that are employed in the analysis in this section, along with a brief description and data sources. The rest of the variables will be discussed in the Practice Questions section at the end of the chapter.

Exporter	ISO country code of the exporter i. ID variable.
Importer	ISO country code of the importer j. ID variable.
Year	Year t, varying from 1990 to 2023. ID variable.
Trade	Aggregate bilateral trade flows between i and j in year t in nominal/current U.S. dollars. Source: UN COMTRADE database, https://comtradeplus.un.org/.
Distance	Population-weighted distance in kilometers between i and j. Source: USITC DGD database, https://www.usitc.gov/data/gravity/dgd.htm.
RTA	Indicator for the presence of a Regional Trade Agreement (RTA) between i and j in year t. Source: Mario Larch's database, https://www.ewf.uni-bayreuth.de/en/research/RTA-data/.
EU	Indicator that countries i and j are both members of the European Union in year t. Source: The author. Based on data from the European Commission.
GDP_Exporter	GDP of exporter i in year t in nominal/current U.S. dollars. Source: World Bank WDI database, https://datacatalog.worldbank.org/home.
GDP_Importer	GDP of importer j in year t in nominal/current U.S. dollars. Source: World Bank WDI database, https://datacatalog.worldbank.org/home.

[5] The practice questions at the end of the chapter include more examples.

Figure 6 provides an excerpt from the GU data for three countries (Canada, Mexico, and the United States) and three years (1993-1995). The data shows that trade between these three countries is large, asymmetric, and it varies over time. By construction, the distance for each pair is symmetric. The RTA variable is equal to one for trade between Canada and the U.S. in all years, because of the Canada-U.S. trade agreement from 1989, while the RTA variable for U.S.-Mexico and Canada-Mexico switches from zero to one in 1994, because of the North American Free Trade Agreement (NAFTA). The EU variable is always zero because none of these countries are EU members. The GDP variable(s) capture the size difference between the countries. Finally, even though you will not have to create and add fixed effects explicitly, Figure 6 includes four examples of fixed effects. 'CAN_exp_1993' is the exporter fixed effect for Canada in 1993, 'USA_imp_1995' is the importer fixed effect for USA in 1995, 'CAN_MEX' is the country-pair fixed effect for Canada's exports to Mexico, and 'MEX_CAN' is the country-pair fixed effect for Mexico's exports to Canada.

Figure 6: The Structure of the Gravity Data

	Exporter	Importer	Year	Trade	Distance	RTA	EU	GDP_Exporter	GDP_Importer	CAN_exp_1993	USA_imp_1995	CAN_MEX	MEX_CAN
1	CAN	MEX	1993	9.882e+08	3472.085	0	0	5.791e+11	5.302e+11	1	0	1	0
2	CAN	MEX	1994	1.620e+09	3472.085	1	0	5.799e+11	5.536e+11	0	0	1	0
3	CAN	MEX	1995	1.374e+09	3472.085	1	0	6.060e+11	3.802e+11	0	0	1	0
4	CAN	USA	1993	1.136e+11	2134.945	1	0	5.791e+11	6.859e+12	1	0	0	0
5	CAN	USA	1994	1.319e+11	2134.945	1	0	5.799e+11	7.287e+12	0	0	0	0
6	CAN	USA	1995	1.483e+11	2134.945	1	0	6.060e+11	7.640e+12	0	1	0	0
7	MEX	CAN	1993	2.785e+09	3472.085	0	0	5.302e+11	5.791e+11	0	0	0	1
8	MEX	CAN	1994	3.274e+09	3472.085	1	0	5.536e+11	5.799e+11	0	0	0	1
9	MEX	CAN	1995	3.901e+09	3472.085	1	0	3.802e+11	6.060e+11	0	0	0	1
10	MEX	USA	1993	4.072e+10	2492.907	0	0	5.302e+11	6.859e+12	0	0	0	0
11	MEX	USA	1994	5.033e+10	2492.907	1	0	5.536e+11	7.287e+12	0	0	0	0
12	MEX	USA	1995	6.275e+10	2492.907	1	0	3.802e+11	7.640e+12	0	1	0	0
13	USA	CAN	1993	8.804e+10	2134.945	1	0	6.859e+12	5.791e+11	0	0	0	0
14	USA	CAN	1994	1.002e+11	2134.945	1	0	7.287e+12	5.799e+11	0	0	0	0
15	USA	CAN	1995	1.098e+11	2134.945	1	0	7.640e+12	6.060e+11	0	0	0	0
16	USA	MEX	1993	4.832e+10	2492.907	0	0	6.859e+12	5.302e+11	0	0	0	0
17	USA	MEX	1994	5.481e+10	2492.907	1	0	7.287e+12	5.536e+11	0	0	0	0
18	USA	MEX	1995	5.397e+10	2492.907	1	0	7.640e+12	3.802e+11	0	0	0	0

Source: The "Gravity for Undergrads" dataset.

The excerpt in Figure 6 includes 18 observations because there are 3 countries and each of them exports to and imports from two other countries in each of the three years, i.e., $(3\ countries) \times (2\ partners) \times (3\ years) = 18$. Thus, if data were available for all pairs and all years, the number of observations in the GU dataset would be $100 \times 99 \times 34 = 336,600$. However, the GU data is not fully balanced because some countries, e.g., the former Soviet republics, were not independent in the early 1990s, and their trade and GDP data are only available for later years. As a result, the observations in the GU dataset are $320,920$.

Before we proceed with the econometric analysis, we need to transform some of the variables (i.e., trade, distance, and GDP) in logarithmic form. This is achieved by the following simple Stata code:

```
generate ln_trade=ln(Trade)
generate ln_dist=ln(Distance)
generate ln_gdp_exp=ln(GDP_Exporter)
generate ln_gdp_imp=ln(GDP_Importer)
```

Estimating Naive Gravity. The following Stata command line obtains the most 'traditional' (but 'naive') gravity estimates from a simple OLS specification, which corresponds to Equation (5):

```
regress ln_trade ln_dist RTA EU ln_gdp_exp ln_gdp_imp if Year==2023
```

The elements of this estimation line include: 'regress' – the standard Stata command to estimate OLS;

'ln_trade' – the dependent variable; 'ln_dist', 'RTA', 'EU', 'ln_gdp_exp', and 'ln_gdp_imp' – the independent variables. Finally, the clause 'if year==2023' ensures that the estimation is performed for a single year – 2023, which can be changed to any other year in the data, i.e., any year from 1990 to 2023.

The gravity estimates from this specification are reported in column (1) of Table 1. Overall, the results are as expected. Five findings stand out. First, the estimate of the effect of distance on trade is negative and statistically significant, implying that a 1 percent increase in distance will lead to 1.179 percent decrease in trade. The estimates of the effect of RTAs and the EU are both positive and statistically significant. The RTA estimate implies that, all else equal, RTAs have led to a 94 percent (calculated as $[exp(0.664) - 1] \times 100 = 94.25$) increase in trade among the RTA members. Third, perhaps not surprisingly, because the EU is a very deep integration effort, the EU effect on trade is even larger. (Using the calculation for the RTA impact as an example, can you calculate the increase in trade due to the EU among the EU members?) Fourth, the large, positive, and statistically significant estimates of the effects of GDP, both on the importer and exporter side, confirm the strong positive relationship between size and trade. Finally, with an $R^2 = 0.64$, the naive gravity model delivers a strong fit even with only 5 standard explanatory variables.

Table 1: Evolution of the Gravity Estimations

	(1) Naive Gravity	(2) Structural Gravity	(3) Panel Gravity	(4) HDFE Gravity	(5) Multiplicative Gravity
ln_dist	-1.179	-1.473	-1.583		
	(0.032)**	(0.034)**	(0.006)**		
RTA	0.664	0.236	0.147	0.103	0.076
	(0.049)**	(0.051)**	(0.010)**	(0.013)**	(0.012)**
EU	0.900	-0.261	-0.526	0.427	0.273
	(0.112)**	(0.117)*	(0.022)**	(0.030)**	(0.024)**
ln_gdp_exp	1.352				
	(0.015)**				
ln_gdp_imp	1.143				
	(0.015)**				
Constant	-37.864				
	(0.600)**				
N	9564	9564	288085	288085	320920
R^2	0.643	0.759	0.772	0.897	

Source: The author. Standard errors in parentheses. $^{+}$ $p < 0.10$, * $p < .05$, ** $p < .01$. See the main text of the chapter for details on each specification.

Estimating Structural Gravity. We make two changes to the previous estimation line in order to estimate the structural gravity model (6):

```
reghdfe ln_trade ln_dist RTA EU if Year==2023, absorb(Exporter Importer)
```

First, we use the command 'reghdfe'.[6] While it is perfectly fine to use the standard 'regress' command from the previous specification, the 'reghdfe' command is preferred here, because it allows for easy handling of many fixed effects, which is the second change that we implement in the estimation line. Specifically, the use of the option 'absorb (Exporter Importer)' means that the estimation is performed with 'Exporter' and 'Importer' fixed effects, which will control for any observable and unobservable characteristics on the exporter side (e.g., the value of output) and the importer side (e.g., expenditure) as well as the multilateral trade costs that may affect bilateral trade. Thus, the new estimation results, which appear in column (2) of Table 1, no longer include the estimates of the effects of GDP on trade.[7]

[6] The 'reghdfe' command and the other fast command for estimations with high-dimensional fixed effects ('ppmlhdfe'), which will be used in this chapter are due to Correia (2016) and Correia et al. (2020), respectively.

[7] Once fixed effects are included, the estimate of the constant loses its standard interpretation. Therefore, it is omitted as well.

There are four notable differences between the estimates in columns (1) and (2). First, the estimate of the impact of distance in column (2) is still negative and statistically significant; however, it is larger in absolute value. Second, the EU estimate is much smaller, in fact negative. This result is consistent with the discussion of the theoretical implications of the multilateral trade costs from the previous section. Specifically, as predicted by theory, once the multilateral trade costs are accounted for, the impact of the EU is much smaller. Third, we see that the RTA estimate is smaller too; however, it remains positive and statistically significant. Finally, the R^2 in column (2) is larger. This, of course, is expected and due to the use of the exporter and importer fixed effects, which account for all observable and unobservable determinants of trade flows on the exporter and on the importer side, respectively.

Estimating Panel Gravity. We make two changes to the previous estimation line in order to obtain panel gravity estimates that correspond to Equation (7):

```
reghdfe  ln_trade  ln_dist  RTA  EU,  absorb(Exporter#Year  Importer#Year)
```

First, we remove the 'if year==2023' statement, because we want to use all years in the data. Second, we allow the fixed effects in the econometric model to also vary over time, i.e., we now use exporter-time fixed effects (Exporter#Year) and importer-time fixed effects (Importer#Year). Note that the dimensions of the fixed effects in the panel setting are consistent with the theory from the panel gravity model (3), where the exporter-specific and the importer-specific variables (i.e., the country-size variables and the multilateral trade costs) also vary over time.

The panel estimates are reported in column (3) of Table 1. First, we note that, as expected, the number of observations (N) in column (3) is significantly larger. On a related note, i.e., due to the use of more data, the standard errors in column (3) are indeed smaller. The estimate of the effect of distance is comparable to the corresponding cross-section result from column (2). The estimate of the effect of RTAs is smaller, but still positive and statistically significant. However, the estimate of the EU in column (3) is large, negative, and statistically significant. This result is very surprising from a policy perspective. However, it is not surprising from a methods perspective, and the explanation for it is that this panel specification omits many potentially important bilateral trade costs. This issue is resolved in the next specification.

Estimating HDFE Gravity. The next adjustment to the panel gravity model is to introduce country-pair fixed effects. This is easily done in the 'absorb' option of our estimation command, where, in addition to the exporter-time and the importer-time fixed effects, we now add country-pair fixed effects – 'Importer#Exporter'. The corresponding estimation line becomes:

```
reghdfe  ln_trade  RTA  EU,  absorb(Exporter#Year  Importer#Year  Importer#Exporter)
```

Note that we no longer see the variable 'ln_dist' in the new estimation line. The reason is that once the country-pair fixed effects are introduced, they will absorb and fully control for all possible time-invariant bilateral trade costs, including distance. To the extent that much of the bilateral trade costs are determined by geography, the country-pair fixed effects are very powerful econometric tools to control for many variables that cannot be easily measured or even observed. This will enable us to focus on the policy variables of interest, e.g., the EU and RTAs in our case, which are also bilateral but vary over time. If the interest is in the effects of time-invariant determinants of trade flows (e.g., distance, contiguous borders, common language, etc.), then the country-pair fixed effects cannot be used as their inclusion precludes such estimation. Such analyses are included in the practice questions at the end of the chapter.

The gravity estimates with pair fixed effects are reported in column (4) of Table 1. This specification leads to a very different estimate of the EU effect. Specifically, the estimate of the impact of the EU is now large, positive, and statistically significant, as expected. The RTA estimate is a bit smaller than before, but still positive and statistically significant. As noted earlier, the explanation for the changes in the policy estimates is that our previous specifications have omitted some key time-invariant independent variables, which are now fully controlled for by the country-pair fixed effects. Thus, as long as the object of interest is the effect of bilateral policy variables (e.g., RTAs, EU membership, tariffs, sanctions, etc.), the specification with country-pair fixed effects is highly recommended. The use of the country-pair fixed effects has led to a further increase in the overall fit of the model, as captured by $R^2 = 0.9$.

Estimating Multiplicative Gravity. In our last specification, we estimate the gravity model in multiplicative form. The corresponding estimation line becomes:

```
ppmlhdfe  Trade  RTA  EU,  absorb(Exporter#Year  Importer#Year  Importer#Exporter)
```

The two simple adjustments that enable us to estimate gravity in multiplicative form are (i) to replace the 'reghdfe' command with the 'ppmlhdfe' command, and (ii) to use trade flows in levels: 'Trade' instead of the logarithm of trade flows 'ln_trade'. The new 'ppmlhdfe' command reflects the fact that we will use the multiplicative PPML estimator instead of the OLS estimator. The PPML estimator has established itself as the standard gravity estimator because it has a number of econometric advantages over OLS, including the fact that, due to its multiplicative form, PPML is able to account for zero trade flows in the data.[8] Importantly, the interpretation of the gravity estimates remains the same when they are obtained with PPML. Moreover, due to recent computation developments, the PPML estimator is robust and fast.

The last set of gravity estimates appears in column (5) of Table 1. The multiplicative gravity model delivers an RTA estimate of 0.076, which implies that, *ceteris paribus*, the impact of the RTAs that entered into force during the period of investigation is an increase of about 8 percent among the RTA members. The EU estimate implies a corresponding effect of about 31 percent. The estimates of both policy variables may seem a bit small, however, one should keep in mind the following. First, given the use of country-pair fixed effects, our estimates only capture the effects of the RTAs that entered after 1990 and those due to the EU members that joined after 1990. Second, the effects of both the EU and the RTAs can be very heterogeneous. For RTAs, this is demonstrated in the Practice Questions section at the end of the chapter. Finally, note that there is no R^2 reported in column (5).[9] Instead, to gauge the overall fit of the model, I refer the reader back to Figure 2, which is obtained based on the estimates from this specification.

A Final Consideration. If you understand the meaning, motivation, and implications of each of the elements in the multiplicative gravity specification that we just discussed, then you have gone a long way and should be very proud of yourself. As will be demonstrated in the Practice Questions section at the end of the chapter, by just adding the corresponding independent variables to this estimation line, you will be able to estimate the effects of many time-varying bilateral policies (e.g., sanctions, Customs Unions, etc.) on trade. Moreover, if you are interested in the effects of any time invariant determinants of trade flows (e.g., distance, colonial ties, common official language, etc.), you can include these variables instead of using the country-pair fixed effects. Overall, by now you are very well equipped to study the impact of many policies on trade flows with a modern empirical gravity model.

Before we conclude, there is one last element of the econometric gravity model that we did not discuss. This element is related to the treatment of the standard errors in gravity estimations, and it was not discussed earlier for three reasons. First, the topic is perhaps too advanced for the current purposes. Second, there are no firmly established practices to address this issue. Third, the estimates of the coefficients on the independent variables will not change, while the corresponding standard errors may increase or decrease, however, often this only leads to small differences in the gravity results. Therefore, without going into details, my recommendation is to follow one of the most widely used practices to treat the gravity standard errors by simply adding the option 'cluster(Exporter#Importer)' at the end of your estimating gravity model. Since prominent academic econometricians have not agreed on this yet, you should not worry about it and, if asked why you made this adjustment, you can simply respond "Blame it on Yotov!".

[8] I refer the interested reader to Santos Silva and Tenreyro (2006) and Larch et al. (2025) for detailed discussions of the attractive properties of PPML for gravity estimations.

[9] The reason is that the non-linear PPML estimator produces a 'pseudo-R^2' that is not comparable to the R^2s from the previous OLS specifications. See https://www.statalist.org/forums/forum/general-stata-discussion/general/1528609-ppmlhdfe-pseudo-r2.

References

Anderson, James E.. 1979. "A Theoretical Foundation for the Gravity Equation." *American Economic Review* 69 (1): 106–116.

Anderson, James E.. 2011. "The Gravity Model." *Annual Review of Economics* 3 133–160.

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

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

Baldwin, Richard E., and Daria Taglioni. 2006. "Gravity for Dummies and Dummies for Gravity Equations." *NBER Working Paper No. 12516*.

Correia, Sergio. 2016. "A Feasible Estimator for Linear Models with Multi-Way Fixed Effects." *unpublished manuscript available at http://scorreia.com/research/hdfe.pdf*.

Correia, Sergio, Paulo Guimarães, and Thomas Zylkin. 2020. "Fast Poisson Estimation with High-Dimensional Fixed Effects." *The Stata Journal* 20 (1): 95–115. 10.1177/1536867x20909691.

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

Frankel, J.A., and D. Romer. 1999. "Does Trade Cause Growth?" *American Economic Review* 89 (3): 379–399.

Krugman, Paul. 1995. "Increasing returns, imperfect competition and the positive theory of international trade." In *Handbook of International Economics*, edited by Grossman, G. M., and K. Rogoff Volume 3. 1st edition, Chap. 24 1243–1277.

Larch, Mario, Serge Shikher, and Yoto Yotov. 2025. "Estimating Gravity Equations: Theory Implications, Econometric Developments, and Practical Recommendations." Working Papers 2025001, Center for Global Policy Analysis, LeBow College of Business, Drexel University, `https://ideas.repec.org/p/drx/wpaper/2025001.html`.

Ravenstein, Ernest George. 1885. "The Laws of Migration: Part 1." *Journal of the Statistical Society of London* 48 (2): 167–235.

Rose, Andrew K. 2000. "One money, one market: the effect of common currencies on trade." *Economic Policy* 15 (30): 08–45.

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

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

Yotov, Yoto V. 2024. "The evolution of structural gravity: The workhorse model of trade." *Contemporary Economic Policy* 42 (4): 578–603.

Yotov, Yoto V., Roberta Piermartini, Jose-Antonio Monteiro, and Mario Larch. 2016. *An Advanced Guide to Trade Policy Analysis: The Structural Gravity Model*. Geneva: UNCTAD and WTO.

Practice Questions

The practice questions in this section are designed with the following objectives in mind. First, to reinforce the material and tools that were covered in the Chapter. Second, to practice and extend your coding abilities. Third, to provide more and realistic policy applications. The questions are ordered based on difficulty, from easier to more difficult. Each question is independent and could/should be solved as such. Some of the questions are split in parts and the solution to each of them requires only a couple of lines of code. The *"Gravity for Undergraduates"* database is the only data that is needed for all questions.

1. **Physics & Trade.** Use the GDPs for the exporter and the importer and the Distance variable for the year 2023 from the GU dataset to construct trade according to the naive trade gravity model (1) from the Chapter. Consistent with the existing trade literature, assume that the elasticity of trade with respect to distance is $\theta = -1$. Moreover, since you will be interested in the correlation between constructed and actual trade flows, assume that the trade gravity constant is $\tilde{G} = 1$. Correlate the trade flows that are constructed from the naive gravity model and the actual trade flows from the data. What is the correlation coefficient? Now, compare the means of the two variables. What can explain the differences?

2. **Cross-section estimates over time.** Use the GU database to estimate the "Naive Gravity" model from Section 4 of the Chapter (i.e., the specification in column (1) of Table 1) for the years 1990, 2000, 2010, and 2020. Discuss the evolution over time of the estimates of the five gravity variables. (Hint: you can estimate 4 separate equations or you can write a loop.)

3. **The standard gravity variables.** The log of bilateral distance, the presence of contiguous borders, common official language, and colonial relationships are the four most widely used proxies for time-invariant trade costs that are used in the trade literature. In this question, you will use the GU database to estimate the effects of these 'standard' gravity variables on trade.

3.a. Start with the "Multiplicative Gravity" model from Section 4 of the Chapter (i.e., the specification in column (5) of Table 1) and add the variables 'ln_dist', 'Contiguous_Border', 'Common_Language', and 'Colonial_Ties' directly to this specification. Can you identify the effects of these variables? Why?

3.b. Now estimate the same specification but without the country-pair fixed effects. What is the element that adds the country-pair fixed effects in your specification? Eliminate it. Do you obtain estimates of the new covariates from part 3.a.? Interpret the new estimates in terms of sign and magnitude.

4. **The effects of 'Customs Unions'.** Customs Unions are a form of RTAs, where the RTA members also adopt a common external trade policy, i.e., for trade with non-member countries. Thus, Customs Unions are 'deeper' RTAs and, as such, Customs Unions are expected to promote trade more than RTAs. In this question, you will use the GU database to test this hypothesis.

4.a. Call the GU data and estimate the "Multiplicative Gravity" model from Section 4 of the Chapter (i.e., the specification in column (5) of Table 1) after adding the variable 'Customs_Union'. Interpret and discuss the estimate on 'Customs_Union'.

4.b. Since, by definition, Customs Unions are RTAs, the RTA variable already includes the Customs Unions. Thus, the estimate on 'Customs_Union' from part 4.a. should be interpreted as a deviation from the estimate on RTAs. To obtain the estimate of the effects of Customs Unions in levels, set the RTA dummy variable to zero when 'Customs_Union' is equal to one. Then, re-estimate the model from part 4.a. Compare the 'Customs_Union' estimates from parts 4.a. and 4.b., and then interpret the new estimate.

5. **The effects of 'Trade Sanctions'.** Economic sanctions have become an integral part of contemporary statecraft. In this question, you will use the GU database to estimate the effects of complete trade sanctions, and you will isolate the effects of the sanctions on Russia.

5.a. Start with the GU data and estimate the "Multiplicative Gravity" model from Section 4 of the Chapter (i.e., the specification in column (5) of Table 1) after adding the variable 'Trade_Sanction'. Interpret and discuss the estimate of the effects of sanctions.

5.b. Start with the previous specification from part 5.a. and isolate the effect of the sanctions on Russia. To this end, use the variable 'Trade_Sanction' to construct a new variable for the sanctions on Russia – 'Russia_Sanction' – which takes a value of one if Russia is part of a trade sanction, and is equal to zero otherwise. Then, set the 'Trade_Sanction' variable to zero when 'Russia_Sanction' is equal to one, and estimate the model from part 5.a. after adding the variable 'Russia_Sanction'. Interpret your results.

5.c. Isolate the effects of the 2022 sanctions on Russia due to its invasion of Ukraine. To this end, split the variable 'Russia_Sanction' from the previous specification into two variables: 'Russia_Sanction_Old' for the sanctions before 2022; and 'Russia_Sanction_New' for the sanctions that were imposed in 2022. Then, estimate the model from part 5.b. after replacing the 'Russia_Sanction' variable with the two new variables for the sanctions on Russia. Discuss and interpret your results.

6. **The heterogeneous effects of RTAs.** The single RTA estimate that was obtained in the Chapter may hide significant heterogeneity in the effects of RTAs. In this question, you will use the GU database to explore several dimensions of the heterogeneity of the effects of RTAs.

6.a. Start with the GU data and estimate the "Multiplicative Gravity" model from Section 4 of the Chapter (i.e., the specification in column (5) of Table 1) after allowing for the effects of RTAs to be different for the periods 1990-1999, 2000-2009, and 2010-2023. (Hint: You need to split the single RTA variable in three RTA variables, one for each period.) Discuss your results.

6.b. Start with the "Multiplicative Gravity" model from Section 4 of the Chapter (i.e., the specification in column (5) of Table 1) and isolate the effects of the Australia-China FTA (ChAFTA) from 2015. To this end, construct an indicator variable ChAFTA, which takes a value of 1 for trade between Australia and China for the years after 2014. Then, add this variable to the "Multiplicative Gravity" model.

6.c. The specification from 6.b. delivers an estimate of ChAFTA as a deviation from the effect of RTAs. The reason is that the RTA variable includes ChAFTA. To obtain an estimate of the full ChAFTA effect, set the RTA variable to zero when ChAFTA is equal to one. Then estimate the specification from part 6.b. again. Compare the ChAFTA estimates from the two specifications and interpret the new ChAFTA estimate.

6.d. Test whether the effects of ChAFTA are asymmetric. (Hint: split the ChAFTA variable into two variables depending on the direction of trade flows, i.e., from Australia to China and from China to Australia. Then, use these two variables instead of the ChAFTA variable in the previous specification and compare their estimates.)