Essays on International Trade and Economic Geography

Uğur Yeşilbayraktar

TESI DOCTORAL UPF / Year 2022

THESIS SUPERVISORS

Giacomo Ponzetto and David Nagy

Department Department d'Economia i Empresa

To my parents, Ayla and Haluk, for always encouraging me to follow my passion for learning and for their unwavering support.

Acknowledgements

I am deeply indebted to my advisors Giacomo Ponzetto and David Nagy for their continuous guidance, patience and support throughout the years. This body of work would not have been possible without their contributions and insights. I am also incredibly grateful to Jaume Ventura for his mentoring throughout the years and for always challenging me with inspiring questions. I was fortunate enough to learn from Fernando Broner, Alberto Martin, Albrecht Glitz, Manuel Garcia-Santana and Joan Monras and I would like to thank them for their countless insights over the years which has helped shaped this thesis and my research ideas. I would like to also thank to all participants of the CREI International Lunch for their helpful comments and suggestions. I thank Marta Araque and Anna Rios for their help over the years and making life easier and allowing me to focus on my research to a greater extent.

I owe everything to my parents, Ayla and Haluk, who has supported me through all the challenges I faced. I am eternally grateful for their sacrifices and for always being there for me.

Abstract

This thesis explores how international trade shapes economic geography and analyses the effects of borders and other trade costs on the pattern of trade. In the first chapter, I study how the reduction in trade costs stemming from the 2004 Enlargement of the European Union contributed to spatial inequality within and across countries. I develop an open-economy model of economic geography to structurally estimate the effect of EU enlargement on differential growth in urban centres. In the second chapter, I employ a newly constructed dataset to estimate the "border effect" on trade in Europe using novel empirical techniques in the trade literature. I find that borders in Europe are still a significant impediment to trade and represent a cost equivalent to a 32.5 percent tariff. In the third chapter, I explore the trade patterns within and across country borders in Europe, focusing on the differences between home, country and foreign trade. I document that European regional trade has a strong home and country bias and that this strong home bias is heterogeneous across geography.

Resum

Aquesta tesi explora com el comerç internacional configura la geografia econòmica i analitza els efectes de les fronteres i altres costos comercials sobre el patró del comerç. En el primer capítol, estudiem com la reducció dels costos comercials derivada de l'ampliació de la Unió Europea de 2004 va contribuir a la desigualtat espacial dintre i entre països. Desenvolupo un model d'economia oberta de geografia econòmica per estimar estructuralment l'efecte de l'ampliació de la UE sobre el creixement diferencial dels centres urbans. En el segon capítol, faig servir un conjunt de dades recentment construït per estimar l'"efecte frontera"sobre el comerç a Europa utilitzant noves tècniques empíriques de la literatura comercial. Trobo que les fronteres a Europa segueixen sent un obstacle important per al comerç i representen un cost equivalent a un aranzel del 32,5 per cent. En el tercer capítol, exploro els patrons comercials dins i a través de les fronteres dels països a Europa, centrant-me en les diferències entre el comerç nacional, nacional i exterior. Documento que el comerç regional europeu té un fort biaix nacional i nacional i que aquest fort biaix domèstic és heterogeni a través de la geografia.

Preface This thesis explores how international trade shapes economic geography and analyses the effects of borders and other trade costs on the patterns of trade. In the first chapter, I study how the reduction in trade costs stemming from the 2004 Enlargement of the European Union contributed to spatial inequality within and across countries. I develop an open-economy model of economic geography to structurally estimate the effect of EU enlargement on differential growth in urban centres. In the second chapter, I employ a newly constructed dataset to estimate the "border effect" on trade in Europe using novel empirical techniques in the trade literature. I find that borders in Europe are still a significant impediment to trade and represent a cost equivalent to a 32.5 percent tariff. In the third chapter, I explore the trade patterns within and across country borders in Europe, focusing on the differences between home, country and foreign trade. I document that European regional trade has a strong home and country bias and that this strong home bias is heterogeneous across geography.

This thesis explores how international trade shapes economic geography and analyses the effects of borders and other trade costs on the patterns of trade. In the first chapter, I investigate to what extent urban divergence in Europe is explained by market integration. I develop a quantitative model of economic geography and structurally estimate it to measure the effects of the 2004 enlargement of the European Union on urban growth, spatial inequality and welfare. I find that market integration caused by EU enlargement accounts for 26 percent of the observed increase in residents of the major cities in Central Europe. In terms of welfare, I find that all countries gain at the aggregate level, and the average increase in real income is 0.21 percent. Overall, new member states gain the most as their real income increases by 0.4 percent, implying that the enlargement entailed a decrease in overall income inequality in Europe. Beyond this, my model also delivers predictions regarding the evolution of within-country spatial inequality. In Germany, the effects are progressive, as real income growth in East Germany is significantly higher than in the West. In contrast to this, long standing economic disparity between the North and South Italy worsens as a result of the eastward expansion of the union. Northern Italian hubs become even more central relative to those in the South within the new single market and gain more.

The second chapter investigates whether country borders are still an impediment to trade flows within Europe. We construct a matrix of bilateral trade flows for 269 European regions across 24 countries using a microlevel survey with more than 3 million annual shipments of goods transported by road freight. This matrix provides the first integrated view of regional trade within Europe. In order to estimate the magnitude of the "border effect" we employ a causal inference framework. Our results show that national borders are still a significant impediment to trade within Europe. In particular, we find that, on average, trade between intranational region pairs is 5.714 times larger than international region pairs. This is equivalent to a 32.5 percent bilateral tariff.

In the third chapter, we use this new dataset to systematically explore for the first time trade patterns within and across country borders in Europe and focus on the differences between home trade, country trade and foreign trade. First and foremost, we document that European regional trade has a strong home and country bias. Next, we find that geographic distance and national borders are important determinants of regional trade but cannot explain the strong regional home bias present in the data. Finally, we show that this strong home bias is heterogeneous across regions and is driven by political regional borders.

Contents

List of Figures

List of Tables

Chapter 1

EUROPEAN INTEGRATION AND THE RISE OF SUPERSTAR CITIES

1.1 Introduction

Spatial inequality within developed countries has been rising over the last few decades. This phenomenon is particularly well documented for the United States, as large cities continue to attract an ever growing number of skilled workers (Berry and Glaeser, 2005) (Moretti, 2012) and diverge from their smaller and more peripheral counterparts in terms of economic performance (Giannone, 2017). Western Europe has experienced similar trends, with rising disparities among metropolitan areas since about 2004 (Ehrlich and Overman, 2020). The diverging fortunes of Europe's cities and concomitant political, social and economic consequences are becoming a growing cause for concern (Rodríguez-Pose, 2018).

Concurrent with this rise in spatial disparities; Europe has also undergone a dramatic process of market integration, that culminated with the 2004 enlargement of the European Union as its keystone. Thereafter, eight

Eastern European countries along with Cyprus and Malta joined the union and became permanent members of the single market.[1] These two major developments in European economic geography may well be connected. Theoretically, the link between market access and agglomeration is a seminal result in economic geography (Krugman, 1991). Empirically, however, the question remains open: How much did European integration contribute to the diverging fortunes of European cities?

This paper investigates to what extent urban divergence in Europe and the growing pre-eminence of regional superstar cities can be explained by market integration. I build a quantitative model of economic geography that connects trade and city growth. The model delivers predictions that are consistent with the data, where I observe differentially higher growth in cities that gain more in terms of market access. I structurally estimate my model to measure the effects of the 2004 enlargement of the European Union, where I treat the accession of new member states into the union as a reduction in border related trade barriers between the old (EU15) and new member states (NMS). I find that the EU enlargement can explain 26% of the rise in population of the major cities in Central Europe.

My model economy features a large set of locations that differ in their amenities and their proximity to one another. Trade takes place only at a subset of locations that serve as trading hubs and generates positive spillovers in the form of local consumer amenities. Labor is mobile across locations, subject to mobility frictions. The spatial equilibrium of my model entails a rich geography both across and within market areas that form endogenously around each hub. As trade opportunities increase, the cost of distance from the hub becomes more salient. As a result, workers move closer to the hub through which they trade. Residents of the hub cities also benefit from

[1] New member states had individual free trade agreements with the EU prior to their accession. However their admission into the EU still had a significant effect on trade costs via reductions in tariffs and non-tariff barriers. Furthermore, abolishing customs checkpoints between countries was a major milestone in increasing market integration and led to efficiency gains in terms of travel time.

consumer amenities that grow larger as trade expands. This endogenous increase in consumer amenities amplifies the rise in *regional urban primacy*, or the concentration of population that reside in the trading hub within each market area.

Besides this overall increase in urban primacy, the model predicts heterogeneous impacts across hubs, depending on the increase in their market access. Hubs located close to the border between the NMS and the EU15 such as Berlin, Vienna and Ljubljana reap the largest gains in market access and thrive. In contrast, Southern Italian hubs such as Naples and Palermo become even more peripheral and suffer from trade diversion. In general equilibrium, they may even lose population as economic activity reallocates across market areas towards those that benefit the most from the eastward expansion of the union.

Identifying the relevant urban hubs is the starting point for my quantitative empirical analysis. I rely on the algorithm developed by Mori et al. (2020) to construct a hierarchical city system for each country by recursively partitioning its geography based on the location and the size of its largest cities. This procedure selects 73 hubs from the total sample of 304 cities across 8 countries.[2] Having selected the hubs, I establish two key facts in the data that are consistent with the predictions of my model. First, in the wake of 2004, annual population growth among hubs increased on average by 0.4 percentage points. Second, this increase in population growth was stronger for hubs that were located within 50 kilometers of the accession border, which experienced an additional 0.5 percentage points increase in annualized population growth.[3]

In light of these empirical facts, I structurally estimate the model to

[2] I use the set of cities defined by the Eurostat as candidates for hubs.

[3] In my quantitative and empirical analysis, I specifically focus on the set of countries that are contiguous to the border between the (EU15) and the (NMS). Accession of the NMS into the union did have transformative effects on the continent as a whole. Even so, shifts in relative market access generated by the integration of Eastern European economies into the single market were strongest within the geography I consider.

quantify the effects of the 2004 Enlargement of the EU on the spatial distribution of economic activity and population. I calibrate my model to match the hub and country population levels before EU enlargement. Although I only match the population levels of the 73 hubs, the model can credibly fit the entire population distribution observed in the data, attested by the high correlation (.48) between model-implied and data-based populations of the 31,538 municipalities in 8 countries that make up my full sample.[4] To isolate the effects of market integration, I then solve for the counterfactual scenario in which all structural parameters of the model are kept fixed except border-related barriers to trade.

My results show how economic integration contributed to the reshaping of European economic geography. In my counterfactual, the aggregate population of all urban hubs grows by .4%, equivalent to 26% of the 1.5% increase observed in the data. However, my counterfactual abstracts from changes in overall population. The share of total population living in all hubs increases by .08 percentage points, which is almost the whole increase observed in the data. Turning to the effects on real income, I find that all countries gain at the aggregate level, and the average increase in real income is 0.21%. Overall, new member states are the biggest winners from the enlargement process. On average, their real income increases by 0.4%. Since Eastern European countries are initially poorer, enlargement entails a decrease in overall income inequality. The standard deviation of real income across locations in my sample falls by 0.5%.

My model also speaks to the evolution of within-country spatial inequality. In Germany, the effects are progressive, as real income growth in East Germany (0.17%) is considerably higher than in the West (0.09%). Integration of Eastern European markets improves the centrality of East German hubs within Europe's trade network relative to the West, allowing them to catch up. Conversely, European integration exacerbates the long-standing economic dichotomy in Italy between the North and South. Southern Italian

[4] Population data are provided by the Europe-wide census conducted in 2001.

hubs become even more peripheral relative to their Northern counterparts within the new single market. As a result, the North continues to pull away.

My work is related to the extensive literature that studies the effects of trade on local economic outcomes such as population (Redding and Sturm, 2008; Bleakley and Lin, 2012) , employment and wages (Brülhart et al., 2012, 2018; Caliendo et al., 2019; Ducruet et al., 2020), regional development and inequality (Dix-Carneiro and Kovak, 2017; Brooks et al., 2018) and country development (Pascali, 2017; Feyrer, 2019). It is also related to the literature on the determinants of urban growth and divergence such as skill composition (Berry and Glaeser, 2005; Shapiro, 2006; Moretti, 2012; Diamond, 2016; Ganong and Shoag, 2017), institutional quality (Henderson and Wang, 2007), technological change (Henderson and Wang, 2005) and trade openness (Krugman and Elizondo, 1996; Bakker, 2020). I bridge these two strands of the literature by studying the effects of European integration on urban concentration. To my knowledge, this is the first paper that focuses on the link between trade and urban primacy within the context of trade liberalization among modern economies.

My paper fits in a broader literature on quantitative spatial economics. Redding and Rossi-Hansberg (2017) provide an extensive review of this rapidly growing literature. One strand of this literature proposes models in which trade contributes to urban growth via different mechanisms (Coşar and Fajgelbaum, 2016; Fajgelbaum and Redding, 2018; Bakker, 2020). My paper contributes to this literature by developing a quantitative spatial economic framework based on Nagy (2020), who studies the effects of the dissolution of the Austro-Hungarian empire on the economic geography of Hungary. My model extends this framework to incorporate multiple countries. I do not assume, as Nagy (2020) did, that borders are prohibitive to trade. Instead, I estimate the magnitude of border-related barriers before and after 2004 and measure the effects of this change on the spatial distribution of economic activity.

Moreover, my model incorporates trade-driven endogenous amenities in urban hubs, which Nagy (2020) did not consider. A growing body of literature highlights the importance of local consumption amenities in spurring urban growth (Glaeser et al., 2001; Rappaport, 2007; Carlino and Saiz, 2008; Diamond, 2016). Prior work mainly related endogenous amenities to a city's population or skill composition. In this paper, I focus on a different determinant: economic activity arising from trade. Intuitively, trade makes a city more cosmopolitan and facilitates its development as a center of culture, art and science. Furthermore, well-connected hubs attract a greater variety of brands and a more diverse mix of people, both of which are not available in smaller, more peripheral locations. Thus, urban hubs provide exclusive access to a wide range of activities and goods, and also contain a richer social fabric. Trade enhances their capacity to provide such services and generates spillovers in the form of consumer amenities.

My analysis also complements studies that focus on the economic implications of the EU enlargement (Baldwin et al., 1997; Henrekson et al., 1997; Dustmann and Frattini, 2012; Kennan, 2017; Caliendo et al., 2021). I further this literature by assessing the impact of European integration on differential city growth.

The rest of the paper is organized as follows. Section 2 develops the quantitative model. Section 3 describes its calibration and structural estimation. Section 4 discusses the model's predictions and presents supportive evidence. Section 5 presents the results of the quantitative exercise. Section 6 concludes.

1.2 A Model of Trade and Urban Growth

In this section, I develop an open-economy quantitative model of economic geography and trade. Trade takes place only at a subset of locations that serve as trading hubs and generates positive spillovers in the form of local consumer amenities. I assume that workers pay a utility cost when shipping their products between their residential location and their trading hub. As a consequence of this, cities will form around hubs, with population density gradually decreasing with distance from the center of the hub. Given this setting, a reduction in trade costs will lead to real income gains and spur urban growth.

In my model, a geography G consists of a finite number of locations $r \in G$ distributed across multiple countries C. Each worker produces a uniquely differentiated good and production of goods requires labor only. Labor endowment of each worker is fixed and is normalized to one. Goods can be traded within a country or between two different countries, however cross border trade is subject to additional border related costs. I assume that goods can only be traded at a subset of locations $\mu_i \in G$, which I refer to as *hubs*. Thus, in equilibrium, workers choose a residential location r and a trading hub h where they sell their produce and engage in trade.

1.2.1 Consumption

Worker i, who resides in location r and trades in hub h, obtains the following utility.

$$u_h(r,i) = a(r,i) + \zeta(h,r)^{-1} \left[\sum_{j=1}^{\bar{L}} c_h^j(r,i)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}} \tag{1.1}$$

Here, $a(r,i)$ denotes the level of amenities that worker i enjoys at her residential location, $\zeta(h,r)^{-1} \geq 1$ is the utility cost workers pay for shipping their goods between their residence and their trading hub, $c_h^j(r,i)$ is the consumption of worker i who trades in hub h of worker $j's$ product and $\sigma \geq 1$ is the elasticity of substitution across goods. Because goods are substitutes, in equilibrium there is trade across all workers and by extension all hubs.

Workers are heterogeneous in their tastes for residential locations. Amenities $a(r,i)$ capture location specific features that are common across all workers as well as idiosyncratic factors that are specific to the individual and take the following form

$$a(r,i) = a(r) + \epsilon(r,i) \tag{1.2}$$

where $a(r)$ captures the location specific component of the amenity that is common across workers whereas $\epsilon(r,i)$ is the idiosyncratic component that reflects heterogeneous tastes. I assume that $\epsilon(r,i)$ is i.i.d across workers and locations and follows a Gumbel distribution:

$$F(\epsilon(r,i) \leq z) = e^{-e^{-\frac{z}{\theta}}} \tag{1.3}$$

where θ is the shape parameter that determines the heterogeneity in idiosyncratic preferences. This heterogeneity in tastes is the main dispersion force in the model. Higher values of θ imply greater heterogeneity in tastes. As $\theta \to \infty$, this heterogeneity in location preferences becomes large enough to dominate both the distribution of location specific amenities $a(r)$ and consumption driven incentives. As a consequence of this, at the limit each location r will host the same number of workers.

The location specific component of amenities that is common across all workers $a(r)$ takes the following form.

$$a(r) = \bar{a}(r) + \tilde{a}(r) \tag{1.4}$$

where $\bar{a}(r)$ represents the exogenous component (such as having a beach or having pleasant weather) and $\tilde{a}(r)$ captures endogenous amenities that are driven by trade.

I assume that endogenous trade related amenities take the following form

$$\tilde{a}(r) = T(r)^{\alpha} \tag{1.5}$$

where T(r) is the total volume of trade conducted at location r and α is the elasticity of trade related amenities. In Section 3, I expand on the structure of amenities and provide a detailed explanation on how I estimate the parameter α that drives the relationship between trade and amenities.

Lastly, I assume that the utility cost of shipping goods between the residential location r and its associated hub $\mu(r) = h$ is exponential in distance.

$$\zeta(\mu(r), r) = e^{\psi \cdot dist_{\mu(r),r}}$$

The presence of shipping costs, in combination with the fact that trade is mediated by specific hubs, generates a force of agglomeration since workers will prefer living closer to their chosen hub in order to save on shipping costs. I assume that shipping costs captured by $\zeta()$ are prohibitive across borders. This implies that in equilibrium workers will always choose a hub that is located within their country of residence.

1.2.2 Production

Each worker is endowed with one unit of labor and producing one unit of a good requires one unit of labor as well. Workers produce their goods at their residence r and ship their products to their hub h of choice. Here, they engage in monopolistic competition. Worker i chooses the price of her product p_h^i taking into account the price index at hub h, P_h and set of transport costs across hubs $\tau(\mu_h, \mu_h) \geq 1$. For simplicity, I assume that transport costs are symmetric between each trading hub pair i.e $\tau(\mu_h, \mu_o) = \tau(\mu_o, \mu_h)$ and take the following form:

$$\tau(\mu_h, \mu_o) = e^{\phi \cdot dist_{\mu_h,\mu_o} + 1_{h,o}\beta_{h,o}}$$

where ϕ is the elasticity of trade costs with respect to distance, $1_{m,o}$ is an indicator that equals 1 if hubs h and o are located in different countries and $\beta_{m,o}$ is the border related barriers that hinder trade. The exponential formulation of transport costs is a common assumption in international trade (Head and Mayer, 2014) and economic geography (Desmet and Rossi-Hansberg, 2014). I augment this functional form by adding border related barriers. As shown by Santamaría et al. (2020) borders are still an important impediment to trade. In addition to tariffs (that existed to some degree between EU15 and NMS countries prior to their accession into the union), border related barriers also capture non-tariff barriers stemming from judicial differences and uncertainty (Turrini and van Ypersele, 2010) and informational barriers due to frictions on the expansion of business and social networks (Combes et al., 2005).

Taken together, shipping costs between residential locations and hubs $\zeta()$ and transport costs between hubs $\tau()$ are responsible for the force of agglomeration in the model. First, more centrally located trading hubs will attract a greater number of workers. Next, workers will prefer to live closer to their designated hub. In equilibrium, these forces of agglomeration are counterbalanced by idiosyncratic location preferences of workers.

1.2.3 Equilibrium

In this section I define what constitutes a spatial equilibrium in the model and present the equilibrium conditions that characterize its structure. In the model, amenities consumed by the worker in her residential location do not depend on the associated trading hub. In addition to this, workers' utility is additively separable in amenities and consumables. Thus in equilibrium all workers who live at the same residential location r will choose the same trading hub. Let me denote this trading hub by $\mu(r)$. Given this, I define a spatial equilibrium of the economy below.

Given a geography G, parameters $\sigma, \theta, \phi, \psi, \beta, \alpha, \bar{L}$ and functions a: G $\rightarrow$R$_+$, $\{\tau, \zeta\} : S^2 \rightarrow \mathbb{R}_+$ an equilibrium of this economy consists of a population distribution $L : G \rightarrow \mathbb{R}_+$, consumption levels $c : [0, \bar{L}] \times G \times [1, ..., M] \rightarrow \mathbb{R}_+$, goods prices and levels of production $p, x : [0, \bar{L}] \times [1, ..., M] \rightarrow \mathbb{R}_+$; and a function that assigns each residential location a trading hub $\mu : G \rightarrow [1, ..., H]$, such that following conditions hold:

1. *Workers choose their consumption, production, price, residential location and trading hub to maximize their utility subject to the production technology and their budget constraint.*

2. *The market clearing for each good clears at every hub, implying*

$$x_h^i = \sum_o \tau_{h,o}^{1-\sigma} (p_h^i)^{-\sigma} P_o^{\sigma-1} p_o L_o \tag{1.6}$$

for any worker i, where x_m^i denotes the workers level of production, h denotes the hub of choice where she sells her product and P_o, L_o denoting the price index and total number of workers trading at destination hub o.

In equilibrium, the number of people residing in location r is given by

$$\log L(r) = v + \theta^{-1} \left[a(r) + \zeta(\mu(r), r)^{-1} MA_{\mu(r)}^{\frac{2\sigma-1}{\sigma(\sigma-1)}} \right] \tag{1.7}$$

where L(r) is the population of location r, v is a constant determined by parameters of the model and MA_m is the market access of trading hub h. Population at r increases with the level of amenities enjoyed at the location $a(r)$ and the market access provided by the associated trading hub $MA_{\mu(r)}$. Provided that shipping costs $\zeta(\mu(r), r)$ increase with distance, equation 1.7 implies that cities with a negative population gradient around their center will form around hubs.[5]

[5] Differentiating equation 1.7 with respect to distance to chosen hub yields $\frac{\partial \log L(r)}{\partial dist_{\mu_r, r}} = -\psi\theta^{-1}\zeta(\mu_r, r) MA_{\mu(r)}^{\frac{2\sigma-1}{\sigma(\sigma-1)}}$ which is strictly negative for any positive value of ψ.

Given their location choice, workers prefer the hub that offers them the best combination of proximity and market access.

$$\zeta(\mu(r), r)^{-1} MA_{\mu(r)}^{\frac{2\sigma-1}{\sigma(\sigma-1)}} \geq \zeta(\mu_m, r)^{-1} MA_m^{\frac{2\sigma-1}{\sigma(\sigma-1)}} \quad \forall m \tag{1.8}$$

Market access of any trading hub h is implicitly defined by

$$MA_h = \sum_o MA_o^{-\frac{\sigma-1}{\sigma}} \tau_{h,o}^{1-\sigma} L_o \tag{1.9}$$

implying that, trading hub h has better market access if it has other large hubs (high L_o) in close proximity (low τ_{ho}). Appendix A.1.2 also shows that the level of real income $\omega_h = p_h/P_h$ at a trading hub is an increasing function of its market access:

$$\omega_h = MA_h^{\frac{2\sigma-1}{\sigma(\sigma-1}} \tag{1.10}$$

Finally, the size of each trading hub h is given by the total number of people who choose to trade there.

$$L_h = \sum_{r:h=(r)} L(r) \tag{1.11}$$

When taken together, equations 7-11 govern the spatial distribution of economic activity in the model. In the presence of shipping costs that increase with distance, workers prefer to live close to trading hubs that are centrally located and have good market access. In equilibrium, this force of agglomeration is balanced by idiosyncratic location preferences, which preclude all workers from locating at the same hub.

Given the set of equilibrium conditions, I construct a model implied metric for urban primacy, analogous to the urbanization index in Nagy (2020). I use this metric to demonstrate how an increase in trading opportunities fosters urban growth. In particular, I define the *urban-primacy index* at location r as the gradient of log population with respect to proximity to its trading hub.

$$UI(r) = \frac{\partial logL(r)}{\zeta(\mu_r, r)^{-1}}$$

If the UI is large, then the gradient of population distribution is steep, indicating that population increases rapidly as one moves closer to the center of the trading hub. Using equation 1.7, I can show that the urban-primacy index increases with the market access of the trading hub.

$$UI(r) = \theta^{-1} MA_{\mu(r)}^{\frac{2\sigma-1}{\sigma(\sigma-1}}$$

This result indicates that hubs with better market access will generate more population concentration around them. It is a consequence of the trade-off between agglomeration and dispersion forces that are present in the model. To save on shipping costs, workers prefer to live closer to their chosen hub. This predilection is counterbalanced by their idiosyncratic location preferences. Naturally, the incentive to live closer to the trading hub is much stronger around hubs that have better market access. For this reason, population distribution is more concentrated around hubs that are more centrally located.

1.3 Structural Estimation

1.3.1 Data

This section provides a brief description of the data used in my empirical and quantitative analysis. First, in order to use the algorithm that selects hubs, calibrate the model and test its predictions regarding economic geography, I need fine-grained population data for each country in the region. Second, to estimate the intensity of border related barriers I need data on international trade flows in Europe and a complete set of bilateral controls that are customary in gravity specifications. Third, in order to provide supporting evidence on the link between trade and amenities, I need data on housing prices, wages and exports at the city level. Since data on exports are not available at the city level, I instead construct a suitable metric that measures the change in export exposure using sectoral employment data.

For population figures, I use the census data provided by Eurostat at the municipality level. This dataset records the population at the start of each decade between 1961 and 2011 for the entire set of European municipalities. This paper focuses on the 8 countries that are contiguous to the border between the NMS and the EU15.

In estimating border effects, I rely on two different datasets depending on the period in question. For the post-enlargement period, I use the regional trade flow data constructed by Santamaría et al. (2020), based on the European Road Freight Transportation Survey prepared by Eurostat. Unfortunately regional trade data is not avaialble for the pre-enlargement period. Therefore I use country level trade data from COMTRADE and bilateral control variables from CEPII.

For the analysis of trade related amenities, I combine two main data sources. First, the data on labor composition of German cities come from the German Regional Statistics Office.[6] This dataset contains information on sectoral employment across 2-digit WZ93 industries at the NUTS3 (district) level. Second, information regarding the housing market at the city level comes from Eurostat.[7] I combine the two datasets by constructing a correspondence between German districts (NUTS3) and German cities defined by Eurostat. For the majority of German cities, the boundaries of the city correspond to those of a single NUTS3 region, which allows me to retain the majority of the cities in the combined dataset.

[6]Dataset: 42111-01-02-4

[7]Dataset: urb$_c livcon$

1.3.2 Determining Hubs

A key assumption of the model is that trade is mediated by only a subset of locations, which I refer to as trading hubs. This assumption is based on the observation that trade is concentrated in large, centrally located cities that act as hubs for their respective hinterlands (Marin et al., 2020; Mori and Wrona, 2021). Identifying the right trading hubs is paramount for the empirical analysis as well as the quantitative exercise since the theoretical predictions of my model indicate that we should observe increased city growth only in a subset of locations in which trade is concentrated and that act as hubs.

Before determining the subset of cities that are trading hubs in Europe, one first needs to identify the set of all cities. Here I replicate Eurostat's methodology in classifying European municipalities into three broad categories in terms of their degree of urbanization, namely 1) Cities, 2) Towns and suburbs and 3) Rural areas. In a subset of countries, municipal boundaries coincide with those of the core city.[8] For this set of countries, I use municipalities in the "Cities" category as the set of candidates for trading hubs. For Italy, Poland and Slovakia, it is possible to aggregate multiple municipalities to form "Greater Cities" for a small number of locations.[9] When applicable, "Greater City" captures the true urban core of these locations. Therefore for this subset of locations, I aggregate municipalities to form the "Greater City".

In order to select the set of cities in my sample that will be classified as hubs, I employ an algorithm developed by Mori et al. (2020), who propose a simple classification process that identifies central hubs and their hinterlands assuming a hierarchical city system. The algorithm successively partitions a country into multiple (K) regions, with each region centered around one of the K-largest cities contained within. By specifying the number of lower-layer

[8] Germany, Slovenia, Hungary, Czech Republic

[9] Milan (IT), Naples (IT), Katowice (PL), Bratislava (SK) and Kosice (SK)

hubs K (here K=3) and a final stopping rule, one can construct a hierarchical city system using only data on population.

I assume that such a hierarchical city system exists with multiple layers within each country. The 0th layer of this system, centered around the country's largest city, is the entire country. Next, the first layer becomes the unique Voronoi-3 partition of the country, with each partitioned region centered around one of the 3 largest cities in the country.[10] The second layer consists of the successive Voronoi-3 partitions of each individual region generated by the first layer, this time centered around the three largest cities within each region outlined by layer-1. Thus at the k^{th} layer, the algorithm will deliver $3^{(k-1)}$ hubs along with their respective regions (hinterlands). Without an additional stopping rule, this process stops if there are not enough cities left to generate the next layer. After I generate this multi-layered system for each country I introduce a selection rule, and assume that in order to be considered as a trading hub, a city must have a population that is greater than 150,000. This population threshold is consistent with the reduced form evidence I present in the next section regarding the presence of trade related amenities. Furthermore, it allows me to construct a consistent sample of hubs in terms of size across this very heterogeneous set of countries.

Figure 1.1 exhibits the partition process of Germany for the 1st and 2nd layers. In part A, we have the first layer of hubs in the country, which corresponds to the 3-partition of Germany, with cells centered around Berlin, Hamburg and Munich. Layer 2 is generated by dividing each of these cells into 3, based on the location of largest cities contained within them, as displayed in Figure 1.1 part B.

[10] A Voronoi partition is a partition of any plane into regions based on the proximity to the elements of a given set of objects. In my exercise, these objects are a finite set of points (cities) located on a plane (country area). For each city, there is a corresponding region, called the Voronoi region, consisting of all points on the plane that are closer to that city then any other.

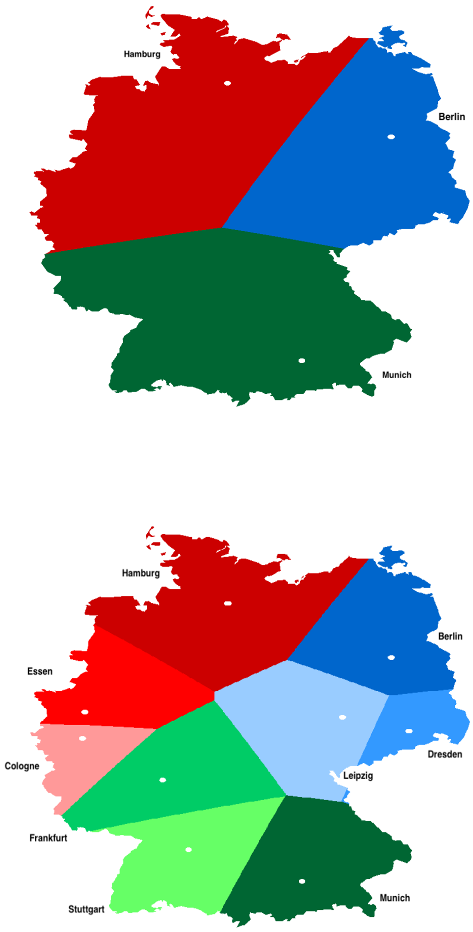

Figure 1.1: Spatial hierarchical 3-partition of Germany: (A) Layer-1 with 3 hubs (B) Layer-2 with 9 hubs.

Using this algorithm for the entire sample of countries generates the distribution of hubs presented in Figure 1.2. A complete list of selected hubs can be found in the Appendix. For smaller countries in the sample such as Austria, Hungary, Slovakia and Slovenia, the algorithm stops after Layer-1. For Germany, Poland and Italy, the algorithm is able to generate a full hierarchical system up to Layer-3. Finally for the Czech Republic, it generates a two layer city system. In total, 73 hubs are selected from a sample of 304 cities.[11]

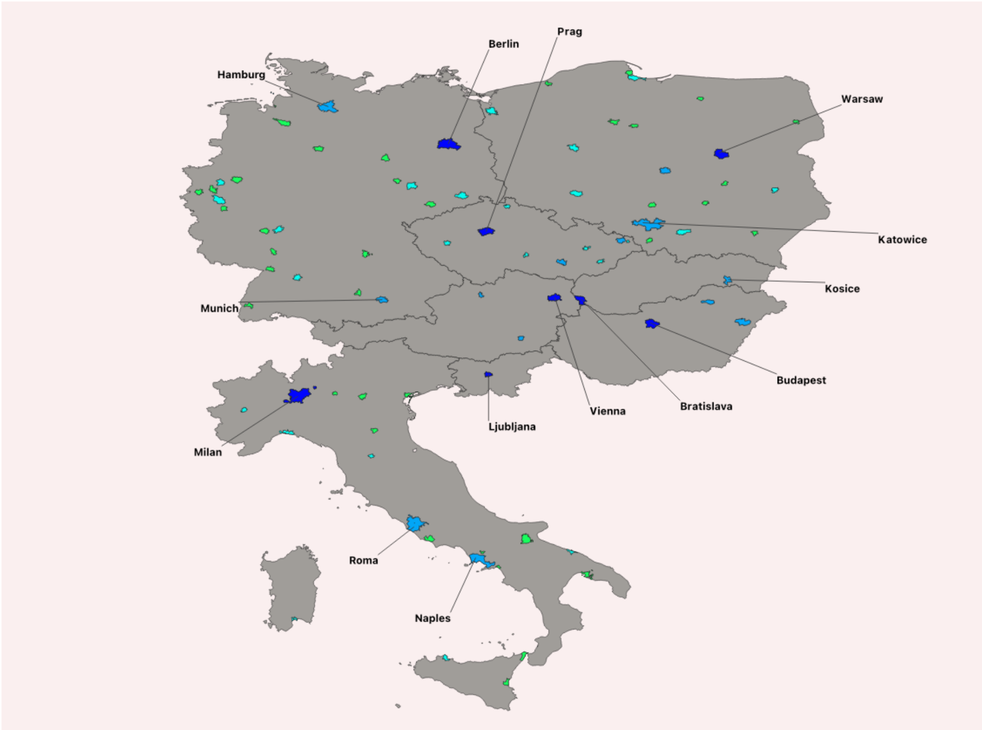

Figure 1.2: Urban Hubs

1.3.3 Model Geography

I use a fine discretization of space when setting up the model geography. I employ a 0.01° by 0.01° spatial grid of the territory that contains the three EU15 countries and five NMS countries that lie along the accession border. I consider each cell that lies within the boundaries of these 8 countries as a location. Having defined the set of all locations, I map hubs on the grid. First, I compute the coordinates of the geographic center and the total area of all hubs using shapefiles provided by Eurostat. Next, I assume that each hub has a circular shape and distribute them around their geographic center, where I equate the area of each circle to the area of the hub coming from the

[11]Some cities in Germany, Italy and Poland and Czech Republic are dropped from the sample of identified hubs after imposing the population threshold.

data. I compute hub-to-hub distances using the *georoute* STATA package developed by Weber et al. (2021).[12] For hub-to-location distances, I use flyover distances computed over the grid. Given the total number of hubs in the sample, this simplification does not alter the partition of the geography into hubs in a significant way.

1.3.4 Calibration of the Structural Parameters

In choosing the value for elasticity of substitution σ, I follow the previous literature in Simonovska and Waugh (2014) and set a value of $\sigma = 5$. I choose the value of ψ, which captures the utility cost paid by workers for shipping goods between residential locations and trading hubs, based on Nagy (2020) and set it to $\psi = 0.102$.

The parameter that governs the heterogeneity in location preferences, θ, plays a central role as it influences the effects of trade on urban primacy. A key property of the model is that, the elasticity of population with respect to income is not constant, and varies with the level of income ω.

$$\frac{\partial \log L(h)}{\partial \log \omega_h}(\bar{\omega}_h) = \frac{\partial \log L(h)}{\partial \log \omega_h}(\bar{\omega}_h) \cdot \frac{\partial \omega_h}{\partial \log \omega_h}(\bar{\omega}_h) = \frac{\partial \log L(h)}{\partial \log \omega_h}(\bar{\omega}_h) \cdot \bar{\omega}_h = \theta^{-1}\bar{\omega}_h$$

meaning that a one unit change in real income at a trading hub h, leads to a θ^{-1} change in its log population $\log L(h)$. Furthermore, the elasticity of population with respect to income increases with the level of income as well. This implies that, ceteris paribus, effects of economic integration will be felt the most in large hubs. Since the elasticity of population with respect to real income is not constant in the model, I choose the value for θ based on the average income elasticity computed across all hubs. The literature hasn't converged on a single value for the elasticity and provides various estimates for late-20^{th} century Brazil (1.91) (Morten and Oliveira, 2018), present day

[12] Exact specifications are provided in the data Appendix.

China (2.54) (Tombe and Zhu, 2019) and for present day United States (3.30) (Monte et al., 2018). Therefore I choose a value of $\theta = 15.25$, which yields an average income elasticity of population to real income that is equal to 2.49, a value that is equal to the average of prevailing estimates in the literature.

To compute the elasticity of cross-hub transport costs with respect to distance (ϕ), I estimate the gravity specification that is delivered by the model, using the ERFT data on truck shipments across European regions. Bilateral trade between hubs h and o is given by

$$T_{ho} = \tau_{ho}^{1-\sigma} p_h^{1-\sigma} L_h P_o^{\sigma} p_o L_o \tag{1.12}$$

where T_{nm} denotes the total volume of trade.

$$\log T_{ho} = \beta_1 \log dist_{ho} + \beta_2 B_{ho} + \delta_h + \delta_o + \epsilon_{ho} \tag{1.13}$$

Given this expression, one can estimate the specification given in Equation 1.13 where the coefficient β_1 will be equal to the elasticity of trade with respect to distance. After estimating β_1, one can compute ϕ for a given value of σ since we have that $\beta_1 = (1-\sigma)\phi$. I find that $\phi = 3.44 \cdot 10^{-4}$.[13]

1.3.5 Trade Related Amenities

The urban economics literature has long established the role of consumer amenities in stimulating urban growth (Rappaport, 2007) and highlighted their growing role in determining agent's location choices (Rappaport, 2008). In the words of Glaeser et al. (2001), *"sovereignty of the consumer is inescapable"* over stimulating growth in cities. Cities benefit from the presence of a rich variety of services, brands, activities as they host a more diverse range of restaurants and cuisines, live performance venues and professional sports. In addition, they also contain an attractive blend of social partners

[13] Here I retain the assumption that $\sigma = 5$.

coming from different backgrounds. Prior literature has emphasized how city size and density engender such consumer amenities. In this paper, I focus on a different underlying determinant of consumer amenities in cities: trade. Trade makes a city more cosmopolitan and facilitates its development as a center of the arts, sciences and other cultural activities as well as enhancing its capacity to attract a greater variety of brands and services. Thus, trade generates positive spillovers in the form of consumer amenities within cities.

In the rest of this section, I expand on the strategy I follow in calibrating the trade elasticity of endogenous consumer amenities, α, which is the key structural parameter that links trade and amenities in the model. I start by specifying a functional form for the total location specific amenities, $a(r)$. Using the fact that location specific amenity function $a(r)$ can take on any form, I look for the functional form that would enable the model to match exactly the city and country population levels observed in the data. To this end, I assume the following structure for amenities:

$$a(r) = \begin{cases} a_h + a_C & \text{if } r \in h \\ a_C & \text{if otherwise} \end{cases} \tag{1.14}$$

implying that location specific amenities $a(r) = a_h$ if location r is within the boundaries of a hub h and $a(r) = a_C$ otherwise. Here, a_h can be viewed as the total level of amenities a hub provides to its residents. Whereas a_C are country-specific amenities that are common across all residential locations within the same country.

If hub-specific and country-specific amenities were treated as freely time-varying parameters, it would be possible to fit perfectly the population of all hubs and countries both before and after enlargement. This auxiliary calibration, in which I solve the model separately for the pre-enlargement (2001) and post-enlargement (2011) periods, imply changes in total hub amenities that are positively correlated with model-implied changes in each hub's trade volume. This finding supports the assumption that amenities are endogenously rising with trade, and allows me to calibrate their elasticity through a moment condition akin to the ones used by Ahlfeldt et al. (2015). With the estimated model at hand for both periods, I can generate a decomposition of the changes in total hub amenities into its exogenous ($\bar{a}_h$) and endogenous ($\tilde{a}_h$) components for any given value of α. Here, I employ a moment condition that parametrizes endogenous amenities so that any change in the residual, exogenous component of amenities is orthogonal to the change in trade volumes induced by the exogenous decline in border related trade barriers. More precisely, the moment condition is given by

$$corr[\Delta\bar{a}_h, \Delta Trade_h^{E,W}] = 0 \tag{1.15}$$

where $\Delta\bar{a}_h$ denotes the change in the exogenous component of city amenities and $\Delta T_c^{E,W}$ denotes the change in aggregate trade that passes through the EU15-NMS border at the hub level between 2001 and 2011.[14] The moment condition in Equation 1.15 states that, by being exogenous at the hub level, the systematic effects of EU Enlargement on trade must be uncorrelated with the changes in exogenous component of amenities in hubs. This identifying assumption requires that the systematic changes in hub amenities is explained by the mechanisms present in the model rather than by changes in the residuals.

Figure 1.3 displays the evolution of the correlation for increasing values of α starting from the benchmark case of $\alpha = 0$ in which trade does not generate amenities in hubs. For $\alpha = 0$, $corr[\Delta\bar{a}_m, \Delta Trade_m^{E,W}] > 0$, indicating that on average, hubs that are closer to the accession border saw a greater increase in total amenities.[15] As the value of α increases, $corr[\Delta\bar{a}_m, \Delta Trade_m^{E,W}]$ goes

[14] Here $\Delta Trade_c^{E,W}$ is equal to the change in total trade done with Eastern cities if the city c is located in the West (Austria, Germany and Italy) and is equal to the change in total trade done with Western cities if the city c is located in the East (Czech Republic, Hungary, Poland, Slovenia and Slovakia).

[15] Although the reduction in border related barriers increases trade between all EU15-NMS hub pairs, this effect is largest for those hub that are close to the border. Note that when $\alpha = 0$, $\Delta\bar{a} = \Delta a$. Combining these two observations, $corr[\Delta\bar{a}_m, \Delta Trade_m^{E,W}] > 0$ imply that increase in overall amenities was larger for hubs that are close to the accession

down, allowing me to identify the elasticity as $\alpha = 0.133$.

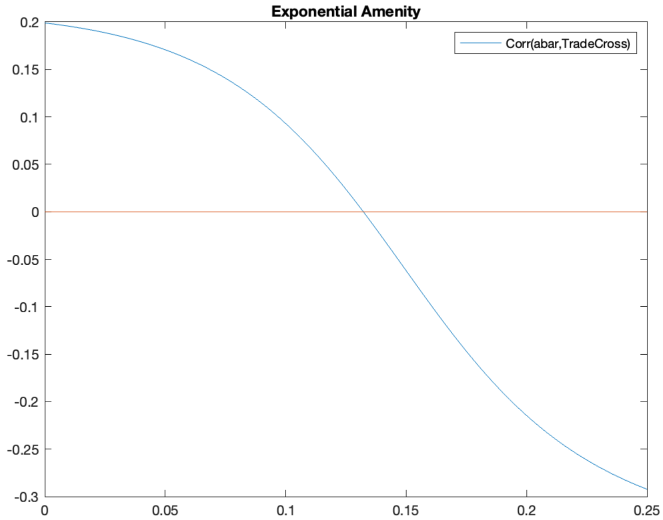

Figure 1.3: Estimation: Elasticity of Trade Driven Amenities

In order to provide supportive empirical evidence for the amenity-inducing effect of trade, I investigate to what extent, differential growth in housing prices vis-a-vis wages depends on the exporting capabilities of a city, following a positive trading shock. In the presence of trade related amenities, a positive trade shock would make cities more amene, and cause housing prices to grow more relative to the increase in wages. To test this theory in the data, first I compute the percentage growth in housing prices $\Delta Rent_c$, as measured by rents, and the percentage growth in wages $\Delta Wage_c$ for German cities. Then I construct my dependent variable Y_{ct} as the difference between the growth in housing prices and wages, or $Y_{ct} = (\Delta Rent_c - \Delta Wage_c)$. Unfortunately trade data at the city level is not available for Europe. Therefore to measure the export exposure of German cities to NMS economies I follow Dauth et al. (2014), and compute what is referred to as the exports impact parameter ExI_c for each city c in Germany. This measure captures a city's potential to benefit from a positive demand shock coming from Eastern

border

European countries based on its initial sectoral employment composition.

$$ExI_{ct} = \sum_j \frac{L_{cjt}}{L_{jt}} \frac{\Delta EX_{jt}^{EAST}}{L_{ct}}$$

where ΔEX_{ct}^{EAST} is the total change in exports from Germany to the East in industry j between time periods t and $t+1$ (in constant Euros of 2004), L_{cjt}/L_{jt} denotes city c's share of national industry employment in industry j in period t and L_{ct} is the total manufacturing employment in city c in period t. In order to control for any potential reverse causality, I use lagged sectoral employment shares from 2001 [16] and thus have

$$ExI_{ct} = \sum_j \frac{L_{cjt-1}}{L_{jt-1}} \frac{\Delta EX_{jt}^{EAST}}{L_{ct-1}} \tag{1.16}$$

Having constructed the export impact measure ExI at the city level, I estimate the effect of an increase in export exposure on the relative growth of housing prices vis-a-vis wages using the following specification:

$$Y_{ct} = \beta_1 ExI_{ct} + \beta' X_c + \epsilon_{ct} \tag{1.17}$$

In essence, I relate the relative growth of housing prices vis-a-vis wages between 2004 and 2014 Y_{ct} to changes in potential export exposure to Eastern Europe during the same period, while controlling for some pre-period control variables at the city level X_c. In the baseline specification, for which the results are presented in Table 1.1, I control for available dwellings per capita at the city level in order to control for potential confounding effects related to city size and housing supply. Column 1 presents the results for the full sample of German cities. Here the size of the coefficient on ExI is very small and it is not significant, indicating that, there is not a discernible link between export exposure and differentially higher growth in housing demand

[16] If employment reacted to anticipated future trade, using contemporaneous employment shares could lead to overestimation.

for the average German city. At a first glimpse, this result may look a little disheartening, however it is consistent with the assumptions of the model regarding the presence of trade related amenities. Not all cities in Germany can be considered as major trading hubs. Therefore the absence of a strong link between trade and amenities is not surprising here. Next, I estimate the baseline specification for the subsample of cities that were selected by the algorithm as hubs. These cities make up the upper echelons of the hierarchical city system where trade is concentrated. Results are presented in column 2, where the coefficient of interest is now much greater in magnitude and also significant at the 95% level. Accordingly, within this sample of 19 German cities which were selected as hubs, a €1,000 euro increase in per worker export exposure corresponds to a .6 percentage points higher growth in housing prices relative to wages.[17] This implies that hubs that benefited more from the increase in demand for German exports coming from Eastern Europe following the 2004 enlargement of the EU, also experienced a more than proportional increase in demand for their housing relative to income.

Table 1.1: Relative Growth in Housing Prices vis-a-vis Wages Across

	All Cities	Hubs
	Δ Rent - ΔWage	Δ Rent - ΔWage
Export Impact (2004-2014)	0.00651	0.597**
	(0.116)	(0.272)
Dwellings/Capita (2001)	-20.44	-44.45
	(31.42)	(39.81)
Observations	64	19
R^2	0.008	0.307

Standard errors in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

[17]In total, 24 German cities are selected as hubs. However due to data limitations only 19 of them are contained within the sample here.

1.3.6 The Evolution of Border Costs in Europe

I model the 2004 Enlargement of the EU as a reduction in border related trade barriers. Therefore, to measure the effects of EU Enlargement, I need an accurate assessment of border related barriers that prevailed across countries before and after 2004. I allow for heterogeneity in border effects based on country blocs. I estimate three different border effects, β^{WW}, β^{EE} and β^{EW}, that denote the average border effect between two EU15 countries, between two NMS countries and between one EU15 and one NMS country respectively.

In estimating border effects, I employ a traditional gravity specification with origin and destination fixed effects, given by

$$\log T_{nm} = \beta_0 + \beta_1 \log dist_{nm} + \beta_2 B_{nm} + \beta' \mathrm{X} + \delta_n + \delta_m + \epsilon_{nm} \tag{1.18}$$

where T_{nm} denotes total bilateral trade between origin n and destination h, $dist_{nm}$ denotes the distance in km between these locations, B_{nm} is the border dummy that takes on the value 1 if origin and destination are located in different countries, X denotes the matrix of bilateral control variables and δ_n, δ_m are origin and destination fixed effects respectively.

For the post-enlargement period, I use the dataset on regional trade constructed by Santamaría et al. (2020) in order to estimate the average border effect. This dataset is based on the ERFT Survey and contains region to region trade flows across 24 European countries at the NUTS2 level between years 2011 and 2017. When available, using regional trade flows leads to more accurate estimates of the border effect. Using this data, I first compute the average annual trade values and then estimate the specification given in equation (1.18). Unfortunately, trade data at the regional level is not available before 2004. Thus for the pre-enlargement level of border effects, I rely on international trade data from COMTRADE (1996-2003) where the geographic unit of observation is at the country level and follow the same empirical methodology.[18]

Table 1.2: Border Effects Before/After the Enlargement of EU

	West-West	East-East		East-West	
	(1)	(2)	(3)	(4)	(5)
	Before & After Accession	Before Accession	After Accession	Before Accession	After Accession
Border Effect	-1.496***	-2.970***	-2.255***	-3.408***	-1.243***
	(0.212)	(0.668)	(0.313)	(0.291)	(0.269)
Controls	Yes	Yes	Yes	Yes	Yes
Origin FE	Yes	Yes	Yes	Yes	Yes
Destination FE	Yes	Yes	Yes	Yes	Yes
N	25633	25	1311	38	6680
R-squared	0.72	0.96	0.81	0.97	0.85
Data Source	ERFT	Comtrade	ERFT	Comtrade	ERFT
Geo. Unit.	NUTS2	Country	NUTS2	Country	NUTS2

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Results are presented in Table 1.2. I set the value of the border effect between two EU15 countries to its post-2004 value and assume that it does not change. In my quantitative analysis, I assume that β^{WW} does not change after the enlargement. It is important to note that even if there was a change in the border effect that prevailed between EU15 countries, this can't be a consequence of the accession of NMS to the union. Therefore, when quantifying the effects of enlargement, I set the value of β^{WW} to its post 2004 value.[19] I find a modest reduction in the border effect between two NMS countries (β^{EE}). These countries were already members of CEFTA prior to their accession into EU.[20] This observation, combined with the likely western-biased trade diversion effect of the EU membership process can explain why changes β^{EE} were rather modest. Lastly, I estimate a sizable decrease in the border effect between EU15 and NMS countries (β^{EW}). As evidenced by columns 4 and 5 of Table 1.2, the average border effect between EU15 and NMS countries decreased from -3.408 to -1.243, implying that after becoming members of the EU, NMS experienced a 70% reduction in border related

[18] Set of bilateral controls as well as bilateral distances at the country level are gathered from CEPII.

[19] Estimation of border effects with regional trade yields more accurate results as it corrects for a number of biases that are endemic to country level data.

[20] Poland, Czech Republic, Slovakia, Slovenia and Hungary collectively left CEFTA and became members of the EFTA following the accession.

trade barriers in tariff equivalent units. The considerable reduction in border related barriers captures the full extent of the accession period and is in line with other estimates in the literature.

1.4 Model Predictions and Reduced Form Empirical Evidence

In this section, I expand upon the predictions of the model within the context of EU enlargement and explore to what extent these predictions are reflected in the data. In my reduced form empirical analysis, I do not impose the structure of the model on the data.

The first main prediction of the theoretical model is that the enlargement of the EU, by reducing border related barriers on trade, will lead to an increase in regional urban primacy at the aggregate level. As trade costs decline, people will move closer to their trading hub in an attempt to reap the benefits of increased trading opportunities. Next, this effect on urban growth will be heterogeneous across space. In particular, the increase in concentration of urban population will be greater around hubs that are located near the accession border since they experience the largest gains in terms of market access. Finally, regions near the border would experience a greater increase in population growth, given that realized real income gains are larger.

To investigate whether the predictions of the model are supported by the data, I adopt a difference-in-differences strategy. First, I compare the population growth trends of municipalities located close to the accession border with the population growth trends of other municipalities located farther away. Following Brülhart et al. (2012), in order to study the effects of EU enlargement, I undertake this comparison before and after the accession of Eastern European countries into the union. My baseline empirical specification is given by:

$$Pop.Growth_{it} = \beta_1(B_{X,i} \times A_t) + \alpha_i + \alpha_t + \epsilon_{it} \tag{1.19}$$

where $Pop.Growth_{it}$ is the annualized population growth rate over the periods 1991-2001 and 2001-2011 in municipality i at time t, $B_{X,i}$ is a dummy that is equal to one if the municipality i is located within the first X kilometers of the accession border, A_t is a dummy that is equal to one for the post-enlargement period (2001-2011) and α_i, α_t are municipality and time period fixed effects respectively. To allow for serial correlation of the error term ϵ_{it} within geographic regions, standard errors are clustered at the NUTS3 level. Here the coefficient of interest is β_1, which measures whether population growth trends evolved differently in border regions.

Next, to test the predictions of the model regarding urban growth, I augment the baseline empirical specification as follows:

$$Pop.Growth_{it} = \beta_1(B_{X,i} \times A_t) + \beta_2(A_t \times Hub_i) + \beta_3(B_{X,i} \times A_t \times Hub_i) + \alpha_i + \alpha_t + \epsilon_{it} \tag{1.20}$$

where Hub_i is a dummy that is equal to one if municipality i is designated as a hub by the algorithm described in Section 3.2. In this specification, the coefficient on the double interaction between accession and hub dummies, β_2, measures how population growth trends of hubs evolved differently after the accession of new member states. Additionally, the coefficient of the triple interaction, β_3, will capture whether population growth evolved differently in hubs that are closer compared to their counterparts that are located farther away. In order to assess the geographic extent of these effects, I estimate both specifications for different border dummies $B_{X,i}$, varying in terms of distance to the accession border that determines the treatment group.

Table 1.3: The impact of EU Enlargement on Population Growth Trends

	Dep: Annualized Pop. Growth		
	(1)	(2)	(3)
	50Km	100Km	150Km
Border x Accession	0.281**	0.487**	0.371**
	(0.121)	(0.192)	(0.175)
Year FE	Yes	Yes	Yes
Location FE	Yes	Yes	Yes
Observations	72514	72514	72514
R^2	0.618	0.622	0.621

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 1.3 presents the results pertaining to the first specification. As evidenced by columns 1, 2 and 3, following the accession of NMS into the union, municipalities located within close proximity of the accession border experienced a larger increase in their annualized population growth rate relative to other municipalities located farther away. This differential increase in population growth rates is largest for the municipalities located within the first 100km of the border (0.487).

Results of the second specification are presented in Table 1.4. First, the coefficient of the double interaction between accession and hub dummies, is positive and significant at the 99% level across all three columns. These results indicate that, in the wake of the 2004 EU enlargement, hubs experienced approximately a .4 percentage points larger increase in their annual population growth relative to other municipalities. Furthermore, this effect is even stronger for hubs that are located within very close proximity of the accession border. The coefficient of the triple interaction term in column 1 implies that hubs that are located within 50 kilometers of the accession border, experienced on average a .5 percentage points larger increase in their annualized population growth relative to other more distant hubs. Overall, these results show that theoretical predictions of the model are consistent with the population changes observed in the data.

Table 1.4: The mpact of EU Enlargement on Population Growth Trends

	Dep: Annualized Pop. Growth		
	(1)	(2)	(3)
	50Km	100Km	150Km
Border x Accession	0.281**	0.488**	0.371**
	(0.121)	(0.192)	(0.175)
Accession x Hub	0.384***	0.471***	0.434***
	(0.124)	(0.103)	(0.101)
Border x Accession x Hub	0.523*	0.245	0.301
	(0.295)	(0.289)	(0.282)
Year FE	Yes	Yes	Yes
Location FE	Yes	Yes	Yes
Observations	72514	72514	72514
R^2	0.618	0.623	0.621

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

1.4.1 Robustness

The previous section presented three important results. First, after the 2004 enlargement of the EU, municipalities that are located closer to the border between the NMS and the EU15 experienced a larger increase in their population growth rates. Second, on average, hubs saw a significant increase in population growth in the wake of the enlargement process. Finally, this increase was stronger for hubs that are located closer to the accession border. In this section, I discuss several robustness exercises to support these empirical findings, whose results are reported in Appendix A.1.

In order to investigate heterogeneity of the effect among hubs depending on distance from the border, I augment the baseline specification by introducing a series of exclusive distance band dummies ranging from 0-50 kilometers to 50-150 kilometers. I restrict my sample to contain only cities and also include the double interaction between accession and hub dummies to check whether all hubs started to grow faster than non-hub cities after the enlargement. Results are presented in column (1) of Table A.1.1. The coefficient on the double interaction between accession and hub dummies is slightly negative and insignificant, implying that not all hubs outperformed non-hub cities after 2004. The estimated coefficients on the triple interactions for 0-50 kilometers and 50-100 kilometers are positive and statistically significant, while the estimated coefficient on the interaction for 100-150 is positive but not significant. These results are consistent with the predictions of the model and the results of the previous specifications. Next, in order to check whether similar patterns were in effect before 2004, I remove fixed effects and introduce double interaction terms between the hub dummy and distance band dummies ranging from 0-50 kilometers to 50-150. Results are presented in column (2) of Table A.1.1. In contrast to the positive and significant coefficients on the triple interaction terms, the coefficients on the double interactions for 0-50 kilometers and 50-100 kilometers are negative and statistically significant. This indicates that, in the decade prior to the enlargement, hubs near the border experienced a decline in their growth trends. However this decline stopped and these hubs experienced larger increases in population growth following the enlargement process.

As a final robustness check, I turn to data from regional data from Poland's national statistics office and show that Polish regions close to the German border experienced larger employment gains after 2004. Polish data is special in the sense that, for each NUTS3 region, both employment and urban employment figures are calculated by the official statistical office. This allows me to separately estimate the treatment effect on the urban fraction of employment and infer whether urban concentration of employment increased relatively more near the border area. Results are presented in A.1.2. The coefficients on the triple interaction for the 0-50 kilometers and 50-100 kilometers are positive for both overall employment (column 1) and urban employment (column 2) and the effect dissipates after the first 100 kilometers. This implies that NUTS3 regions near the border experienced both a higher growth in overall employment and urban employment.

1.5 Measuring the Effects of 2004 EU Enlargement

In this section I present the results of the structural estimation. In section 5.1, I describe the quantitative exercise I perform in order to isolate the effects of the enlargement process alone. In section 5.2, I explore the model's fit to municipality level population data from the 2001 and 2011 censuses. In section 5.3, I present model predictions regarding trade, urban growth and real income. In section 5.4, I analyze the effects of EU enlargement on spatial inequality. Finally, Section 5.5 presents the decomposition of income and urban primacy effects between different channels that are present in the model.

1.5.1 Isolating the effect of 2004 enlargement of EU

The purpose of this quantitative exercise is to measure the effects of EU enlargement, as seen through the lens of increased market integration, on urban primacy and welfare. To this end, I model the integration process as a reduction in border related barriers that hinder trade between NMS and EU15 countries. Accession of new member states into the union is a multi-faceted process that is likely to have affected the economic geography through various other channels. In addition, there were other economic and demographic forces at play during this period that had an effect on the spatial distribution of economic activity as well. One example is the change in total population within and across countries in the sample. To measure the effects of enlargement via trade in isolation, one has to shut down other potential mechanisms when solving the model for the counterfactual. To this end, the quantitative exercise presented in this section employs the following strategy:

1. Solve the model for the pre-enlargement period, matching hub amenities and country specific amenities to 2001 census population data.

2. Compute the endogenous component of hub amenities related to trade.

3. Back out the exogenous component of hub amenities ($\bar{a}_h$).

4. Update border related barriers between NMS-EU15 (β^{EW}) and NMS-NMS (β^{EE}).

5. Solve the model for post-enlargement period, with reduced border related barriers, keeping the exogenous component of hub amenities ($\bar{a}_h$), country specific amenities ($\bar{a}_C$), total population in the entire sample and all of the structural parameters in the model fixed.

By keeping the total population in the sample fixed, I control for macro level demographic changes that don't stem from market integration and possibly affect urban growth. Keeping country amenities fixed controls to what extent workers will be allowed to move between countries, but does not shut down cross border labor mobility completely.[21] There were considerable restrictions to labor mobility even after 2004 between the EU15 countries present in my sample and the NMS countries. Germany and Austria did not ease pre-enlargement restrictions until 2011, while Italy did so only after 2006. Thus implementing a strict but not complete form of mobility friction in the model is a realistic way of capturing the policies in effect during this period. Keeping the exogenous component of hub amenities fixed controls for any residual change that could affect the population at the city level.

1.5.2 Model fit to municipality population data

In this section, I assess the ability of the model to fit the population distribution in the data, before and after the enlargement of the EU, as captured by the 2001 and 2011 censuses. I compute the model implied population of

[21] Here, matching country populations with the new set of border effects, and computing a new set of country specific amenities would amount to shutting down cross border mobility completely.

each municipality in the sample. I assume that each settlement has a circular shape and distribute them onto the grid centered around their geographic center. Then, I aggregate the total predicted population of grid cells that fall within the boundaries of each municipality, in both 2001 and 2011. Next, I calculate the change in population of each municipality, as implied by the model. As a final step, I calculate the correlation between population vectors implied by the model and census, in order to measure to what extent the model can fit the population distribution in the data.

Overall, the model's ability to match population levels is good. The correlation between model implied and real non-hub municipality populations are 0.47 in 2001 and 0.46 in 2011 respectively. Considering the very large number of settlements in the sample and the fact that only hub populations in 2001 are matched during calibration, the model can be said to capture, to a good extent, the spatial distribution of population across the geography in question.[22] Regarding percentage changes in municipality populations, the correlation between model implied values and census data is 0.26 for the sample of non-hub municipalities.

1.5.3 Effect of Enlargement on Regional Urban Primacy and Welfare

Figure 1.4 presents the map of real income gains experienced across the geography. All locations that trade in the same hub have the same gain in real income in percentage terms. On average, the model predicts a 0.21% increase in real income with a standard deviation of 0.31. The large standard deviation of income gains vis-a-vis its average indicates significant heterogeneity in real income gains across countries, as evidenced by Figure 1.4. On average, income gains are higher for small countries that become centrally located within this new more interconnected geography such as Slovenia (2.7%), the Czech Republic (0.61%), Slovakia (0.66%) and Austria (0.65%). On the other

[22]In total there are 31,538 non-hub municipalities and 73 hubs in the sample.

hand, large economies of the West such as Germany (0.12%), report relatively modest real income gains, with Italy (0.06%) benefiting from enlargement the least.

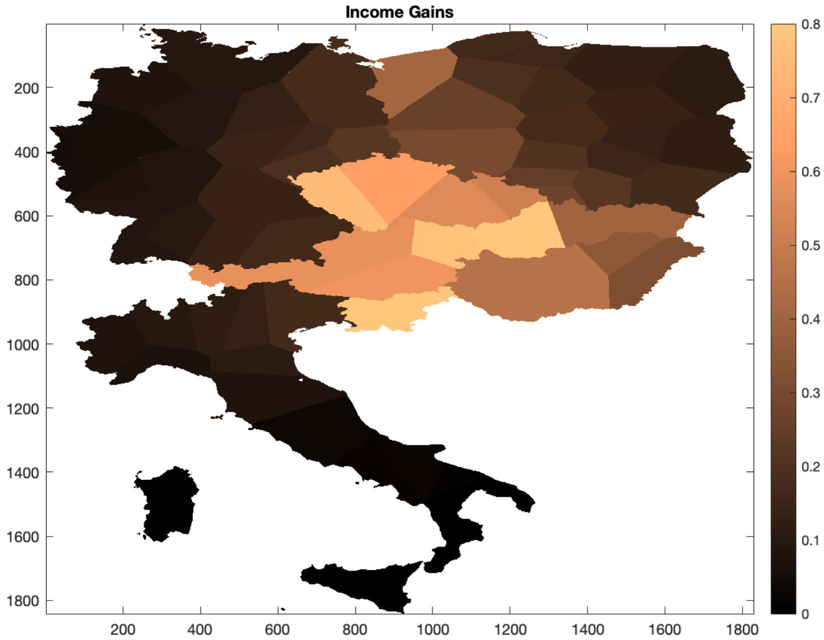

Figure 1.4: Map of Real Income Gains

At the aggregate level, the model predicts a 0.39% increase in the concentration of population living in urban hubs. According to census figures, this number is equal to 1.51%, meaning that the 2004 enlargement of the EU, by facilitating trade between new and old member states, can explain 26% of the increase in population in top urban hubs. Next, I calculate the aggregate increase in urban primacy, defined as the percentage of total population living in urban hubs. I find that urban primacy increases from 21.51% to 21.59%, or by 0.08 percentage points. Here, the model is able to explain 96% of the actual increase observed in the data.

In reality urban growth is not a phenomenon that is limited to the core parts of cities alone and propagates beyond the physical boundaries of the core city. This observation has pushed the literature to define concepts such as the greater city and the functional urban area (FUA) (Dijkstra and Poelman, 2014). Concept of greater cities, where applicable, captures the urban core of a city accurately. Unfortunately, they are defined only for a subset of locations in Europe. Just as municipal boundaries suffer from capturing too little of a city, functional urban areas, instead suffer from capturing too large an area.[23] For example, the FUA of Berlin, extends beyond the boundaries of multiple regional governments (landkreise); and covers an area that surely can't be considered as an urban hub in its entirety. In this paper, I constrain my definition of the city to its urban core as captured by the municipal boundary.[24] Therefore, when interpreting the increase in urban primacy predicted by the model, and its explanatory power over figures in the real data, I don't claim that model is able to capture reality in its entirety. Instead, due to its built-in limitations due to geographic definitions of a hub, the model is able to capture the portion of urban primacy that involves the very cores of cities. To support this claim, I project the FUA's of all the hubs in my sample onto the population grid and explore to what extent the model can explain demographic changes across these units. When measured over FUA's, the model predicts a 0.08 percentage points increase in urban primacy and is able to capture 16 percent of the effect observed in the data (0.49). The fact that increase in urban primacy is higher when calculated over FUA's supports my previous interpretation of results.

Figure 1.5 maps the geography of population change across countries that lie along the accession border. Increases in urban primacy are strongest in those regions that are close to the border, which is consistent with the reduced form empirical results presented in Section 4. The effects of EU enlargement, despite being very heterogeneous across space, imply an increase in urban

[23]Current standard in the literature includes municipalities that have at least 15% of their population commuting in a city, part of the FUA of that city. This threshold is too low for the purposes of the quantitative exercise conducted here since it captures way beyond the urban core of a city, where trade related amenities are present.

[24]Except in those cases in which a Greater City is defined.

primacy at the aggregate level, which is also consistent with the positive shift in population growth trends observed in the data. Here, the cases of Germany and Italy deserve special attention, as they provide valuable insights regarding the effect of enlargement on spatial inequality, which I discuss in the next section.

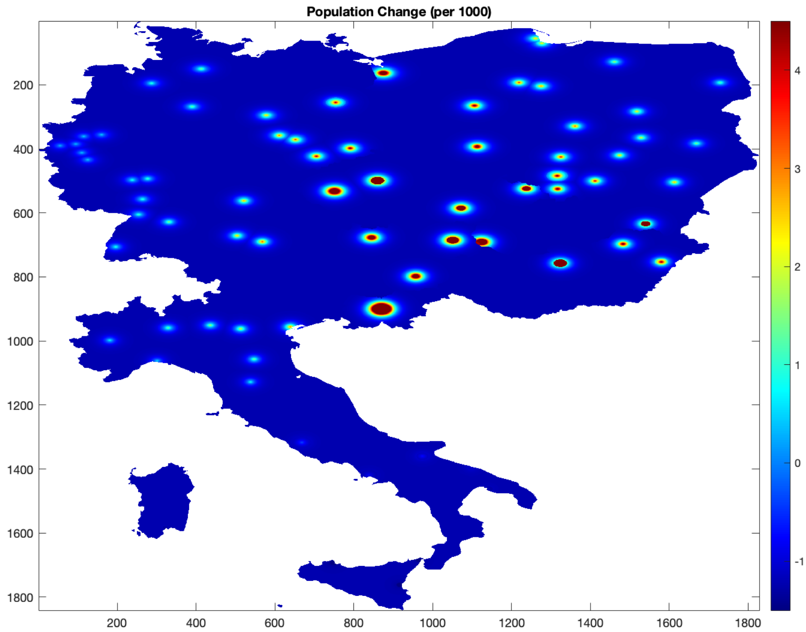

Figure 1.5: Map of Changes in Population

1.5.4 EU enlargement and spatial inequality

As evident in results presented in the previous section, on average, a reduction in border related barriers has a positive effect on real income levels. However, these effects are not uniform across space and therefore have effects regarding spatial inequality. The integration of new markets moves the economic center of gravity of the EU. The heterogeneity in shifts in market access across hubs affects their relative attractiveness as a trading hub.

In the case of Germany, previously peripheral hubs in the East became more central after the accession of Eastern European countries into the union. Thus, when compared to hubs in the West, they experience larger income gains. In particular, the model predicts an average income effect that is twice as large for East Germany (0.17%) compared to West (0.09%). This observation suggests that enlargement of the EU towards the East has had a progressive impact on regional inequality in Germany. Turning to Italy, the same mechanism has the opposite effect. When compared to the hubs in the North, southern hubs become even more peripheral in the post-enlargement period. As a consequence of this, income gains in the North (0.11%) are much greater than in the South (0.04%). Thus, integration of Eastern European economies into the single market exacerbates the spatial inequality present between north and south Italy.

1.5.5 Shutting down trade related amenities

In this section, I investigate to what extent the effects on urban growth and real income gains depend on the presence of endogenous amenities driven by trade. In order to decompose the total effect presented in section 4.3 into "market access driven" and "endogenous amenity driven" components I employ the following strategy. First, I solve the model for the pre-enlargement period, matching total hub amenities and country specific amenities to 2001 census population data. Then, I shut down the channel propagated by endogenous amenities by assuming that total hub amenities a_h is equal to the exogenous component alone $\bar{a}_h$.[25] Finally, I solve the model for the post-enlargement period, with updated border related barriers, keeping the exogenous component of hub amenities, country specific amenities, total population and rest of the structural parameters of the model fixed. The counterfactual equilibrium solved here is one in which hubs only benefit from a reduction in trade costs via changes in market access. I find that, in the ab-

[25]This is equivalent to modifying the definition of location specific component of amenities given in 1.4 as $a(r) = \bar{a}(r) + a_C(r)$.

sence of endogenous amenities, the model predicts a .19% increase in urban population and a .04 percentage points increase in urban primacy. Comparing these values with their counterparts presented in Section 4.2, I conclude that approximately half of the effect on urban growth predicted by the model is generated via the presence of endogenous amenities related to trade. On the other hand, the average increase in real income in the absence of endogenous amenities is 0.21%. This shows that the presence of endogenous amenities acts as an amplifying mechanism in determining population movements between hubs and their hinterlands. Yet, this mechanism is not strong enough to generate significant real income effects at the aggregate level.

1.5.6 Robustness

In this section, I establish the robustness of the main findings of the paper, namely that the 2004 enlargement of the EU led to a .39% increase in the aggregate population of urban hubs, a .08 percentage points increase in urban primacy and a .21% increase in average real income level in Central Europe. I consider alternative values for the model's structural parameters and estimate the trade elasticity of endogenous amenities α for these values. In particular, I individually allow for each of the structural parameters θ, ψ, ϕ and σ to vary by 10% and replicate the quantitative exercise in which I measure the effects of the 2004 enlargement of the EU for these values.

Table A.1.3 presents the point estimate for α as well as the average increase in income, aggregate hub population and urban primacy for each of the robustness exercises. Increasing the value of the taste heterogeneity parameter θ has negligible effects in terms of real income gains. Higher values of θ imply greater heterogeneity in idiosyncratic location preferences. Thus in equilibrium idiosyncratic location tastes start to dominate economic incentives and market integration induces less growth in hubs. It is possible to interpret the hub-to-residence shipping cost ψ as a commuting cost. Higher values of the shipping cost parameter ψ and the long distance transport cost parameter ϕ imply less trade between locations. Therefore, the increase in urban primacy is weaker in response to a positive trade shock since an increase in these parameters diminishes the potential gains from integration. A higher value of the elasticity of substitution (σ) decreases the gains from trade and I observe less urban growth. As substitutability across goods increases, locations rely less on trade with other places. Therefore market integration generates less gains and a weaker incentive to concentrate closer to the hub. Overall, changing θ, ψ and ϕ has little effect on the estimated real income gains and urban primacy. On the other hand changing σ has much larger effects.

1.6 Conclusion

This paper provides novel evidence on the relationship between trade and regional urban primacy in Europe. It develops a quantitative model of economic geography with novel features to measure the effects of European integration on the spatial distribution of economic activity and more specifically on urban growth.

I develop an open-economy quantitative spatial framework in which large and centrally located cities take on the role of trading hubs and also benefit from increased trading opportunities in the form of endogenous trade driven amenities. I also provide new evidence regarding the trade-amenity nexus at the city level. Despite its ability to accommodate different geographical structures with great flexibility, the model is able to provide straightforward predictions on the effects of trade on urban growth with few parameter requirements. Next, I present reduced form evidence regarding population growth rates of Central and Eastern European cities. I show that cities in close proximity to the new markets, and gain relatively more in terms of market access, experience a sharper increase in their population growth trends. Finally, I quantify the role of trade and market integration on urban growth by structurally estimating the model and find that it can explain a significant fraction of the trends observed in contemporary Europe. Results show that while all of Europe gained from EU enlargement, its effects were more prominent for the developing economies of Eastern Europe. Urban primacy increases at the aggregate level but effects are very heterogeneous across space. Urban growth is strongest in regions close to the border and within smaller economies, since they benefit the most from trade liberalization.

Understanding the role European integration plays in shaping the urban landscape has important implications regarding policy. The diverging fortunes of European cities and the growing urban-rural divide in many countries lead to social and political tensions that may curtail the development of the union. The future success of the European Project calls for policies that can complement the growth observed in successful hubs with more economic participation of the periphery.

Future work might aim at exploring the link between trade and amenities more. Consumer amenities are one of the key drivers of city growth. Therefore a more comprehensive understanding of their relationship with trade will yield valuable insights regarding city growth. In addition, the quantitative framework developed here can be extended to understand specific urban growth patterns we see in Europe today. A considerable fraction of European urban growth is propagated via the suburban expansion of cities into surrounding municipalities. This paper abstracts away from this pattern by focusing on the "core city" in order to unveil the foundational link between trade and urban growth. Building on this theoretical framework, future research can paint a more detailed picture by analyzing growth patterns within regions with different degrees of urbanization.

Chapter 2

BORDERS WITHIN EUROPE

Joint with Jaume Ventura (CREi, Universitat Pompeu Fabra and Barcelona School of Economics) and Marta Santamaria (University of Warwick).

2.1 Introduction

How do country borders affect trade flows within Europe? Using a newly constructed data set of regional trade in Europe, Figure 1 shows sales from Catalonia (shown in grey) to 268 European regions as a share of total spending in each destination region. A striking aspect of these market shares is their national bias. Catalonia's total share of Spanish markets, excluding Catalonia, is 5.8 percent; while its total share of non-Spanish markets is only 0.26 percent. Catalonia is not special in this regard, though. A similar national bias emerges when we examine market shares for other European regions. For the average region (whose size is about 25 percent that of Catalonia) the intranational and international market shares are 2.2 and 0.08 percent respectively.

To what extent is this bias caused by country borders?[1] Comparing in-

[1] We say that there is a border between two regions if they belong to different countries. Thus, we adopt a purely political view of borders, i.e. having a border means not sharing

Figure 2.1: Market shares of Catalonia in Europe

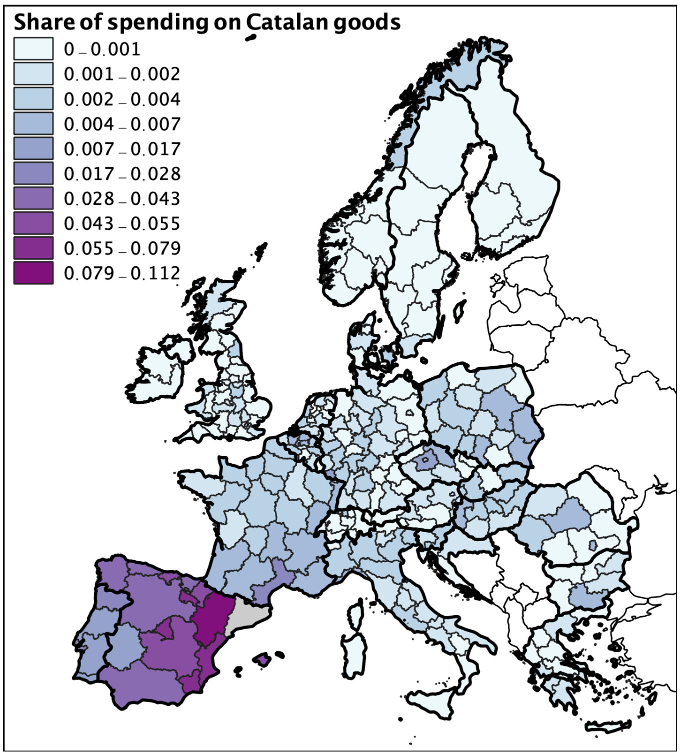

Notes: The figure shows the share of spending on Catalan goods in each European region. The shading represents the value of the market share, with darker shares representing larger market shares. The spending shares come from our newly built regional trade dataset (see Section 2.2).

tranational and international trade could be misleading. As Figure 1 shows, Spanish regions are on average closer to Catalonia than non-Spanish regions. Since geographical distance raises transport costs and reduces trade, this creates an identification problem. A cleaner strategy would be to compare neighbouring regions. For instance, the market share of Catalonia in Languedoc-Rousillon (in France just north of Catalonia) is almost three times smaller than the market share of Catalonia in Valencia (in Spain just south of Catalonia). Is this difference caused by the French-Spanish border or the Pyrenees mountain range that coincides with it? We need to make comparisons that control for factors, such as distance and mountain ranges, that influenced the

a country government.

placement of borders in the past and may influence trade outcomes today.

To search for these confounding factors, normalize market shares by their average and think about them as deviations from the predictions of a naïve gravity model:[2]

$$\ln\left(\frac{n\text{'s share of market } m}{n\text{'s share of all markets}}\right) = \ln\left(n\text{'s sales to } m\right) - \ln\left(\frac{n\text{'s total sales} \times m\text{'s spending}}{\text{spending in all markets}}\right)$$

where n and m are the origin and destination regions, respectively. The LHS is the (log) normalized market share, while the RHS is the difference between the actual (log) sales and the predicted (log) sales using a naïve gravity model. Naïve gravity applies if (i) regions produce differentiated products; (ii) regions have common homothetic preferences, and (iii) trade costs are negligible. Under these assumptions, all regions purchase the same proportions of all goods and, as a result, these proportions must be the average ones:

$$\frac{n\text{'s sales to } m}{m\text{'s spending}} = \frac{n\text{'s total sales}}{\text{spending in all markets}}$$

Since assuming that regions produce differentiated products is uncontroversial, our search for confounding factors must focus on differences in preferences and trade costs.

There is a national bias in preferences if, for a common set of prices across regions, spending falls disproportionally on national goods, i.e. a violation of assumption (ii). One reason for such a bias is the behavior of governments. Eager for political support, governments prefer to award procurement contracts to expensive domestic suppliers instead of cheap foreign ones.[3] Another reason for a national bias in preferences is the behavior of individuals, who often prefer expensive domestic goods than cheap foreign

[2]To see this relationship, simply note that (i) n's share of market m equals n's sales to m divided by m's spending; and (ii) n's share of all markets equals n's total sales divided by spending in all markets.

[3]Herz and Varela-Irimia (2020) examine 1.8 million European public procurement contracts awarded from 2010 to 2014 and published in the EU's Tenders Electronic Daily

ones. Over the last couple of centuries, national governments have made massive efforts aimed at creating a common national identity. Policies such as the adoption of a single official language, the advancement of shared interpretations of history and traditions, the homogenization of educational systems and the promotion of internal migration, have all contributed to the creation of a national culture and, together with it, a preference for national goods. We treat this behavior of governments and individuals as endogenous to the border, as channels through which country borders affect trade.

There is a national cost advantage if trade costs are lower for intranational than for international trade, i.e. a violation of assumption (iii). Although tariffs have been eliminated and technical regulations have been de jure harmonized within Europe, many de facto trade barriers remain. National courts ruling on contract disputes tend to favor national firms, raising the costs of foreign firms to operate in the domestic market. National regulators tend to impede conformity assessments of foreign products to favor domestic firms. National agencies create infrastructure systems that favor intranational mobility, often at the expense of international mobility. These factors are endogenous to the border, additional channels through which country borders affect trade.

There is an important part of the national cost advantage, however, that is due to geography and cannot be attributed to country borders. The cost of transporting goods grows with distance and the presence of geographical obstacles, such as mountain ranges or seas; and it shrinks with the presence of geographical advantages, such as navigable rivers or plains. Individual spending falls disproportionally on goods with low transport costs, and these tend to be lower for intranational trade than for international trade. Interestingly, geography might also contribute to the national bias in preferences. Even if technological improvements were to eliminate transport costs, the

database. The probability that a firm located in the same region as the contracting authority obtains a contract is 900 times larger than that of a firm located abroad, but only 2 times larger than that of a firm located in another region of the same country.

effects of geography would still be felt as past transport costs interact with habit formation to shape present individual preferences. Since geography precedes borders and causes them (as we shall show formally later), we need an empirical strategy that effectively controls for geographical factors and produces an unbiased estimator of the causal effect of country borders on trade.

The first step in our empirical strategy is to find the appropriate dataset to work with. Measuring the border effect essentially amounts to comparing trade within and across national borders. Although there is plenty of data on trade across national borders, there is a surprising scarcity of reliable data on trade within national borders. A first contribution of this paper is to build a dataset of trade in goods for 269 regions from 24 European countries, using the European Road Freight Transport survey collected by Eurostat. This survey annually records around 3 million shipments of goods by road across Europe. For each shipment, we observe its origin and destination regions, the industry of the goods shipped, the weight of the shipment and the distance covered. We aggregate these shipments and impute export prices to build matrices of bilateral trade flows for 12 industries covering the period 2011 to 2017. This dataset provides the first integrated view of regional trade within Europe. Figure 1, for instance, was simply not known or available before.

The second step in our empirical strategy is to use the causal inference framework (see Imbens and Rubin (2015)) to design a credible identification strategy. We first estimate the probability of having a border (or propensity score) as a function of distance, insularity, remoteness and the presence of mountain ranges and river basins. These covariates explain almost half of the border assignment. Figure 2.2 shows the distribution of propensity scores for Catalonia (again shown in grey). Interestingly, we find regions in Spain, Portugal and France that have similar propensity scores, i.e. for which the border assignment was equally likely ex-ante even though ex-post some have a border with Catalonia and some do not.

We want an estimator that is not only unbiased, but also has a small sampling variance. Imbens and Rubin (2015) argue that there are two factors that reduce the sampling variance: (i) the number of observations (region pairs); and (ii) the balance or overlap of propensity scores between treated (region pairs separated by a border) and control (region pairs not separated by a border) groups. We first examine the entire sample and find that it is too unbalanced to produce reliable estimates. This should be apparent by looking at Figure 2.2. For almost all non-Spanish regions the probability of a border with Catalonia is higher than 90 per cent. Thus, we trim the sample, eliminating extreme observations with propensity scores close to zero or one, to achieve a much better overlap of propensity score distributions between treated and control pairs. We then use the trimmed sample to construct a blocking estimator. That is, we build subsamples or blocks of region pairs with similar propensity scores, we estimate the border effect within these blocks and we weight the block estimates to produce an average border effect. Since the probability of having a border is similar between treated and control pairs within each block, the difference in trade between them can be interpreted as the causal effect of the border.

Take two similar region pairs, the first one containing regions in different countries and the second one containing regions in the same country. The main result of this paper is that the market share of the origin region in the destination region for the international pair is only 17.5 percent that of the intranational pair. We refer to this estimate as the average border effect, and we say that country borders cause reductions in market shares of 0.175. This estimate is quite precise and remarkably similar across blocks, i.e. at different levels of the propensity score. Thus, the specific weighting scheme chosen for the blocking estimator has little effect on the final estimate. We do find some variation, though, when we estimate the border effect for each industry separately. In particular, we find that borders cause reductions in market shares that range from 0.123 to 0.389.

How should one interpret and use our estimate of the border effect? Im-

Figure 2.2: Probability of having a border with Catalonia

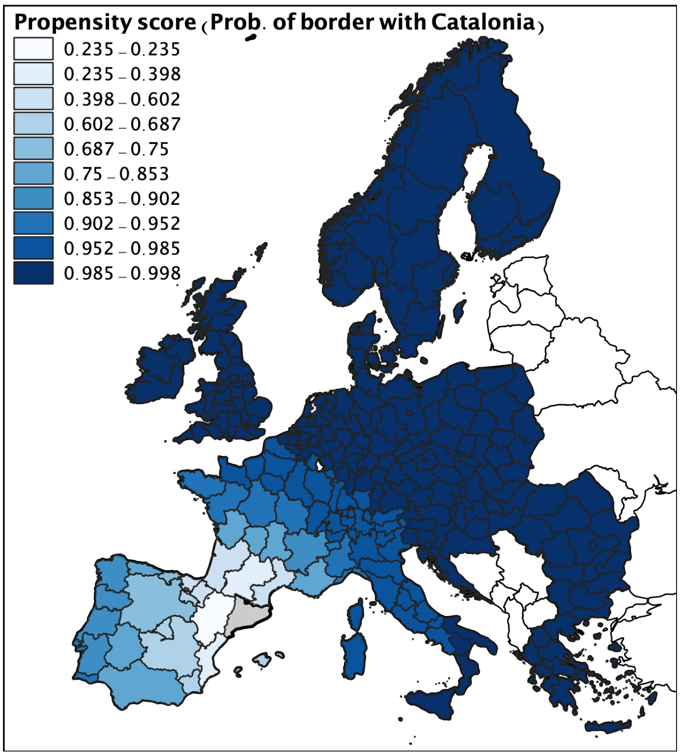

Notes: The figure shows the probability of finding a border between Catalonia and each European region based on a set of geographical covariates (propensity score). The shading represents the value of the market share, with darker shares representing probabilities closer to one.

portantly, it should be treated as a "partial-equilibrium" estimate, i.e. as the effect of changing one border *keeping all other borders constant*. This partial-equilibrium clause, which is standard in micro studies that use the causal inference framework, has an added force in this context. It still contains the standard requirement that region pairs be small so that "treating" one of them does not have general equilibrium effects on European trade. But this is not enough. The units of observation are region pairs, but borders are not bilateral variables. It is not possible in general to "treat" one region pair only, leaving all other pairs "untreated". For instance, consider a counterfactual scenario in which the French-Spanish border were southwest of Catalonia rather than north. This produces 37 border changes affecting

22 French regions and 15 Spanish regions. Since these border changes affect only 0.001 percent of all European region pairs, it seems safe to assume they would have a minor impact on European trade and the partial-equilibrium assumption holds. Thus, we can use our estimate to say that, if history had been such that Catalonia were a French region today, its market shares in other French regions would be $100/17.5 = 5.714$ times larger, while its market shares in Spanish regions would be $17.5/100 = 0.175$ times smaller.[4]

Is our estimate of the border effect large? The answer to this question naturally depends on one's own priors. But we can gain some intuition by being more specific about the counterfactual. After the War of Spanish Succession (1701-1714), the first Bourbon king of Spain Philip V incorporated Catalonia as a province of the kingdom of Spain. What would have happened if, instead, it would have been the French Bourbon king Louis XIV who incorporated Catalonia as a province of the kingdom of France? It is not too far-fetched to think that this would have made Catalonia quite different from what it is today. French would co-exist with Catalan and Spanish would be considered a foreign language, Catalans would exhibit a taste for French goods and traditions rather than Spanish ones, transport systems would foster mobility north rather than south, many Catalans would have their origins and family ties in other French regions rather than in Spanish ones, and so on. Is it surprising to find that, in this scenario, Catalonia would be trading 5.714 times more with other French regions and 0.175 times less with Spanish regions today?

An important observation is that our estimate should be treated as an "average" border effect. One potential source of heterogeneity is the age of the border. It takes a long time to build a common national identity, or an infrastructure system aimed at promoting internal interactions. It takes less time to implement a procurement system that favors domestic firms or

[4]As we explain in Section 3, our estimate is also conditional on the number of borders that regions have. In this counterfactual scenario, the number of borders in Catalonia would drop by 7, and we should adjust our estimate to take this into account. Orders of magnitude do not change, though.

to enact laws and regulations that protect them from foreign competition. Thus, borders with different ages might have different effects. Fortunately (at least for our purposes!), since 1910 Europe has experienced a process of political fragmentation. Indeed, about one third of the region pairs that shared a government in 1910 no longer share a government in 2010. Using the methodology explained above, we find that post-1910 borders reduce market shares to 28.3 percent of their potential. This estimate is still large, but substantially smaller than our estimate of 17.5 percent obtained by pooling pre- and post-1910 borders.

The paper is organized as follows. Section 2.2 describes how we construct the dataset. Section 2.3 explains our identification strategy. Section 2.4 presents our results. Section 2.5 concludes. Before all of this, we review previous efforts to estimate the border effect.

Literature review: In his pioneering study, McCallum (1995) estimated a gravity equation (that is, a linear regression of bilateral trade on economic size and distance) extended to include a border dummy. The estimated coefficient indicated that, after controlling for economic size and distance, trade between Canadian provinces was on average 22 times larger than trade between Canadian provinces and US states. Although the notion that borders hinder trade was not surprising, the magnitude of the effect came as a shock, as model-based explanations based on conventional trade barriers seemed unable to account for the size of the border coefficient.

A first reaction to McCallum's result was mostly methodological, and it centered on how to estimate gravity equations that are consistent with the theory. In an influential paper, Anderson and Van Wincoop (2003) showed that controlling for differences in price levels, something that McCallum (1995) had not done, reduced McCallum's estimate from 22 to 5. The estimation procedure used by Anderson and Van Wincoop (2003) was somewhat burdensome and model-dependent. Feenstra (2002) proposed a much simpler fixed-effects strategy that soon became the standard to estimate gravity equations. This did not affect, though, the finding that controlling for price levels reduces McCallum's estimate from 22 to 5. The methodology to estimate gravity equations evolved rapidly over the next few years.[5] But this has not led to a revision of the effect of the US-Canadian border.

The first contribution of our paper is to shift the focus away from the gravity framework, and towards the causal inference framework. The gravity equation is a relationship between endogenous variables that holds in an interesting class of models that share some assumptions about functional forms. It is useful and reassuring to know that this relationship holds both in the data and in the models. But the coefficient of a border dummy in a gravity equation cannot be interpreted as causal. Borders reduce the spending on goods produced by a region, lowering its income. And yet gravity equations include incomes as independent variables alongside the border dummy. This creates a classic "bad-control" problem when we try to interpret the coefficient of the border dummy as causal.[6] A similar problem applies to bilateral variables that are typically thrown into gravity equations, such as dummies indicating a common language or a common currency. The causal inference framework prescribes specific conditions under which observational data can be used as if it came from an experimental setting, and it forces us to be explicit about the assumptions needed to estimate the causal effects of borders on trade. Moreover, by abandoning gravity (only for this purpose!) our estimates do not rely on specific functional forms or models.

A second reaction to McCallum's result was to go beyond the US-Canadian border and look at the effects of other borders. A major obstacle, though, was the absence of readily available datasets on regional trade for other country

[5] The use of log-linear OLS came under scrutiny due to concerns regarding its performance in the presence of heteroskedasticity (Silva and Tenreyro, 2006) and its inability to incorporate zero trade flows (Helpman et al., 2008). As a consequence, more flexible estimation methods such as Poisson-Pseudo Maximum Likelihood and Gamma-Pseudo Maximum Likelihood became customary. Head and Mayer (2014) provide a review of these developments.

[6] This problem cannot be solved by using origin and destination fixed effects, which are precisely designed to capture economic size and other factors that are endogenous to the border.

pairs. Wei (1996) and Nitsch (2000) computed intranational trade as national production minus exports and compared it to international trade for OECD and European countries, respectively. Later studies measured intranational trade using data at the region-region level and international trade using data at the region-country level (See, for instance, Gil-Pareja et al. (2005) and Coughlin and Novy (2021)). This was indeed an improvement, although comparisons between different units are still far from ideal.[7]

The second contribution of our paper is the construction of a new dataset of bilateral regional trade for 269 regions in 24 European countries that allows region-region level comparisons.[8] As we show next, this dataset constitutes a major leap forward in terms of data quality and coverage. We are not aware of any other dataset with similar characteristics that could be used to reliably measure the causal effect of country borders on trade.

2.2 European regional trade: a new dataset

The European Road Freight Transport survey (ERFT) is a micro-level survey of freight road shipments collected by the statistical office of the European Union, Eurostat. The ERFT data is collected from a survey of shippers in the industry, and is therefore similar in nature to the Community Flow Survey data available for the United States that has been used in a number of empirical studies. This section describes the main features of the ERFT

[7]The problem is aggravated because working with the wrong units also makes it difficult to measure distance. Head and Mayer (2009) showed that accurate measurement of distance is critical to having a precise estimate of the border coefficient. Moreover, Hillberry and Hummels (2008) and Coughlin and Novy (2021) have shown that using large geographical units overlooks the non-linear effect of distance on trade, generating an upward bias on the border coefficient.

[8]Gallego and Llano (2015) is the only study we have found that uses region-region level data to measure both types of trade and focuses on a border other than the US-Canadian one. This study uses a road transport survey to construct a dataset of flows from each Spanish region to itself, other Spanish regions and to the regions of Spain's 7 main trade partners in the EU. The paper however follows the gravity methodology and does not attempt to estimate the causal effect of the border.

survey and shows how we use it to build our dataset.

A natural question is whether freight road shipments are representative of all trade flows. According to Eurostat's own statistics, between 2011 and 2017 road freight accounted for about 49 percent of all intra-EU trade in tonne-km terms, while the share of maritime short-sea shipping and rail transport were 32 percent and 11 percent respectively (the other modes of transportation reported are inland waterways 4, pipelines 3 and air 0.1). Thus, we think that our dataset measures a sizeable fraction of intra-European trade.

2.2.1 From road shipments to regional trade weights

The ERFT survey covers shipments by road aggregated every year from micro-data collected by a total of 29 European countries, all European Union members except for Malta plus Norway and Switzerland.[9] Each participating country chooses a stratified sample of vehicles from the national register of road freight vehicles, following Eurostat guidelines.[10] The operators of the sampled vehicle are required to report, for a limited number of days in a month, the characteristics of all the shipments completed.

The survey requests information at the level of the vehicle, the journey and the specific goods shipped. At the level of the vehicle, the survey records vehicle characteristics such as age, type of vehicle and ownership. At the journey level, the questionnaire records whether the journey is loaded

[9]The European Union adopted in 1998 regulation to provide a legal base for the collection of a wide range of data on road freight transport ((EC) 1172/98), laying the emphasis on quality and comparability of statistical information. This regulation has introduced major changes in the data collected in order to describe the regional origin and destination of intra-European Union transport on the same basis as national transportation (Road Freight Transport methodology, 2016 edition).

[10]The selection of the sample is made to ensure that the raw survey results are representative of the total numbers recorded on the vehicle register. In countries where such a registry is not available or sufficiently reliable, a register of persons licensed to operate as road hauliers (company/registered owner for private hauliers) or a business register of companies could be considered. In this case, the sampling unit could be the vehicle operators or transport companies. (Road Freight Transport methodology, 2016 edition) Further details are provided in the ERFT survey documentation.

or unloaded, the type of transport (hired or own account) and the type of journey.[11] At the goods level, the record includes the shipment's weight (kg), the type of goods carried according to the 2 digit NST 2007 classification, the region of origin and destination (at NUTS3 level), the actual shipping distance covered and a sampling weight for each shipment.[12] Eurostat aggregates the origin and destination of each shipment into larger regions (at NUTS2 level) for anonymity reasons. The ERFT survey is available for the period 2011 to 2017. Using this micro-dataset has several advantages relative to using aggregate trade data. It also requires us to make some adjustments.

A first advantage of the survey is that it allows us to overcome one of the main challenges to estimate the border effect: the lack of subnational trade data. The ERFT survey allows us to distinguish between flows within a region and flows between regions in the same country for all countries surveyed except for five one-region countries: Cyprus, Estonia, Latvia, Lithuania and Luxembourg. For this reason, we drop these countries from the dataset. This leaves us with 24 countries in our sample: the remaining 22 European Union countries plus Norway and Switzerland.

A second advantage of the survey is that it is collected from a stratified sample of actual shippers rather than imputed from different aggregated data sources. This means that our data captures, with higher accuracy, the movement of goods within countries. The survey includes two types of flows: shipments that move goods between producers and consumers and shipments that move goods from a producer to an intermediary or from intermediary to intermediary. What the survey actually captures is the region to region distribution of goods. In most cases, these shipments will take goods from the origin to the destination region. Yet, in other cases, these shipments will

[11] The type of journey records whether the journey involved one single transport operation, several transport operations or a collection/distribution of goods, with many stopping points for loading and/or unloading in the course of a single journey.

[12] The weight of shipments is calculated by multiplying reported estimates by the inverse of the sampling weight. The industry classification followed in the survey is the NST 2007 classification, the "statistical classification of economic activities in the European Community".

be a middle step in a longer distribution chain across European regions, not coinciding with the observed origin and destination of the trade flow.

To address this limitation, we restrict our sample in three ways. First, we use the detailed information in the survey to drop journeys that are classified as distribution journeys. These journeys are characterised by the existence of several stops between the origin and the destination to load and/or unload goods. Dropping these journeys seeks to bring our shipment data closer to trade data.

Second, we restrict the number of industries in the analysis. The shipments are classified into 20 industries enumerated in Table A.2.1 in the Appendix. We adopt two criteria for industry coverage: (i) the industry must be unambiguously associated with trade; and (ii) transport by road must be an important mode of transport for the industry. The first criterion leads us to discard eight industries.[13] The second criterion leads us to discard one additional industry.[14] Thus, we are left with twelve industries.

Finally, we want to make sure that the survey on road shipments is representative of aggregate trade. This would not be the case for regions with a very small share of shipments traveling by road. To ensure this, we restrict the number of regions by dropping insular regions very far from continental Europe. For these small and far away regions, shipments by road are not likely to be representative.[15] Table A.2.2 in the appendix provides a list of all regions.

[13]These industries are: 14 Secondary materials, municipal wastes and other wastes; 15 Mail, parcels; 16 Equipment and materials utilized in the transport of goods; 17 Goods moved in the course of household and office removals, 18 Grouped goods; 19 Unidentifiable goods; and 20 Other goods n.e.c. It is unclear to us what fraction of the shipments included in these categories can be safely classified as trade in goods. For instance, disposing of waste, distributing mail or moving furniture is clearly not associated with trade.

[14]This industry is: 2 Coal and lignite, crude petroleum and natural gas. A large fraction of trade in this industry is transported by railways or through pipelines.

[15]We keep large, close-by islands like Sardinia or Sicily. The survey includes shipments taken by truck when the truck is loaded on a ship and unloaded after crossing to an island. Therefore, we can include these larger islands since their trade is well represented in the survey.

After all these adjustments, our dataset contains 269 regions (in 24 countries) and 12 industries. We use the dataset to construct a set of industry-year matrices:

$$W^{it} = \left[W^{it}_{nm}\right]_{269\times 269}$$

where W^{it}_{nm} is the weight (kg) from industry i shipped from region n to region m in year t. Since our dataset contains 12 industries and 7 years, we have 84 such matrices.

Figure 2.3 plots exports (kg) across the countries in our sample in the Y-axis against bilateral shipments (kg) obtained by aggregating the survey data at the country level on the X-axis. As we can see, most observations concentrate along the 45 degree line (Rsq=0.55), showing that our data is very correlated with aggregate exports data from Eurostat. Figures A.2.1, A.2.2 and A.2.3 in the appendix plot the same relationship, year-by-year and industry-by-industry. These figures show that this correlation is also strong when we use data disaggregated by industry and/or year.

2.2.2 From trade weights to trade values

The survey provides trade weights, and we would like to convert weights into values. Thus, we look for other data sources. The statistical agencies of France, Germany, Spain and United Kingdom release data of exports from individual regions to foreign countries in value and volume. These data allows us to observe export flows from 66 regions in our sample (belonging to the four countries mentioned above) to all the remaining countries in our sample. For these export flows, we observe the value in euros and the quantity in kilograms of export flows, allowing us to compute the price per kilo of exports. Unfortunately, similar data could not be collected for the remaining countries in our sample. The reason why such regional level data on exports is not available for other countries is unknown to us and, hopefully, not systematically related to the price of exports in those regions. Therefore, we think of our data as incomplete data in which the price of exports is

Figure 2.3: Correlation with aggregate international trade data

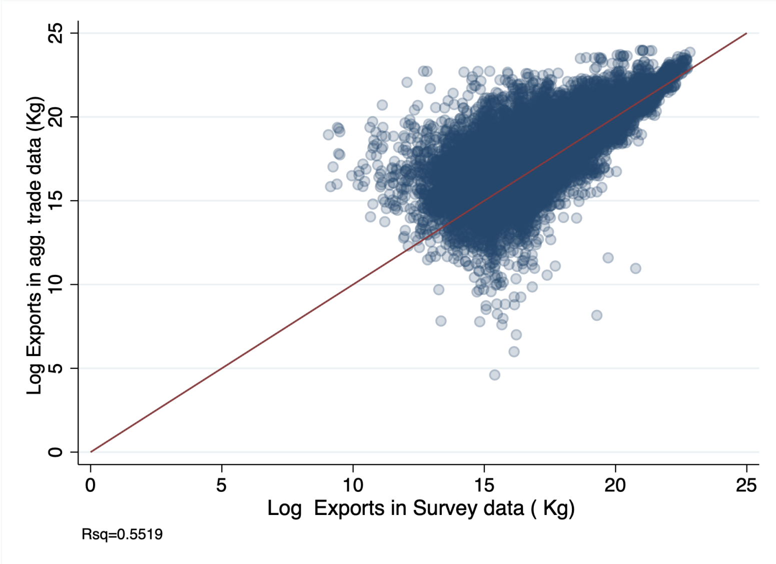

Notes: The figure shows the correlation between exports and shipments in the ERFT survey in kilograms. The Y-axis represents (log) bilateral trade (kg) by country-pair-industry-year using international trade data from Eurostat. The X-axis shows bilateral shipments by road (kg) aggregated by country-pair-industry-year obtained from the ERFT survey.

missing for part of the sample.

Imputation methods replace missing values by suitable estimates and then apply standard methods to the filled-in data. Imputations are means or draws from a predictive distribution of the missing values, and require a method for creating a predictive distribution for the imputation that is based on the observed data. We choose an explicit modelling approach, where the distribution is based on a formal statistical model. In particular, we use regression imputation, a standard choice of conditional mean imputation. First, the regression of the variable with missing values on other covariates is estimated from the complete cases, and then, the resulting prediction equation is used to impute the conditional mean of the missing values. Regression imputation is a plausible method, particularly when the chosen covariates explain most of the variation of the variable with missing values.

Our preferred specification is to pool all time periods and industries to estimate a linear regression for the (log) of the price of exports, calculated as the ratio between the value of exports and the weight of exports for each industry, origin, destination and year. As explanatory variables, we use a vector of origin and destination characteristics. The only bilateral variable that we use is distance.[16] We also include industry-time dummies to allow for different time trends in prices across industries. Table A.2.9 in the appendix contains the full list of variables included in the price regressions.

Our regression model seems to perform well, as shown in Table A.2.3 in the Appendix. The R-squared in the above specifications is higher than 50 percent. Since the collected variables explain a large share of the variation in export prices in the subsample with no missing values, we can use the estimated coefficients from the linear regression to impute the values that are missing.[17]

With our estimated prices per unit, we can finally construct the trade value data for each industry i and year t as follows:

$$V^{it} = \left[V^{it}_{nm}\right]_{269\times 269} \qquad \text{where } V^{it}_{nm} = P^{it}_{nm} \cdot W^{it}_{nm}$$

where V^{it}_{nm} is the value (euros) from industry i shipped from region n to region m in year t.

[16] As shown in Hummels and Skiba (2004), the presence of transport costs leads firms to ship high-quality goods abroad while keeping low-quality goods for the domestic market. This is known as the "Alchian and Allen conjecture" (see Alchian and Allen (1964)). Another reason why export prices per kilogram could increase with distance is transport costs. However, our export prices are Free On Board (F.O.B), meaning that they are net from transport and insurance costs.

[17] In order to further assess the accuracy of our imputed prices we perform two sets of checks. First, we perform a series of out-of-sample estimations where we drop one of the four countries for which we observe regional export prices and we predict export prices for this dropped country. We then compare our out-of-sample estimates with the actual regional prices (See Figure A.2.4 in the Appendix). Second, we collect export value and weights from Eurostat for all European countries and compute unit export prices for every country-pair at the industry and year level. We aggregate our region-pair estimated prices to a country-pair level and compare them to the country-pair price of exports from international trade data (See Figure A.2.5 in the Appendix). Both tests suggest that our imputed prices are reasonable.

Figure 2.4: Correlation with aggregate international trade data

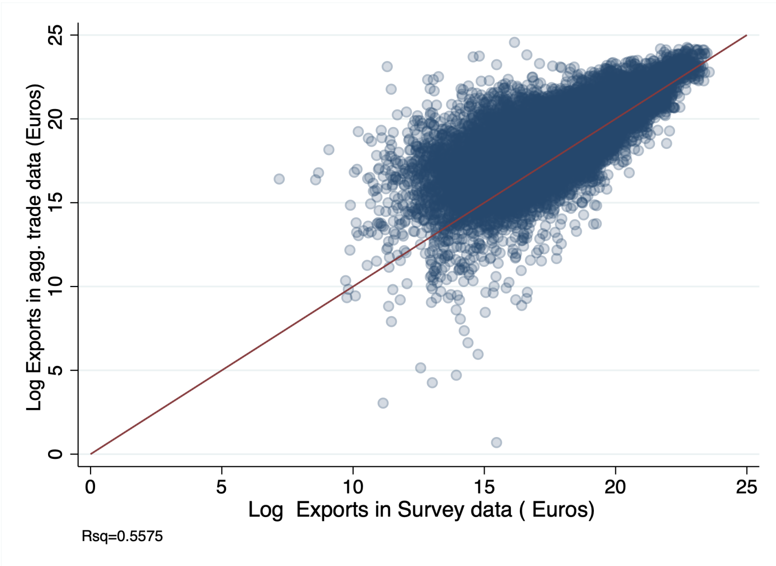

Notes: The figure shows the correlation between exports and shipments in the ERFT survey in euros. The Y-axis represents (log) bilateral trade (euros) by country-pair-industry-year using international trade data from Eurostat. The X-axis shows bilateral shipments by road (euros) aggregated by country-pair-industry-year obtained from the ERFT survey after imputing missing prices.

Figure 2.4 plots exports (euros) across the countries in our sample in the Y-axis against bilateral shipments (euros) obtained by aggregating the survey data at the country level on the X-axis. As we can see, most observations concentrate along the 45 degree line (R-squared = 0.55), showing that our data is very correlated with aggregate exports data that come from Eurostat when we use values. Figures A.2.6, A.2.7 and A.2.8 plot the same relationship, industry-by-industry and year-by-year. These figures show that this correlation is also strong when we use data disaggregated by industry and/or year.

2.2.3 European Regional trade: A first look at the data

Our dataset contains region pairs such that: (i) origin and destination regions belong to the same country; and (ii) origin and destination regions belong

Table 2.1: Summary statistics

Trade type	Intranational trade Mean	International trade Mean
Panel A: Unconditional		
Value (Mill. euros)	553.52	18.61
Weight (Mill. Kg)	601.49	9.98
Normalized Market share	10.87	0.27
Panel B: Zero trade observations		
Region pairs	4958	67134
Region pairs with no trade	157	25699
Regions pairs with positive trade	4801	41435
Panel C: Conditional on positive trade		
Value (Mill. euros)	571.62	30.15
Weight (Mill. Kg)	621.15	16.17
Normalized Market share	11.22	0.44

Notes: This table reports the (unweighted) average bilateral trade flow (euros and kilos) and the (unweighted) average normalised market share in our new European regional dataset. Column 1 reports the average flow between intranational region pairs (origin and destination in regions in the same country) and column 2 reports the average flow between international region pairs (origin and destination regions in different countries). Panel A reports unconditional statistics. Panel B repors the number of region pairs that display positive trade and zero trade. Panel C reports statistics conditional on trading.

to different countries. We refer to these two types of trade as intranational and international, respectively.[18] Out of a total of 72,092 region pairs in our sample, 4,958 are intranational, and 67,134 are international.[19]

Panel A of Table 2.1 shows the average values of the two types of trade at the region-pair and annual level. We see that the average value of trade among intranational pairs is almost 30 times larger than among international pairs. This average is unweighted, and one might think that it could be affected by differences in economic size between groups. We obtain a similar picture, however, when we look at normalized market shares.

Panel B of Table 2.1 shows another important feature of our data, the prevalence of region pairs that do not trade. Among intranational pairs, 96.8 percent exhibit positive trade. The picture is quite different when we look

[18] We exclude from our sample pairs for which the origin region is the same as the destination region. Therefore, intranational trade does not include trade within a region.

[19] These numbers take into account origin and destination. Thus, we count region pair (n, m) as different than (m, n).

at international pairs. Among them, only 61.7 percent of pairs trade with each other. Taking this into account, Panel C of Table 2.1 shows the same statistics as in Panel A but now conditional on observing a positive flow of goods. Not surprisingly, this increases the average trade values among international pairs, without affecting much the average trade values of the other group. The main takeaway is that the national bias manifests itself both on the intensive and the extensive margins.

2.3 Identifying the border effect

The causal relationship of interest is the effect of country borders on trade. In this section, we describe our empirical strategy to identify this effect which draws heavily from the causal inference framework (see Imbens and Rubin (2015)). We use as an outcome variable, the normalized market share:

$$S_{nm} \equiv \frac{V_{nm}/E_m}{Y_n/E} \tag{2.1}$$

where $Y_n = \sum_m V_{nm}$ are the total sales or income of region n; $E_m = \sum_n V_{nm}$ are the total purchases or spending of region m, and $E = \sum_m E_m$ is total spending by all regions. The variable S_{nm} measures region n's share of region m's market normalized by region n's share of all markets, including its own. If market m has an average importance to producers of region n, i.e. $V_{nm}/E_m \approx Y_n/E$; the market share is one. If instead market m has a larger (smaller) than average importance, the market share is above (below) one. Unlike trade values, normalized market shares are not affected mechanically by the economic size of origin and destination regions.[20] This makes them more

[20]To see this, assume trade is balanced, i.e. $E_m = Y_m$ and E = Y. Then, we have that:

$$lnS_{nm} = lnV_{nm} - lnY_n - lnY_m + lnY$$

Since $Y_n = \sum_m V_{nm}$, one might think that ln S_{nm} is obtained by taking out fixed effects from ln V_{nm}. This is close, but not quite right. To construct ln S_{nm}, we subtract and add the logs of the means to ln V_{nm}, and not the means of the logs.

helpful than trade values to infer preference biases and trade costs.

2.3.1 The border effect

The French-Spanish border runs across Catalonia and Languedoc-Roussillon, and not across Catalonia and Valencia. Catalonia's average market share in all the 269 regions in our sample is 1.5 percent. Given how close Catalonia is geographically and culturally to Languedoc-Roussillon and Valencia, it is not surprising that these two markets be specially important for Catalan exporters. Indeed, the normalized share of Catalonia in the Languedoc-Roussillon market is well above one, 1.79, implying that $1.79 \times 1.5 = 2.7$ percent of all the spending of Languedoc-Roussillon is on products that come from Catalonia. Yet Catalonia's normalized share of the Valencia market is almost three times larger than this, 5.21, implying that $5.21 \times 1.5 = 7.9$ percent of all the spending of Valencia is on products that come from Catalonia. To what extent is this difference caused by the French-Spanish border? What would have happened if this border were southwest of Catalonia instead of north? How much would Catalonia's share of the Languedoc-Roussillon market grow? How much would Catalonia's share of the Valencia market shrink?

Answering these questions involves comparing observed market shares with the counterfactual market shares that would have occurred if the French-Spanish border were southwest of Catalonia. More formally, let (n, m) be a region pair, and let $B_{nm} \in \{0, 1\}$ be a dummy variable that takes value one if the regions in the pair belong to different countries, and zero otherwise. Let S_{nm} be the observed market share for region pair (n, m) in our sample. We define two potential market shares as follows:

$$S_{nm} = \begin{cases} S_{nm}(1) & \text{if } B_{nm} = 1 \\ S_{nm}(0) & \text{if } B_{nm} = 0 \end{cases} \tag{2.2}$$

where $S_{nm}(1)$ and $S_{nm}(0)$ are region n's share of market m with a border (active treatment) and without a border (control treatment), respectively.

For each region pair, we observe only one potential outcome. For instance, we observe $S_{CAT,L-R}(1) = 1.79$ for the pair (Catalonia, Languedoc-Roussillon) and $S_{CAT,VAL}(0) = 5.21$ for the pair (Catalonia, Valencia). Unfortunately, we do not observe $S_{CAT,L-R}(0)$ or $S_{CAT,VAL}(1)$.

We define the border effect β_{nm} as the log change in market shares caused by the border:

$$\beta_{nm} = \ln \frac{S_{nm}(1)}{S_{nm}(0)} \tag{2.3}$$

Since one potential outcome is unobserved, we cannot observe border effects. It is tempting however to assume that, if the French-Spanish border were southwest of Catalonia, the roles of these two markets for Catalan exporters would reverse, that is, $S_{CAT,L-R}(0) = S_{CAT,VAL}(0)$ and $S_{CAT,VAL}(1) = S_{CAT,L-R}(1)$. This identification assumption allows us to estimate a common border effect for the two region pairs as follows:

$$\beta = \ln \frac{S_{CAT,L-R}(1)}{S_{CAT,VAL}(0)} = -1.07 \tag{2.4}$$

That is, the French-Spanish border reduces Catalonia's share of the Languedoc-Roussillon market to a third of its potential: $100 \times e^{-1.07} = 34.3$ percent. Should we take this estimate very seriously? How good is the identification assumption that underlies it? The main challenge we face in this paper is to construct samples for which this type of comparisons can be interpreted as causal.

There are a couple of assumptions embedded in our notation worth mentioning explicitly. The first one is that the unobserved potential outcome is unique. As mentioned, moving Catalonia to France would remove the border between Catalonia and Languedoc-Roussillon. But so would moving Languedoc-Roussillon to Spain, or creating a new country containing both regions. Our framework implies that $S_{CAT,L-R}(0)$ is the same in all these cases and, indeed, in any other possible case. This assumption captures the view that, to a first-order approximation, what matters is whether there is a border or not. The specific type of border only matters to a second or third-order approximation. We think this is quite a reasonable view.

Our notation also embeds the notion that the difference in potential outcomes measures the effect of changing the border for one region pair, *keeping all other borders constant*. This partial-equilibrium clause, which is standard in micro studies that use the causal framework, has an added force in this context. It still contains the standard requirement that region pairs be small so that "treating" one of them does not have general equilibrium effects on European trade. But this is not enough in this context. The units of observation are region pairs, but borders are not bilateral variables. It is not possible in general to "treat" one region pair only, leaving all other pairs "untreated". Consider again moving the French-Spanish border southwest of Catalonia. This experiment would remove the border between Catalonia and 22 French regions and create a border between Catalonia and 15 Spanish regions. Thus, it would produce 37 border changes. Since these border changes affect only 0.001% of all region pairs, it seems safe to assume they would have a minor impact on European trade and the partial-equilibrium assumption holds.

Since we cannot experiment with borders, we must rely on observational data to estimate an average border effect. In particular, we define the average border effect β as the average log change in market shares caused by the border as:

$$\beta = E\left(\ln S_{nm}\left(1\right) - \ln S_{nm}\left(0\right) \middle| \, S_{nm}\left(1\right) > 0, \; B_{nm} = 1\right) \tag{2.5}$$

The value of β is expected to be negative since the border is expected to reduce trade. The larger is $|\beta|$, the larger is the average reduction in market shares caused by the border. Throughout, we assume that there are no region pairs such that $S_{nm}\left(1\right) > 0$ and $S_{nm}\left(0\right) = 0$. Obviously, this cannot be verified.

The causal inference framework shows that we can use observational data as if it came from an experiment if the assigment of treatment is (i) probabilistic, (ii) individualistic and (iii) uncounfounded. If the assignment mechanism satisfies these conditions, the comparison of units with different treatments but identical pre-treatment covariates can be given a causal interpretation.

We believe that the first two conditions hold in our setting. Probabilistic assignment requires a nonzero probability for each treatment value, for every unit. The probability that two far-away regions belong to the same country might be very small, but it is not zero. Individualistic assignment requires limited dependence of a particular unit's assignment probability on the values of covariates and potential outcomes for other units. This is the partial-equilibrium clause mentioned above, which we argued is a reasonable one.

The last condition, unconfounded assignment, deserves much more attention. Under unconfoundedness, all the assignment probabilities are free from dependence on potential outcomes, after conditioning on a vector of pre-treatment covariates. This assumption is often referred to as the Conditional Independence Assumption (see Dawid (1979)) and written as $B_{nm} \perp S_{nm}(0), S_{nm}(1) | X_{nm}$. In our setting, unconfoundedness means that the assignment of borders must be independent of potential trade outcomes across regions, after conditioning on a vector of pre-treatment geographical covariates X_{nm}. We describe this vector and explain our control strategy in the next couple of sections.

Let us assume for now that we have a vector of pre-treatement geographical covariates X_{nm} such that, after conditioning for them, the border assignment is unconfounded. This allows us to interpret comparisons between units with different treatments as causal. Does this mean that we can estimate the border effect by simply comparing the average market shares of international and intranational pairs with the same covariate values $X_{nm} = x$? The answer, unfortunately, is negative. The following estimator makes exactly this comparison:

$$\begin{aligned}\hat{\beta} \quad = \quad & E\left(\ln S_{nm}(1) \middle| \mathcal{S}_{nm}(1) > 0, B_{nm} = 1, X_{nm} = x\right) \\ & -E\left(\ln S_{nm}(0) \middle| \mathcal{S}_{nm}(0) > 0, B_{nm} = 0, X_{nm} = x\right)\end{aligned} \tag{2.6}$$

It is straightforward to see that $\hat{\beta}$ suffers from two potential sources of selection bias:

$$\begin{aligned}\hat{\beta} - \beta = & \underbrace{\begin{array}{c} E\left(\ln S_{nm}(0) \middle| \mathcal{S}_{nm}(1) > 0, B_{nm} = 1, X_{nm} = x\right) \\ -E\left(\ln S_{nm}(0) \middle| \mathcal{S}_{nm}(1) > 0, B_{nm} = 0, X_{nm} = x\right) \end{array}}_{\text{Selection bias due to the number of borders}} \\ + & \underbrace{\begin{array}{c} E\left(\ln S_{nm}(0) \middle| \mathcal{S}_{nm}(1) > 0, B_{nm} = 0, X_{nm} = x\right) \\ -E\left(\ln S_{nm}(0) \middle| \mathcal{S}_{nm}(0) > 0, B_{nm} = 0, X_{nm} = x\right) \end{array}}_{\text{Selection bias due to changes in participation}}\end{aligned} \tag{2.7}$$

Consider first the selection bias due to the number of borders, which is the first term of Equation (2.7). It might seem surprising that we condition on the border after assuming that the border assignment is unconfounded. But there is a subtle source of selection bias that arises from any random border assignment, including those that are unconfounded. To understand its nature, consider a world with 6 regions and 2 countries. The six regions are identical in any possible way, except for the border assignment. The latter is random, with all regions being equally likely to belong to any country. Let us assume that the realization of the border assignment is such that regions 1 and 2 belong to country A, while regions 3, 4, 5 and 6 belong to country B. This introduces the only source of asymmetry in this world: regions in A have four borders, while regions in B have only two borders. Assume there are no trade costs other than those caused by the border, which result in the same percentage reduction in market shares for all pairs:

$$\beta = \ln \frac{S_{nm}(1)}{S_{nm}(0)} \quad \text{for all } n, m \tag{2.8}$$

Let S_A^D and S_B^D be the market share of any region in A and B in a domestic market (including itself), respectively. Symmetry and the absence of non-border related trade costs ensure that, within each country, these shares are identical for all relevant pairs. Let S_A^F and S_B^F to be the market share of any region in A and B in a foreign market, respectively. Symmetry and the absence of non-border related trade costs also ensure that, within each country, these shares are identical for all relevant pairs. By construction, normalized market shares must add to one. Thus, we have that

$$2S_A^D(0) + 4S_A^F(1) = 4S_B^D(0) + 2S_B^F(1) = 1 \tag{2.9}$$

It is straightforward to show that Equations (2.8) and (2.9) imply that:

$$\frac{S_A^D(0)}{S_B^D(0)} = \frac{S_A^F(1)}{S_B^F(1)} = \frac{2+e^{\beta}}{2e^{\beta}+1} > 1 \tag{2.10}$$

for any value of $\beta < 1$. That is, regions with many borders have larger market shares. The key observation is that region pairs with many borders tend to be over-represented among international pairs and under-represented among intranational pairs. This creates a positive selection bias that makes the observed difference in average market shares smaller (in absolute value) than the true average border effect.[21]

Fortunately, there is a simple solution to this problem, namely, to estimate border effects conditioning on the number of borders. We shall show later that this type of selection bias is important empirically. But one can already suspect this by looking at Figure 2.5, which shows average intranational and international market shares in panels A and B, respectively. The color of a region represents the value of the average normalized share, with dark blue

[21]The existence of this type of selection bias was noted first by Anderson and Van Wincoop (2003). In their sample, however, the group of intranational pairs contained only Canadian provinces, i.e. regions with many borders; while the group of international region pairs contained mostly US states, i.e. regions with few borders. Thus, they found that this type of selection bias leads to overstating the average border effect. Here, with a balanced sample, this selection bias leads to understating the border effect.

shades representing the smallest values and dark red shades representing the highest values. In countries with many regions, such as United Kingdom or Germany, regions have smaller than average intranational and international market shares (predominantly blue shades). In countries with few regions, such as Belgium, Slovenia, or Portugal, regions have larger than average intranational and international market shares (predominantly red shades).

Consider next the selection bias due to changes in participation, which is the second term of Equation (2.7). This type of selection bias arises because some region pairs trade without a border, $S_{nm}(0) > 0$, but would not trade with a border, $S_{nm}(1) = 0$. Let us refer to these pairs as switchers. Average market shares for intranational pairs include switchers, while average market shares for international pairs do not. If average market shares for switchers and non-switchers were the same, there would be no selection bias and the second term in Equation (2.7) would be zero. But it is reasonable to expect average market shares for switchers to be lower than those of pairs that always trade. This creates a positive selection bias that makes the observed difference in average market shares smaller (in absolute value) than the true average border effect.

The importance of this bias depends on the fraction of switchers in the sample. Without this information, we must treat $\hat{\beta}$ as a lower bound for the border effect. We show later, however, that the fraction of switchers must be quite small in the samples we work with. This means that the bias due to changes in participation cannot be important quantitatively and, as a result, $\hat{\beta}$ provides a good estimate for the border effect.

To sum up, if the border assignment is probabilistic, individualistic and unconfounded, we can compare intranational and international pairs and be confident to obtain a good estimate of the border effect if (i) we condition on the number of borders; and (ii) we check that the fraction of switchers is small.

2.3.2 Understanding the border assignment

Geography affects trade costs and market shares. Since geography precedes borders, this poses an identification problem if the border assignment is also affected by geography. But it is easy to see that this is indeed the case. Our comparison of the (Catalonia, Languedoc-Rousillon) and (Catalonia, Valencia) region pairs shows how difficult it is to escape from this conclusion. Both pairs are contiguous, continental and located on the Mediterranean coast. Thus, comparing their market shares already 'controls' for some of the most relevant geographical factors. But even then, we cannot conclude that the location of the French-Spanish border is unrelated to geographical factors that also affect trade. On its north, Catalonia is separated from Languedoc-Roussillon by the Pyrenees mountain range. On its south, Catalonia shares the Ebro river basin with Valencia. This geographical difference, which affects trade costs, might have also contributed to the French-Spanish border being north of Catalonia rather than southwest.

To satisfy the unconfoundedness condition, causal inference must be conditional on those factors that precede and influence both the treatment assignment and the outcome variable. In our framework, these are the geographical covariates that affect the border assignment and trade outcomes simultaneously. With this idea in mind, we collect the following set of covariates for each region pair:

1. *Distance.* Length of the curve linking the central point of the origin region (centroid) and the central point of the destination region, in kilometers. We use a curve since we take into account the curvature of earth's surface.

2. *Insularity.* Dummy variable taking value one if there is the need to cross a sea to reach from one region to the other, and zero otherwise.

3. *Mountain ranges.* Largest altitude difference between two regions, computed as the difference between the highest altitude point and the low-

est altitude point along the straight line that joins the centre the origin region (centroid) and the centre of the destination region.

4. *River basin*. Dummy variable taking value 1 if both regions belong to the same river basin. We consider the largest rivers in Europe. A map of the areas covered by each river basin can be found in figure A.2.19 in the Appendix.

5. *Remoteness*. We calculate the remoteness of a region as the sum of the bilateral distance from that region to every other region in the sample. Then, we calculate the remoteness of a pair as the average remoteness of both regions.

All these covariates are known to affect bilateral trade, and they can be treated as pre-treatment covariates when considering the border assignment. The next question is whether these covariates also affect the border assignment. Unlike the theory of bilateral trade, which is quite sophisticated and developed at this time, the theory of borders is rough and underdeveloped. Thus, we are forced to rely on some basic conjectures about how these geographical factors affect the costs and benefits of sharing a government.[22]

It seems reasonable to think that distance, insularity and the presence of mountain ranges all raise the costs and lower the benefits of sharing a government. Thus, we would expect these variables to raise the probability of a border assignment. It is less clear however to predict the effects of sharing a river basin. Rivers could be a geographical obstacle such as mountain ranges, but they could also provide a geographical mobility advantage or create externalities that raise the benefits of a shared government. Thus, we do not know a priori whether being in the same river basin raises or lowers the probability of a border assignment. Unconditionally, we would expect remote region pairs to have more borders because they are farther away from each

[22]The relevant costs and benefits are those borne by whomever makes the decision. The decision-maker(s) might be regions in the pair, or other regions elsewhere. Admittedly, the discussion here is quite superficial.

Table 2.2: Covariate distributions across treatment groups

	Treatment group mean	Control group mean	Difference (t-stat)
Distance	1213.62	315.64	-898.0
			(-71.79)
Insularity	0.32	0.06	-0.258
			(-27.23)
Mountain Ranges	1473.66	496.08	-977.6
			(-37.95)
River Basin	0.04	0.19	0.153
			(35.81)
Remoteness	1157.47	1075.85	-81.62
			(-17.19)
N	33567	2479	36046

Notes: This table reports the average value of each geographical covariate in the treatment group (column 1) and in the control group (column 2). The last column reports the difference in means (defined as control minus treated). The t-statistics in parentheses.*Distance* is bilateral distance between origin and destination in kilometers, *Insularity* takes value 1 if one of the regions is an island. *Mountain Ranges* is the highest difference in elevation between two regions in metres (difference between highest point and lowest point), *River Basin* takes value 1 if the region pair shares a river basin and *Remoteness* is the average remoteness of the origin and the destination regions.

other. Conditioning on distance, however, we would expect the probability of a border assignment for a region pair to increase with their remoteness because they have fewer alternative partners to share a government.

Table 2.2 provides summary statistics of these geographical covariates in the treatment and control groups. Intranational pairs are closer to each other, less likely to be insular or separated by a mountain range, more likely to share a river basin, and on average less remote. These differences are significant, and have the expected sign.[23]

To obtain a more convincing assessment of the role of geographical covariates on the border assignment, we estimate the propensity score.[24] In particular, we estimate a logistic regression model, where the log odds ratio

[23]The positive sign on the river basin variable is not informative. International pairs are more distant than intranational ones, making it unlikely that the former be located in the same river basin. One needs to control for distance to determine how sharing a river basin affects the border assignment.

[24]The propensity score at covariate values x is the average probability of border assignment for region pairs (n, m) with covariates $X_{nm} = x$.

of receiving the treatment is modeled as linear in a number of the geographical covariates, with unknown coefficients. We estimate the coefficients by maximum likelihood. To choose how many of our geographical covariates to include in the logistic regression, we follow the recursive procedure recommended in Imbens and Rubin (2015). We find that all the covariates described above should be included.

Table 2.3, column (1) presents the estimation results from the logit model. The coefficients of the covariates are all significative at the 1 percent level and the model has an R-squared of 0.476. As expected, distance, insularity and mountain ranges raise the probability of a border assignment, while remoteness lowers it. Interestingly, we find that being in the same river basin raises the probability of having a border. It seems thus that rivers promote borders rather than the opposite.

By its own nature, the unconfoundedness assumption cannot be proved formally. But economic theory identifies as potential confounding factors a set of geographical covariates that precede the border assignment and affect trade costs. We have shown that, indeed, these covariates affect the border assignment. Thus, comparisons of units with different treatments can be given a causal interpretation only if we condition for these pre-treatment covariates. The next step is to find the right way to do this necessary conditioning.

2.3.3 Constructing the 'right' samples

To measure the border effect we estimate a linear regression model of normalized market shares on the border dummy, controlling for the number of borders and the set of geographical covariates:

$$\ln S_{nm} = \alpha + \beta \cdot B_{nm} + \gamma \cdot N_{nm} + \lambda' \cdot X_{nm} + u_{nm} \tag{2.11}$$

Table 2.3: Propensity Models

Dependent Variable: Border	Full sample (1)	Trimmed sample (2)
Distance	2.998	1.893
	(0.056)	(0.078)
Insularity	1.096	1.059
	(0.096)	(0.128)
Mountain Ranges	0.179	0.283
	(0.030)	(0.031)
River Basin	0.767	0.420
	(0.089)	(0.089)
Remoteness	-3.857	-3.341
	(0.155)	(0.168)
Constant	9.129	11.180
	(0.992)	(1.029)
N	36046	6110
Pseudo R^2	0.476	0.143

Notes: This table reports the estimation of the logistic regression model, where the log odds ratio of receiving the treatment (having a border) is modeled as linear in a number of the geographical covariates. *Distance* is (log) bilateral distance between origin and destination in kilometers, *Insularity* takes value 1 if one of the regions is an island. *Mountain Ranges* is the highest difference in elevation between two regions in metres (difference between highest point and lowest point, in logs), *River Basin* takes value 1 if the region pair shares a river basin and *Remoteness* is the log of the average remoteness of the origin and the destination regions.

where N_{nm} is the log of the number of borders faced by the region pair, and u_{nm} is a zero-mean error term uncorrelated with the regressors.[25] Since this regression controls for both the number of borders and the pre-treatment covariates, we can use the estimated value $\hat{\beta}$ as a lower bound for the border effect. If we are also able to show that the fraction of switchers is small, then $\hat{\beta}$ is an unbiased estimate of the border effect.

[25]The number of borders of a given region equals to 268 minus the number of regions within its country plus 1. The smallest number of borders corresponds to the 38 regions of Germany, with 231 borders. The largest number of borders corresponds to the 2 regions of Slovenia and Croatia, with 267 borders. The variable N_{nm} is the (log) sum of the borders of the region pair. Thus, the values of N_{nm} lie between $\ln(231 \times 2) = 6.1355$ and $\ln(267 \times 2) = 6.2804$.

The question we address now is that of choosing the right sample to estimate the regression model in Equation (2.11). One might initially think that we should use the entire sample. After all, using all the information available is a principled way to proceed. However, Imbens and Rubin (2015) show that the sampling variance of the estimator $\hat{\beta}$ will be large if the population distribution of covariates is unbalanced between treated and control units. Before using regression methods on the entire sample, one needs to ensure that there is enough balance or overlap in the two covariate distributions.

To determine whether there is sufficient overlap in our entire dataset, the left panel in Figure 2.6 plots the distribution of the estimated propensity score for control units (empty bars) and for treated units (blue shaded bars). The overlap of the propensity score distribution for treated and control units is small. Thus, we trim the data to drop units with extreme values for the estimated propensity score, following the procedure recommended by Crump et al. (2009). This trimming procedure amounts to dropping all observations for which the propensity score is above or below a threshold determined following a variance criterion.[26] We apply this methodology to our sample and obtain a value of the threshold equal to 6.5 percent. We trim the sample accordingly and re-estimate the propensity score. Column (2) in Table 2.3 presents the results. The R-squared is now smaller, showing that our covariates explain now a smaller fraction of the variation in the border assignment, as expected after dropping observations in the extremes of the propensity score distribution. The right panel of Figure 2.6 shows that the distribution of the propensity score across control and treated pairs has a much higher overlap after trimming the initial sample.

[26] The idea in Crump et al. (2009) is to choose a subset A of the covariate space X so that there is substantial overlap between the covariate distribution for the treated and control units. Crump et al. (2009) use the asymptotic efficiency bound for the efficient estimator for the treatment effect in subset A to choose the trimming threshold. The intuition is that if there is a value of the covariate space such that there are few treated units relative to the number of controls, for this value the variance for an estimator for the average treatment effect will be large. Therefore, excluding units with such covariate values should improve the asymptotic variance of the efficient estimator.

Table 2.4: Summary statistics of covariates by block

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd
Distance	154.36	186.07	240.35	298.82	349.83	383.02	440.94	480.01	446.70
	61.03	74.23	93.43	121.79	143.55	143.03	161.45	136.84	61.64
Insularity	0.01	0.01	0.01	0.02	0.04	0.07	0.08	0.12	0.22
	0.08	0.12	0.12	0.15	0.20	0.25	0.28	0.33	0.42
Mountain Ranges	208.38	291.05	351.19	466.84	533.75	549.99	596.98	735.32	1244.59
	232.38	320.38	376.25	457.99	528.13	545.14	561.71	681.78	888.16
River Basin	0.29	0.28	0.21	0.19	0.17	0.14	0.12	0.10	0.06
	0.45	0.45	0.41	0.39	0.37	0.35	0.32	0.31	0.24
Remoteness	1169.05	1097.32	1092.09	1087.40	1081.35	1051.59	1038.82	1002.73	938.72
	307.02	268.01	273.50	276.93	275.84	249.16	229.19	187.51	140.79
Propensity score	**0.20**	**0.31**	**0.44**	**0.57**	**0.66**	**0.72**	**0.78**	**0.84**	**0.89**
	0.04	0.04	0.04	0.04	0.02	0.02	0.02	0.02	0.01
N	323	408	515	698	507	660	1062	1582	354

Notes: This table reports the mean and standard deviation of each geographical covariate and the propensity score in each block. *Distance* is bilateral distance between origin and destination in kilometers, *Insularity* takes value 1 if one of the regions is an island. *Mountain Ranges* is the highest difference in elevation between two regions in metres (difference between highest point and lowest point), *River Basin* takes value 1 if the region pair shares a river basin and *Remoteness* is the average remoteness of the origin and the destination regions.

There are two possible methods to perform inference using the propensity score that are recommended by Imbens and Rubin (2015): matching and blocking. In our setting, we think a blocking estimator, based on grouping region pairs with similar propensity score values, is more appropriate. Thus, we build subsamples of pairs such that the border probability is similar. We call these subsamples blocks. To create them, we follow the procedure recommended by Imbens and Rubin (2015), using the algorithm in Becker and Ichino (2002). This algorithm starts by splitting the sample into 5 equally spaced intervals of the propensity score and then testing whether the average propensity score of treated and control units does not differ much within blocks. If it does, the algorithm splits the interval in half and tests again, until the average propensity score of treated and control units no longer differs within blocks. Starting from the trimmed sample, this procedure delivers nine blocks. We have ordered these blocks such that the propensity score is increasing.

Table 2.4 reports the summary statistics of the covariates and the propensity score by block. Recall that there are two factors that reduce the sampling variance of the estimates: (i) the number of observations; and (ii) the balance between treated and control groups. The number of observations varies substantially across blocks, ranging from 323 in Block 1 to 1582 in Block

Table 2.5: Balancing test of covariates by block

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Distance	-22.24	8.207	5.049	4.693	17.13	-11.79	-24.16	-33.09	28.87
	(8.077)	(8.126)	(8.290)	(9.269)	(13.42)	(12.40)	(12.07)	(9.763)	(9.636)
Insularity	-0.00990	0.0206	0.0166	0.0187	0.0302	0.0125	-0.00573	-0.0613	-0.00663
	(0.0105)	(0.0132)	(0.0103)	(0.0110)	(0.0190)	(0.0216)	(0.0208)	(0.0234)	(0.0660)
Mountain Ranges	-31.46	25.23	-16.62	-120.6	-148.1	-114.3	-96.62	-139.1	45.43
	(31.06)	(35.09)	(33.39)	(34.56)	(49.01)	(47.07)	(41.96)	(48.70)	(140.6)
River Basin	0.0528	-0.0328	-0.00768	0.0366	0.0247	-0.0101	-0.00522	-0.0304	0.0285
	(0.0608)	(0.0495)	(0.0362)	(0.0296)	(0.0350)	(0.0303)	(0.0242)	(0.0219)	(0.0382)
Remoteness	-109.7	51.41	30.83	20.83	44.75	-5.539	-21.53	-54.87	59.66
	(40.65)	(29.26)	(24.24)	(21.06)	(25.75)	(21.61)	(17.15)	(13.36)	(22.06)
N	323	408	515	698	507	660	1062	1582	354

Notes: This table reports the difference in means between treated and control region pairs for each geographical covariate by block (defined as control minus treated). Standard errors in parenthesis. *Distance* is bilateral distance between origin and destination in kilometers, *Insularity* takes value 1 if one of the regions is an island. *Mountain Ranges* is the highest difference in elevation between two regions in metres (difference between highest point and lowest point), *River Basin* takes value 1 if the region pair shares a river basin and *Remoteness* is the average remoteness of the origin and the destination regions.

8. Blocks also vary substantially in terms of their propensity score, ranging from 20 percent in the first block to 89 percent in the ninth one. Blocks 3, 4 and 5 are the most balanced ones with a propensity score of 44, 57 percent and 66 percent, respectively.

Table 2.5 reports the t-statistic from a difference in means test between treated and controls (test is defined as control mean minus treatment mean). Covariates are well balanced within blocks, with only small differences in means that do not seem to follow a systematic pattern. If the covariates were perfectly balanced within blocks, we could estimate causal effects as if assignment was random within each block. That is, we could compare the means of the international and intranational pairs controlling only for the number of borders. Since three out of five covariates are continuous, however, it is unavoidable to have some small variation in covariates within blocks. In this case, Imbens and Rubin (2015) recommend that these comparisons also control for covariates. Thus, we shall estimate the regression model in Equation (2.11) for each of the blocks.

To give a sense of the composition of the blocks in terms of regions, Figure 2.7 shows the frequency with which each region appears (as a part of a pair) within the control and treated groups in block 4. This block has an average propensity score of 57 percent. That is, region pairs within this block had roughly an equal chance of having a border than not having one. In this block we find regions from all around Europe both in the treated and in the control units. The composition of regions changes across blocks. As we would expect, blocks 1 and 2 source mostly from region-pairs that are at short distances while blocks 7, 8 and 9 contain regions located in the largest countries, since region-pairs are, on average, further away. The figures for all the blocks can be found in the Appendix.

Let us go back to our example of Catalonia, Languedoc-Roussillon and Valencia. Figure 2.8 shows all the pairs that contain Catalonia (shown in grey) in our sample. The color of each region represents the block in which the corresponding pair is located. White-colored regions are pairs that have been dropped after trimming, for which the probability of a border was close to 1. There is no pair that includes Catalonia in block 1, indicating that the probability of Catalonia having a border with any of its neighbours was always 20 percent or larger. Languedoc-Roussillon is in block 5, where the average probability of a border is about 66 percent; and Valencia is in block 3, where the average probability of a border is about 44 percent.

Figure 2.8 allows us to illustrate our identification strategy, and the motivation behind our approach. Notice that block 7 contains intranational pairs, in Spain, as well as international pairs, in France and Portugal. The former will be used as control units, while the latter will be used as treated units. Region pairs in block 7 have a probability close to 78 percent of being separated by a border. Given that this probability is very similar across treated and control units, the difference in trade between them can be interpreted as the causal effect of the border.

We have now constructed the samples we needed to estimate the border effect. Before using them, though, we need to assess how important is the participation bias in these samples (recall Equation (2.7) and the discussion

Table 2.6: Participation rate: Control vs. Treated

	All	Trimmed	Blocks								
			(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Part. rate control	0.968	0.976	1	.997	.993	.987	.968	.968	.936	.952	.915
Part. rate treated	0.617	0.946	.993	.996	.996	.969	.947	.957	.95	.928	.894
N	72092	12220	646	816	1030	1396	1014	1320	2124	3164	710

Notes: This table reports the share of region pairs that engage in positive trade in our regional trade dataset (participation rate) for the region pairs in the treated and control groups.

after it). Table 2.6 shows how participation rates differ between treated and control groups in the entire sample, the trimmed sample and in each of the blocks. Participation rates among control units are high in all the samples. In the entire sample, however, the participation rate among treated units is only 61.7 percent. This must be due to the fact that many international pairs are far away and likely to have a border. Indeed, participation rates in the trimmed sample increase dramatically among the treated, becoming quite close to those in the control group. The participation rates within blocks are even more balanced. Thus, we conclude that the participation bias cannot be large within these blocks. Remarkably, our construction of blocks has achieved an almost perfect balance in participation rates without using any outcome variables in the procedure. This provides additional support for our chosen empirical strategy.

2.4 Causal effect of borders on trade

Finally we are ready to present our results. We show first our estimation of the average border effect and we continue with the estimation of the border effect across industries. Finally, we present our estimation of the effect recent borders.

2.4.1 Average Border effect

Table 2.7 shows the results of estimating Equation (2.11) for each of the blocks. Recall that the estimated coefficient on the border dummy is the

Table 2.7: Average border effect

Dep. Var: $ln(S_{n,m})$	Block 1 (1)	Block 2 (2)	Block 3 (3)	Block 4 (4)	Block 5 (5)	Block 6 (6)	Block 7 (7)	Block 8 (8)	Block 9 (9)
Border	-1.786	-1.721	-1.699	-1.768	-1.686	-1.796	-1.687	-1.754	-1.858
	(0.182)	(0.178)	(0.175)	(0.175)	(0.238)	(0.289)	(0.268)	(0.290)	(0.201)
Number of Borders	7.058	6.695	7.041	10.779	11.294	11.833	9.234	8.091	0.420
	(1.756)	(1.970)	(2.034)	(1.730)	(2.064)	(2.783)	(2.792)	(3.063)	(2.944)
Geographic covariates	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N	645	813	1024	1364	968	1267	2011	2948	637
R^2	.572	.533	.501	.47	.375	.388	.31	.285	.299

Notes: Standard errors clustered at the country-pair level, are in parentheses. Dependent variable is the (log) normalized market share of n in m. *Border* is a dummy for international border. *Number of borders* is the (log) sum of the number of borders that are faced by n and m.

log reduction in the normalized market share caused by the border, that is, the average border effect within the block. This effect is large, statistically significant at the one percent level, and it varies little across blocks. The border effect ranges from a minimum of -1.686 in block 5 to a maximum of -1.858 in block 9, which indicate that borders reduce trade to somewhere between 18.5 $(= \exp\{-1.686\})$ and 15.6 $(= \exp\{-1.858\})$ of their potential.

Table 2.7 also shows the effect on normalized market shares caused by the number of borders. Recall that the coefficient on this variable measures the elasticity of the normalized market share with respect to the number of borders. This elasticity varies across blocks, ranging from 6.695 in block 2 to 11.833 in block 6. Since $N_{nm} \in [6.1355, 6.2804]$ in our sample, we have that the difference in market shares caused by differences in the number of borders might be substantial. To put an upper bound to this difference, compare the region pair containing the two Slovenian regions, which is in block 1, with a region pair containing two German regions in the same block. According to our estimates, the normalized market share for the Slovenian pair is about 2.78 $(= \exp\{7.058 \times 0.1449\})$ larger than that of the German pair. Thus, our estimates reveal an additional important channel through which the border assignment affects trade. It is not only whether a border is assigned to a specific region pair that matters, but also how many borders are assigned to each region in the pair.

Let us now use these results to be a bit more precise about the counterfactual scenario discussed in the introduction, in which the French-Spanish border is southwest rather than north of Catalonia. Recall that the region pair (Catalonia, Languedoc-Roussillon) is in block 5, and that the change in the French-Spanish border reduces the number of borders of Catalonia by 7 and for Languedoc-Roussillon by 1. Then, we can compute the effect of this change in the border as the product of two separate effects: (i) the average border effect which increases the market share by a factor 5.398 $(= \exp\{1.686\})$; and (ii) the number-of-borders effect which lowers the market share by a factor 0.839 $(= \exp\{11.294(-0.0155)\})$. Thus, our estimates indicate that Catalonia's market share of the Languedoc-Roussillon market would be $4.530(= 5.398 \times\ 0.838)$ larger than it is today. Since the region pair (Catalonia, Valencia) is in block 3 and the change in the French-Spanish border increases the number of borders of Valencia by 1, Catalonia's share of the Valencia market would be 0.165 $(= \exp\{-1.699 + 7.041(-0.0119)\})$ smaller than it is today. These numbers are a bit different from those we showed in the introduction because the latter did not take into account the number-of-borders effect.

Table 2.8 reports the average border effect, after aggregating our regression results by block. We present two possible average treatment effects, weighting the coefficients by the size of the block (row 1) and weighting by the number of treated units in each block (row 2) (see Imbens and Rubin (2015)). The average effect of the border is negative and large in magnitude,

Table 2.8: Average Border Effect (Average treatment effect)

	Estimated β^{ATE}	
	All controls	Without number of borders
Weights: Size of blocks	-1.744	-1.299
Weights: Treated pairs	-1.747	-1.303

Notes: Average treatment effect calculated by computing the weighted average of the estimated coefficient of the *Border* dummy. The first row uses the number of observations in each blocks as weights, while the second row uses the number of treated units in each block.

Table 2.9: Average Border effect using the full and trimmed samples

Dependent variable: $log(S_{nm})$	Full sample (1)	Trimmed sample (2)
Border	-1.968	-1.716
	(0.211)	(0.184)
Number of Borders	7.944	8.346
	(1.807)	(1.647)
Geographic Covariates	Yes	Yes
N	46236	11677
R^2	.482	.642

Notes: Standard errors clustered at the country-pair level, are in parentheses. Dependent variable is the (log) normalized market share of n in m. *Border* is a dummy for international border. *Number of borders* is the log of the total number of borders that are faced by n and m.

and the weighting method does not make much of a difference. Our findings suggest that the border reduces trade between two regions to 17.5 percent of what they would trade without the border ($\exp\{-1.744\} = 0.175$).

A key step in our identification strategy is to control for the number of borders. This matters not only in itself as argued already, but also to avoid a selection bias problem when estimating the average border effect. As discussed in section 3.1, region pairs with many borders tend to have larger market shares and tend to be over-represented among international pairs and under-represented among intranational pairs. This creates a positive selection bias that makes the observed difference in average market shares smaller (in absolute value) than the true average border effect. To show that this source of selection bias is relevant, the second column of Table 2.8 reports the estimated average border effect that we would obtain if we failed to control for the number of borders. This biased estimate of -1.299, would lead us to believe that the border reduces normalized market shares to 27.3 percent of its potential instead of the true estimate of 17.5 percent.

Another key step in our identification strategy is trimming the data set. Table 2.9 shows the results of running Equation (2.11) with the entire sample and the trimmed sample. For the full sample we obtain an estimate of -1.968, which would lead us to believe that the border reduces normalized market shares to 14 percent of its potential. For the trimmed sample we obtain an estimate of -1.716 which is essentially the same as the one provided by the blocking estimator. This is consistent with our finding that the average border effect varies very little across region pairs with different propensity scores.

2.4.2 Border effect across industries

The average border effect may hide some cross-industry heterogeneity.[27] We report now the results of estimating Equation (2.11) industry by industry. Importantly, we can use the estimated propensity score and the same blocks, since both are constructed from region-pair covariates that are constant across industries.

Table 2.10 presents the results for all industries. The border effect is negative and statistically significant in all blocks in all industries (coefficients represented with confidence intervals in figure A.2.17 in the Appendix). As we could anticipate, the average border effect masks some heterogeneity. The industry "Food, Beverage and Tobacco", in column (10) of row (3), has a weighted coefficient of -2.095, meaning that the border effect is 0.123. The industry "Textiles", in column (10) of row (4), has a weighted coefficient of -.945, implying that the border effect is 0.389.

Our industries are very aggregated and it is difficult to say much about these differences in the border effect. But we do notice that lower border effects, of around -1.4 are estimated in Chemicals, Metals and Vehicles. While higher border effects, of around -1.6, are found in Wood and Cork Products and Paper, Non Metals, Machinery and Agriculture. This is suggestive of an increasing border effect for more differentiated or more transformed goods.

[27] Using total trade flows misses the fact that industries have varying trade cost elasticities (Chen and Novy, 2012) and select into geographies taking into account border related costs. Therefore estimates that employ aggregated data at the industry level risk suffering from compositional bias (Hillberry, 1999).

Table 2.10: Border effect across industries and blocks

INDUSTRY	Block 1	Block 2	Block 3	Block 4	Block 5	Block 6	Block 7	Block 8	Block 9	ATE: W	ATE: T
1. AGRI	-1.851***	-1.813***	-1.659***	-1.384***	-1.241***	-1.611***	-1.413***	-1.620***	-1.995***	-1.578	-1.559
2. MINE	-1.714***	-2.017***	-1.607***	-1.592***	-1.413***	-1.374***	-1.160***	-1.019***	-2.054***	-1.471	-1.395
3. FBT	-2.488***	-2.464***	-2.163***	-2.084***	-2.034***	-2.024***	-1.977***	-1.954***	-2.196***	-2.095	-2.047
4. TEX	-1.333***	-1.195***	-0.714***	-1.053***	-0.830***	-0.915***	-0.714***	-0.839***	-1.307***	-0.945	-0.904
5. WOOD	-1.532***	-1.641***	-1.366***	-1.429***	-1.369***	-1.360***	-1.488***	-1.588***	-1.828***	-1.499	-1.505
6. COKE/PET	-2.025***	-1.314***	-1.221***	-0.787***	-0.702***	-0.776***	-0.507***	-0.601***	-1.592***	-0.995	-0.866
7. CHEM	-1.373***	-1.278***	-1.206***	-1.388***	-1.080***	-1.267***	-1.298***	-1.308***	-1.249***	-1.282	-1.280
8. NON-MET	-1.936***	-1.975***	-1.850***	-2.030***	-1.767***	-1.951***	-1.739***	-1.834***	-2.122***	-1.886	-1.874
9. MET	-1.239***	-1.254***	-1.372***	-1.514***	-1.400***	-1.363***	-1.218***	-1.459***	-1.719***	-1.384	-1.400
10. MACH	-2.260***	-1.841***	-1.834***	-1.698***	-1.286***	-1.511***	-1.364***	-1.619***	-1.430***	-1.627	-1.565
11. VEH	-1.545***	-1.303***	-1.366***	-1.406***	-1.091***	-1.210***	-1.233***	-1.338***	-1.762***	-1.330	-1.321
12. OTHER	-2.029***	-1.589***	-1.361***	-1.494***	-1.372***	-1.283***	-1.272***	-1.165***	-1.716***	-1.406	-1.348
Aggregate BE	-1.786	-1.721	-1.699	-1.768	-1.686	-1.796	-1.687	-1.754	-1.858	-1.744	-1.747

Notes: This table reports the estimated border effect (coefficient on dummy Border, in regression equation (2.11)) by industry (rows) and block (column).The last two columns report the average border effect computed using as weights the size of the block (ATE: W) and the number of treated region pairs (ATE: T). The last row (Average BE) reports the average border effect across industries, as reported in table 2.7.

The last row of Table 2.10 reports the average border effect estimated in the previous subsection. In all industries but two this average effect is larger than the industry border effect. In the first blocks, columns 1 to 4, the estimates of the border effect for some industries are below the average and some are above. However, in blocks 5 to 8 we see that the estimates of the border effect for almost all industries are below the average. At first sight, this seems puzzling, since the average border effect is estimated by aggregating the industry-level data. The explanation for this observation is the imbalance in participation rates between treated and controls in this second set of blocks. As explained in the previous section, this generates a participation bias that leads to an underestimation of the border effect.[28]

2.4.3 Effects of post-1910 borders

We next examine whether the border effect varies with the age of the border. Our sample contains borders that were created several centuries ago, such as the French-Spanish border, together with borders that were put in place only some decades ago, like the border between the Czech Republic and Slovakia that was established in 1993. It is plausible to think that effects of these

[28]Figure A.2.18 in the Appendix plots the differences in participation (share of trading pairs) between treated and control units in each industry and block. As expected, participation rates are very similar in all industries in blocks 1 to 3, but much larger for control pairs in other blocks.

borders might be quite different.

Figure 2.9 shows borders in Europe in 1910 and 2010. The 1910 set of borders is the culmination of a process of political integration that included, for instance, the unification of Italy and Germany. After 1910, this trend reversed. The 2010 set of borders shows the effects of a process of political disintegration which included, for instance, the collapse of the Austro-Hungarian empire and the former Yugoslavia and Czechoslovakia. Indeed, about one third of the region pairs that shared a government in 1910 no longer share a government in 2010.

We take 1910 as our reference year and split our sample of region pairs into four groups, according to their border history. The largest group consists of regions that are in different countries both in 1910 and in 2010, and contains 90 percent of our observations. The second largest group consists of regions that have always been in the same country, and contains 6.3 percent of our observations. The third largest group consists of regions that were in the same country in 1910, but are no longer in the same country in 2010. This group contains about 3.1 percent of our observations. The final and smallest group consists of regions that were in different countries in 1910 and now are in the same country. This group contains only 0.5 percent of our observations.

To measure the effects of adding a new border, we compare outcomes between the groups that were in the same country in 1910. As mentioned, about a third of the regions who shared a country in 1910, no longer do so in 2010. Thus, we have a good balance between treated and controls to perform inference. It would be interesting also to measure the effects of removing an old border by comparing outcomes between the groups that were in a different country in 1910. Unfortunately for our purposes, almost none of the regions in these two groups share a country today. There is simply too much imbalance between treated and controls to perform inference.[29]

[29]Previous studies in the literature have found persistent effects of bygone borders on trade. Nitsch and Wolf (2013) find persistence of the former inner German border on current intra-German trade by road, although the estimated border effect has been declining

Table 2.11: Propensity Models for region pair with border 1910=0

Dependent Variable: Border	Full sample (1)	Trimmed sample (2)
Distance	2.254	2.414
	(0.108)	(0.134)
Insularity	0.270	0.257
	(0.186)	(0.192)
Mountain Ranges	0.071	-0.007
	(0.052)	(0.055)
Same River Basin	1.835	1.914
	(0.120)	(0.136)
Remoteness	-2.293	-2.215
	(0.272)	(0.299)
Constant	1.127	0.065
	(1.844)	(1.965)
N	3422	2630
Pseudo R^2	0.222	0.139

Notes: This table reports the estimation of the logistic regression model, where the log odds ratio of receiving the treatment (having a border) is linear in the geographical covariates. *Distance* is (log) bilateral distance between origin and destination in km, *Insularity* takes value 1 if one of the regions is an island. *Mountain Ranges* is the highest difference in elevation between two regions in metres (difference between highest point and lowest point, in logs), *River Basin* takes value 1 if the region pair shares a river basin and *Remoteness* is the log of the average remoteness of the origin and the destination regions.

We start with a sample containing the two groups that were in the same country in 1910. Starting from this sample, we repeat the steps explained in section 3. We re-estimate the propensity score and we trim the sample to achieve a good overlap between treated and control units. Table 2.11 reports the estimation of the propensity score model for the full sample and the trimmed sample, whereas Figure 2.10 shows the distribution of the propensity score among treated and control units. We then create blocks and report the summary statistics of the covariates and the balancing test in Tables A.2.6 and A.2.7 in the Appendix. This procedure now generates 6 blocks.

Table 2.12 reports the results of estimating Equation (2.11) with this subsample. We find a negative and significant border effect for post-1910 borders, albeit smaller than the average border effect without conditioning on historical borders. The average border effect is -1.261 (-1.221) weighting

over time. Beestermöller and Rauch (2018) explore how the trading capital accumulated between members of the Astro-Hungarian empire still drives preferential trade between European countries even after the Fall of the Iron Curtain.

Table 2.12: Average border effect when Border in 1910=1

Dep. Var: $ln(S_{n,m})$	Block 1 (1)	Block 2 (2)	Block 3 (3)	Block 4 (4)	Block 5 (5)	Block 6 (6)
Border	-1.439	-1.165	-1.129	-1.290	-1.169	-1.189
	(0.259)	(0.305)	(0.301)	(0.415)	(0.322)	(0.405)
Number of borders	7.503	7.325	7.714	7.239	8.364	14.124
	(2.120)	(2.765)	(3.502)	(4.762)	(4.696)	(4.240)
Geographical covariates	Yes	Yes	Yes	Yes	Yes	Yes
N	1530	1082	894	703	554	298
R^2	.612	.505	.432	.443	.353	.418

Notes: Standard errors clustered at the country-pair level, are in parentheses. Dependent variable is the (log) normalized market share of n in m. *Border* is a dummy for international border. *Number of borders* is the (log) sum of the number of borders that are faced by n and m.

by size of block (treated). This means that the border reduces the market share to 28.3 percent (29.5 percent) of its potential. These findings show that borders that have been in place for less than a century have large trade reducing effects, although smaller than those of older borders.

2.5 Concluding remarks

In this paper we have built a European regional trade dataset and we have estimated the average border effect on trade flows using a new identification framework. Our results show that the effects of country borders on trade flows within Europe are large. Take two similar region pairs, the first one containing regions in different countries and the second one containing regions in the same country. The market share of the origin region in the destination region for the international pair is only 17.5 percent that of the intranational pair. We refer to this estimate as the average border effect. It seems, then, that we are still far from having a single market in Europe. Country borders have created a national bias in preferences and a national cost advantage that penalize international trade and foster intranational trade. How do country borders affect trade flows? What are the welfare implications? Providing satisfactory answers to these questions is a major research goal on its own, one which is likely to deliver important policy implications for Europe.

We view our contribution as part of a broader research program on the effects of country borders within Europe. To start with, we are currently using our new dataset and the empirical framework developed here to measure the effect of regional governments. In this paper we have focused on the effects of country governments. Yet, regional governments also make decisions about procurement, infrastructure, laws and regulations and so on. What is the effect of regional borders on trade? This project will allow us to obtain a more detailed and precise picture of the effects of different types of political borders.[30]

The broader research program we envision should go beyond estimating the size of border effects, and also try to disentangle the relative importance of the different channels through which country borders affect trade.[31] Some insight can be obtained by looking at differences in the estimates across industries and between new and old borders provided here. But this only scratches the surface. One would like to have precise answers to questions such as: How much would the border effect be reduced if the European Union were able to eliminate the large observed national bias in government procurement? How much would the border effect be reduced if the European Union were able to build a truly European transportation network? Answering these and related questions is only possible with a reliable empirical strategy that addresses the endogenous assignment of borders such as the one developed in this paper.

The research program we have in mind should also go beyond trade flows and examine the effects of country borders on other economic and social

[30] There are a few papers that have looked at the effects of regional borders using the gravity framework. For instance Wolf (2000), Coughlin and Novy (2012) and Garmendia et al. (2012).

[31] There are some papers that have explored a few these channels: Turrini and van Ypersele (2010) explore the effects of judicial systems, Bailey et al. (2020), Combes et al. (2005) and Fukao and Okubo (2004) explore the role of social and business networks, Schulze and Wolf (2009) focus on ethno-linguistic factors, and Chen (2004) analyzes technical barriers to trade and product-specific information costs increase the effect of borders on trade.

interactions. Country borders have implications that go far beyond trade flows. The approach developed here could also be used to measure the effect of borders on migration and investment flows, cultural values, travel and tourism, cooperation in research projects, joint sports activities, and so on. It would be useful to have a broader picture of how country borders within Europe affect economic and social interactions among its regions.

Carrying out this project also made it clear to us that we need a richer theory. Our results suggest that modeling borders is crucial to understand the patterns of intranational and international trade. We have wonderful quantitative theories of trade that realistically model the incentives and constraints faced by consumers and firms. But these quantitative theories rarely include a realistic description of the incentives and constraints faced by governments. If modeled at all, governments either act mechanically or solve some unrealistic social planner problem. How are procurement decisions made? How are infrastructures chosen? How are laws and regulations decided and enforced? Only a realistic and detailed modeling of the behavior of governments can shed light on the channels through which political borders affect trade and welfare. Fortunately, there is a lot of excellent work on the political economy of trade policy to draw upon for this purpose (See, for instance, Grossman and Helpman (2001)).

Much less developed is the theory of country borders. It is here where we have felt more at sea when working on this project. Understanding the border assignment is key to develop a sound identification strategy. And yet there does not exists a theory of borders that is developed at the same level of sophistication, say, than the theory of international trade. There exist some classic approaches to modeling and understanding country formation (see Spolaore and Alesina (2003)); and some recent ones too (see Cervellati et al. (2019) and Gancia et al. (2020)). But these theoretical frameworks can only be seen as promising prototypes, much work is needed to develop them into a fully fledged theory capable of guiding quantitative and empirical research.

Figure 2.5: Average market share and number of borders

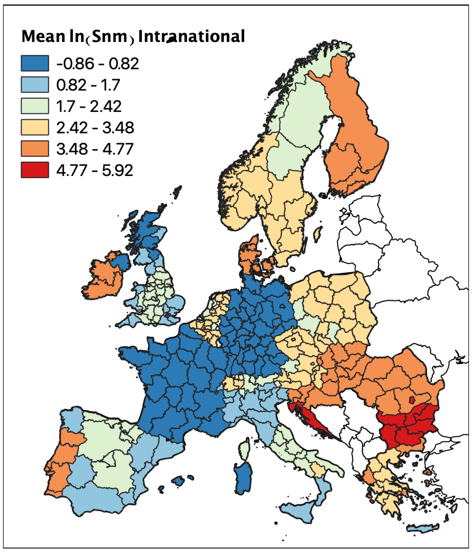

A) Intranational market share

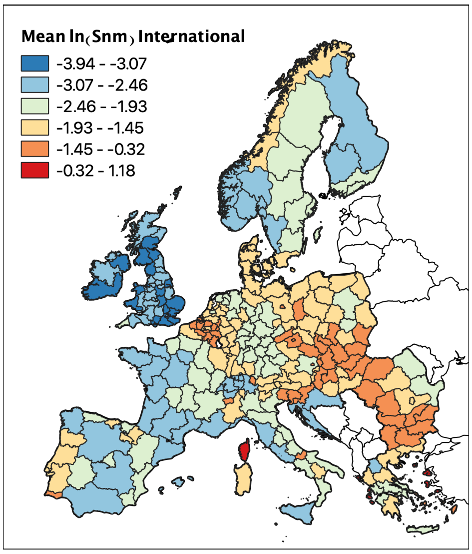

B) International market share

Notes: The figure shows the average market share of each region with its intranational partners (panel A) and with its international partners (panel B). The color shading represents the value of this average, with cooler colours representing lower market shares and warmer colors representing higher market shares

Figure 2.6: Histogram of propensity score

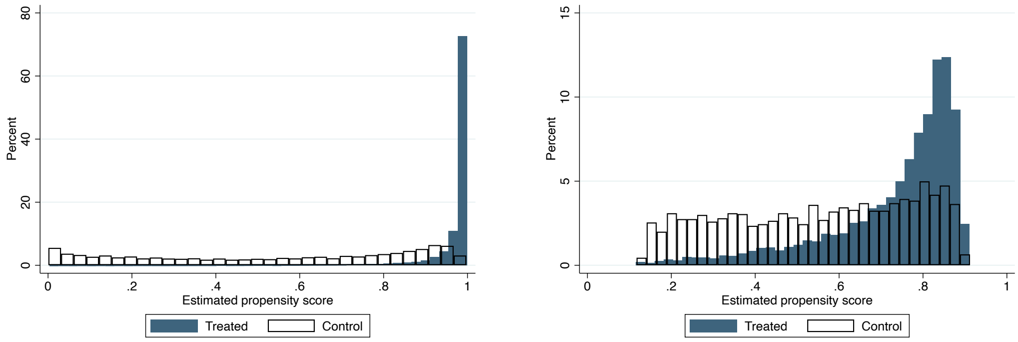

A) All region pairs **B) Trimmed sample**

Notes: This figure shows the distribution of the estimated propensity score, probability of having a border, for control units (empty bars) and for treated units (blue shaded bars). Panel A reports the results using the full sample while panel B reports the results using the trimmed sample (dropping region pairs with extreme estimated probability of having border).

Figure 2.7: Composition of regions in block 4

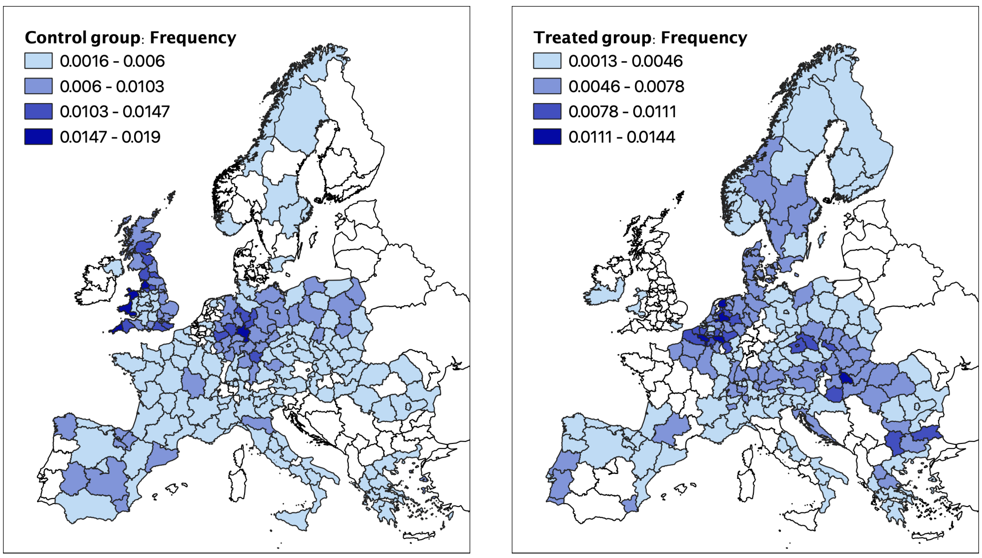

A) Control group **B) Treated group**

Notes: This figure shows the regions that are part of the block. The shading represents the frequency with which each region appears (as a part of a pair) in the control group (first panel) and treated group (second panel) in the block.

Figure 2.8: Distribution of Blocks for region-pairs with Catalonia

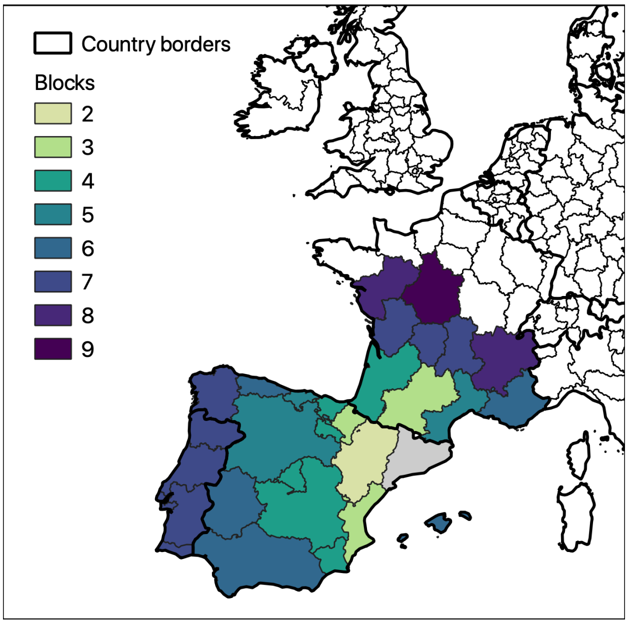

Notes: This figure shows the regions that are part of a pair that includes Catalonia in the trimmed sample. The colors represent the block in which each region pair is included. The blocks are ordered as increasing in the propensity score. Darker shading represents higher probability of having a border.

Figure 2.9: Recent and old borders

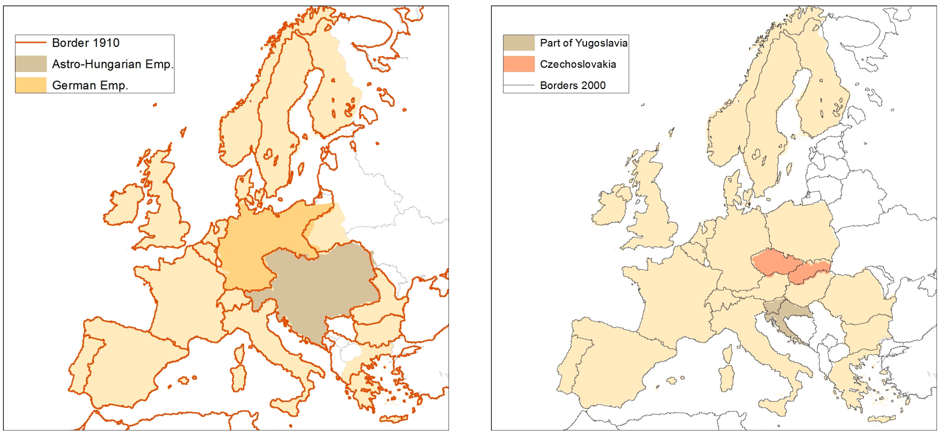

A) Borders in 1910 **B) Borders in 2010**

Notes: This figure shows European borders in 1910 (panel A) and in 2010 (panel B).

Figure 2.10: Histogram of propensity score

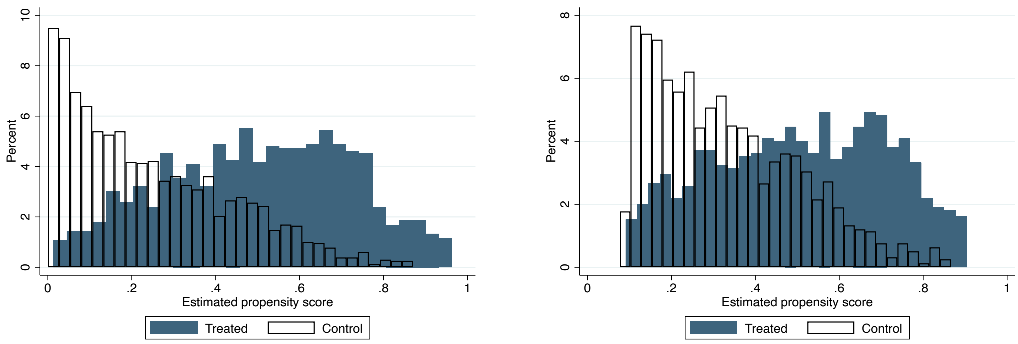

A) All region pairs **B) Trimmed sample**

Notes: This figure shows the distribution of the estimated propensity score, probability of having a border, for control units (empty bars) and for treated units (blue shaded bars). Panel A reports the results using the full sample while panel B reports the results using the trimmed sample (dropping region pairs with extreme estimated probability of having border).

Chapter 3

EXPLORING EUROPEAN REGIONAL TRADE

Joint with Jaume Ventura (CREi, Universitat Pompeu Fabra and Barcelona School of Economics) and Marta Santamaria (University of Warwick).

3.1 Introduction

How do regions trade with each other? We know much about trade across countries thanks to the availability of detailed customs data. We know much less about trade within countries. In this paper, we use the dataset we constructed in Santamaría et al. (2020) to systematically explore for the first time trade patterns across and within European regions.

Europe is a great laboratory to explore regional trade flows. One reason is that Europe is large, as it contains more than 500 million people and it produces about 20 percent of world GDP. Another reason is that European regions exhibit a lot of heterogeneity, as shown in Figure 3.1 using data from 2011 (the starting period of our dataset). The top panel shows the distribution of per capita GDP, which ranges from a low of 3,200 euros in Northwestern Bulgaria to a high of 85,330 euros in Central London. The middle panel shows the distribution of populations, which ranges from a low of 126,761 inhabitants in Valle d'Aosta to a high of 11,852,851 inhabitants in Île de France. The bottom panel shows the distribution of geographical areas, which range from a low of 160 Km^2 in Brussels to a high of 226,716 Km^2 North/East Finland.

Figure 3.1: Heterogeneity across European regions

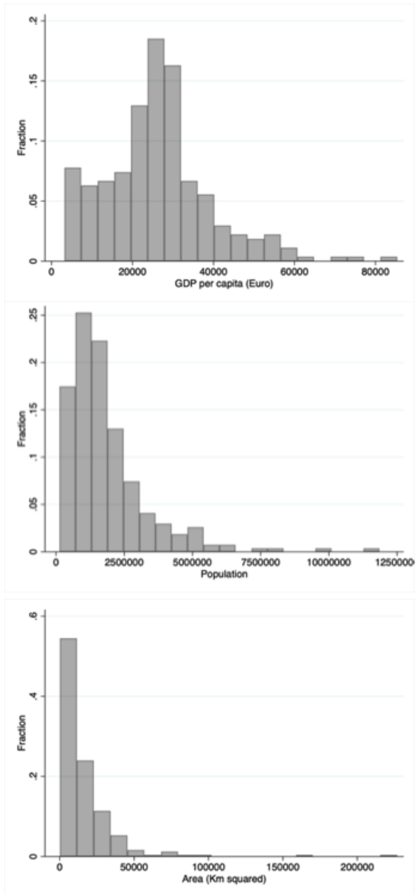

The dataset constructed in Santamaría et al. (2020) is based on the European Road Freight Transport survey which collects data on truck shipments of goods in agriculture, manufacturing and mining. Thus, the dataset covers trade in goods by road, which according to Eurostat is about half of all European trade in goods. The dataset covers 269 regions from 24 European countries between 2011 and 2017 disaggregated into 12 different industries.

An important aspect of this dataset is that it allows us to measure trade flows both across and within regions. Thus, for each year/industry, we have a complete matrix of bilateral trade including the diagonal entries.

The first and more salient aspect of European regional trade is that it has a strong home and country bias. Consider a shipment originating from a randomly selected European region. The probability that this shipment has a destination inside the origin region (i.e. home trade) is 40 percent. The probability that this shipment has a destination outside the origin region but inside the country of the origin region (country trade) is 41 percent. The probability that this shipment has a destination outside the country of the origin region (foreign trade) is therefore only 19 percent. To evaluate these numbers, one must recognize that the size of the destination markets is quite different. The home market is smaller than the country market, and the latter smaller than the foreign market. When we correct for size,[1] we find enormous differences in the magnitudes of these types of trade. In particular, home, country and foreign trade are 469.5, 11.22 and 0.44 times what one would predict knowing only the sizes of the origin and destination markets.

The second salient aspect of European regional trade is the importance of geographic distance and national borders. The ranking of home > country > foreign trade suggests that these factors are important. Foreign trade involves sellers and buyers that are farther away and do not share the same government. Both of these factors are known to have negative consequences for trade. We show that a parsimonious gravity model that uses only national borders and distance can explain about two-thirds of the variation in European regional trade. Obviously, a model with these elements is designed to create a bias towards home and country trade. But there is more to this. The importance of borders generates a small-country effect, namely, that regions in small countries trade more within and outside their country. The importance of geographical distance generates a remoteness effect, namely, that regions that are geographically remote should trade more with other

[1] That is, by dividing by the product of the sizes of the origin and destination markets.

regions inside their country, and less with regions outside. We observe that both the small-country and remoteness effects are present in the European regional data.

We consider increasingly sophisticated versions of the gravity model that allow for more flexible specifications of distance and border effects. First, we allow for a variable elasticity of trade to distance. This does not make much of a difference, however. Second, we allow border effects to be different for region pairs that have a common language or currency. We find that both sharing a language and a currency reduce the border effect. Finally, we estimate a different border effect for each country pair. We observe that the border effect is quite heterogeneous. Even though the data suggests that all these refinements are capturing some aspects of the data, they do not add much to the model's ability to explain the variation in the data.

A third salient aspect of European regional trade is that the strong home bias in trade cannot be explained by geographical distance and national borders. There are few observations of home trade, 269 out of 73,361, but these observations stand out for their size since they add to 40 percent of all trade. To determine the source of this home bias, we exploit a special feature of the data. Due to government structure differences, in some countries the regions in our dataset are only statistical regions created for the purpose of sharing data with Eurostat, while in other countries the regions in our dataset coincide with political divisions with different levels of self government. This allows us to test whether the home bias effect emerges in all regions, or whether the home bias effect emerges only when it coincides with political borders. We separate region-pairs by the type of border that divides them, either statistical or political, and show that it is the later and not the former that exhibit a large home bias in trade. Thus, it seems that the home bias is cause by political border. In terms of magnitude, these borders matter as much as national borders.

There is an abundance of papers that use the gravity framework to study trade flows. Head and Mayer (2014) provide an extensive review of this literature and the improvements in the methods since being introduced by Tinbergen (1962). Due to the scarcity of data at the subnational level, most of these studies have focused exclusively on international trade. Among the most notable exceptions are papers that use the commodity flow survey to study intranational trade flows in the United States such as Hillberry and Hummels (2008) and Coughlin and Novy (2012). There exist also other papers that look at intranational trade in other countries, for instance Head and Mayer (2009) for France, Nitsch and Wolf (2013) for Germany and Mori and Wrona (2021) for Japan. All these papers focus exclusively on intranational trade. One contribution of this paper is to provide an integrated view of intranational and international trade, and their interactions for Europe, which includes 24 countries and 269 regions.

Our findings suggest that political borders, both national and regional, are an important determinant of trade. Thus our paper is closely related to a large literature that aims at measuring border effects. The seminal papers in this literature are McCallum (1995), Anderson and Van Wincoop (2003), Chen (2004). Two recent papers that also focus on Europe are Santamaría et al. (2020), from which we borrow the data, and Head and Mayer (2021). The final contribution of this paper is to show that border effects apply to political borders but not statistical ones.

3.2 A first look at the data

In this section we describe our dataset and provide a first look at the patterns of regional trade in Europe. The bottom line is simple: regions trade with themselves much more than with other regions within the same country, and regions trade with regions within the same country much more than with regions in other countries. This ranking of home > country > foreign trade is not surprising, but the magnitude of the differences might be.

3.2.1 The dataset

We use the dataset of regional trade flows across European regions constructed by Santamaría et al. (2020) using the European Road Freight Transport survey. This dataset covers trade in goods among 269 regions from 24 European countries between 2011 and 2017. This trade is disaggregated into 12 different industries that cover essentially all of agriculture, mining and manufacturing.

The European Road Freight Survey collects data adhering to the geographic divisions presented by the Nomenclature of Territorial Units for Statistics (NUTS) classification. The NUTS classification is a hierarchical system for dividing up the economic territory of the European Union, the United Kingdom and the EFTA member countries for the purposes of collection, development and harmonisation of European regional statistics. Our regions are defined by the NUTS2 classification.

The European Road Freight Survey collects data on truck shipments between European countries. One limitation of our data is that it covers trade by road only but not other modes of transportation. With respect to this, we note that about 70 percent of all European trade in goods is inland trade and 30 percent is sea trade. Road trade accounts for about 75 percent of inland trade. This means that we cover about 52 percent of all European trade. There are 13 industries in the European Road Freight Survey that cover all of agriculture, mining and manufacturing. Except for one (*Coal and lignite, crude petroleum and natural gas*), road trade is by far the most prevalent mode of inland transportation. This is why Santamaría et al. (2020) dropped this industry and the dataset contains the remaining 12 industries.

The second limitation of our data is that it does not cover trade with non-European partners. To understand the implication of this, consider a shipment from China to Switzerland that goes through the port of Rotterdam. In country level statistics this would be recorded as a shipment from China to Switzerland. In our survey this would be recorded as a shipment from the Netherlands to Switzerland. This should not cause a problem for a researcher using this data as long as she is aware of this discrepancy. In any case, a simple back-of-the-envelope calculation shows that this discrepancy cannot be too large. We add all international trade in our data and find that it is 44 percent of all the international trade computed with country statistics from Eurostat.

The dataset contains the value of goods shipped among all region pairs for all industries and years. We refer to the region where a shipment starts as the seller and to the region where the shipment arrives as the buyer. We do not know the identities of the specific parties involved in the shipments. Some of them might entail moving goods between establishments of a given firm, while others might entail moving goods from establishments of one firm to those of a different one. We do not know either how the parties obtained the goods and what they do with them. Some firms shipping goods might be the original producers of these goods, while other might be intermediaries. Some firms receiving the goods might be the final consumers of the goods, while others might be intermediaries. Having this additional information would be useful to test alternative trade theories, but it is not crucial to provide an accurate description of how goods flow within and across European regions.

Since these flows vary little between 2011 and 2017, we use averages over the entire period and ignore the time dimension. Here we mostly focus on the aggregate bilateral trade matrix that also averages across industries. Whenever relevant, we discuss the most notable differences between the results obtained with the average matrix and the industry matrices.

Each of these bilateral trade matrices takes the following form:

$$\mathrm{X} = \begin{bmatrix} X_{11} & X_{12} & \cdots & X_{1N} \\ X_{21} & X_{22} & \cdots & X_{2N} \\ \vdots & \vdots & \ddots & \vdots \\ X_{N1} & X_{N2} & \cdots & X_{NN} \end{bmatrix} \tag{3.1}$$

where X_{nm} is the total value of shipments of goods from origin n to destination m. We measure shipments as a share of total shipments: $\sum_n \sum_m X_{nm} = 1$. Thus, X_{nm} is the probability that a shipment of goods has origin n and destination m.

Figure 3.2 shows a heat map of the matrix of bilateral trade. We refer to the entries in the main diagonal as *home* trade because they record trade within regions. Despite being a small set of entries (269 out of 72,361), each of them contains a lot of trade. Adding them, we find that home trade constitutes 40% of all European regional trade. We refer to the off-diagonal entries such that origin and destination regions are in the same country as *country* trade. Since regions within a country have been listed together, these entries can be identified in Figure 3.2 as the squares centered around the diagonal (without including the latter). Larger squares refer to countries with more regions, such as Germany or France. Smaller squares refer to countries with fewer regions such as Portugal or Ireland. Country trade entries tend to contain less trade than home trade entries. But there are many more country trade entries (4,958 out of 72,361) and, adding them, we find that country trade constitutes about 41% of all European regional trade. Finally, we refer to the remaining off-diagonal entries as *foreign* trade. We can identify these entries in Figure 3.2 as the off-diagonal entries outside the squares. Though most of the entries are foreign trade (67,134 out of 72,361), each of them contains little trade. This is why adding them we find that foreign trade constitutes only 19% of all European regional trade. There is therefore a strong bias towards home and country trade in our data.

The matrix in Figure 3.2 contains a fair amount of zeros. Not surprisingly, there are no zeros for home trade. But there are a few zeros for country trade: 157 out of 4,958 region pairs. And there are many more for zeros for foreign trade: 25,699 out of 67,134 region pairs. This distribution of zeros is also consistent with a strong home and country biases in European regional trade.

What explains these biases? A prime suspect is distance. The distance

Figure 3.2: Bilateral trade matrix for European regions

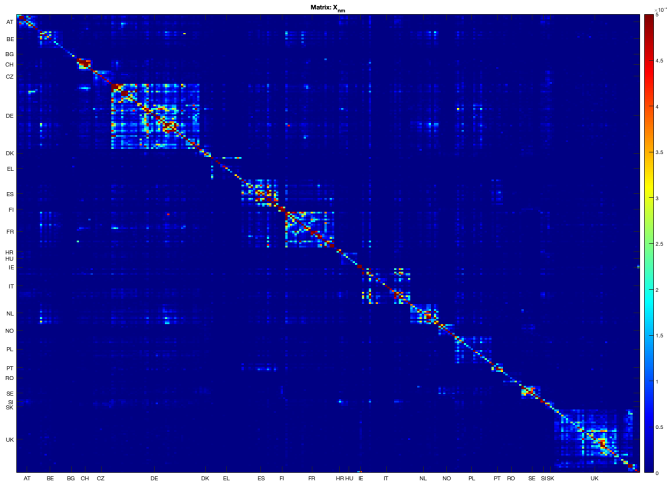

traveled by shipments classified as home, country and foreign trade is not the same. Fortunately, the European Road Freight Survey survey provides the actual distance traveled by each individual shipment, including shipments within and across regions. Figure 3.3 shows the histograms for distance traveled for home, country and foreign trade separately. The average distance traveled for the different types of trade is 21.2 Kms, 223.0 Kms and 631.9 Kms, respectively. There is little overlap, for instance, between the histograms for home and foreign trade.

3.2.2 Normalized market shares

Our goal is to understand the shape the matrix of bilateral trade. Which region pairs have strong trading relationships? Which ones have weak trading relationships? What are the factors that shape the trading relationship of a given region pair?

To answer these questions, we need a benchmark that is size free. To see this, consider the case of Catalonia and La Rioja, two regions in Spain. The

Figure 3.3: Home, country and foreign distances

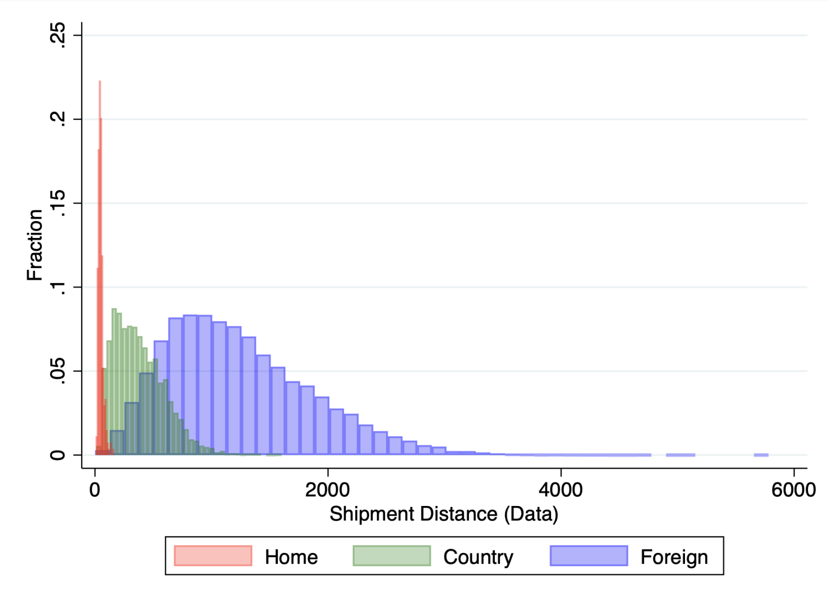

probabilities of a sale to the Basque Country, another region in Spain, for Catalonia and La Rioja are 0.000226 and 0.0000542, respectively. The probabilities of a purchase by Catalonia and La Rioja from the Basque country are 0.0004281 and 0.0000601, respectively. Catalonia's trade probabilities are one order of magnitude larger than those of La Rioja. Does this mean that Catalonia has a more intensive trade relationship with the Basque Country than La Rioja? This would be an absurd conclusion, we think, since Catalonia's population is 7.6 million while La Rioja's is 0.3 million. It is therefore almost inevitable that Catalonia trades more with the Basque Country than La Rioja. The size of origin and destination regions matters and we need to correct for this.

To determine how to correct for size, let us define two events: (i) $O_n =$ a shipment has origin n, and (ii) $D_m =$ a shipment has destination m. The probability of these two events are $X_n^O \equiv \sum_l X_{nl}$ and $X_m^D \equiv \sum_k X_{km}$, respectively. Let us now propose this independence benchmark: "the probability of a shipment from origin n to destination m should be $X_n^O X_m^D$." This benchmark essentially says that the events O_n and D_m are pairwise independent.

One can interpret this benchmark as a theoretical assertion or as a forecast with limited information. A theory asserting that all sellers have the same probability of trading with a given buyer and all buyers have the same probability of trading with a given seller implies that $X_{nm} = X_n^O X_m^D$. If we only know the sizes of regions n and m, the best forecast for their trade probability is $X_{nm} = X_n^O X_m^D$. In both interpretations, the independence benchmark captures the idea that bilateral trade is independent of how far the trading partners are in terms of geographical distance, political institutions, factor endowments, tastes, and so on. Thus, we can use deviations from this benchmark to learn about the role that these factors play in shaping trade relationships.

Figure 3.4 plots $\ln(X_{nm})$ against $\ln\left(X_n^O X_m^D\right)$. Not surprisingly, size shows its weight and pairs containing large regions trade more than pairs containing small regions. A simple regression of $\ln(X_{nm})$ on $\ln\left(X_n^O X_m^D\right)$ delivers an R-squared of 0.22 and a slope coefficient of 0.69. The result that size explains close to a quarter of the total variation in trade probabilities is not very interesting, though, since this relationship is somewhat mechanical. How could the trade probabilities involving a given region not be related to the region's size, which is defined as the sum of the trade probabilities of the region?

What is really interesting about Figure 3.4 is that more than three quarters of the variation in trade probabilities cannot be explained by size. This is the variation we care about. Home trade observations are located well above the 45 degree line, confirming that regions trade with themselves much more than what their sizes suggest. The same applies to country trade observations, although to a lesser extent. The counterpart is that most foreign trade observations are below the 45 degree line. European regions have intense trading relationships with themselves and with other regions within their country, and mild trade relationships with regions in other countries.

To make this idea precise, we measure the intensity of the trade relation-

Figure 3.4: Actual vs predictd trade (log) probabilities

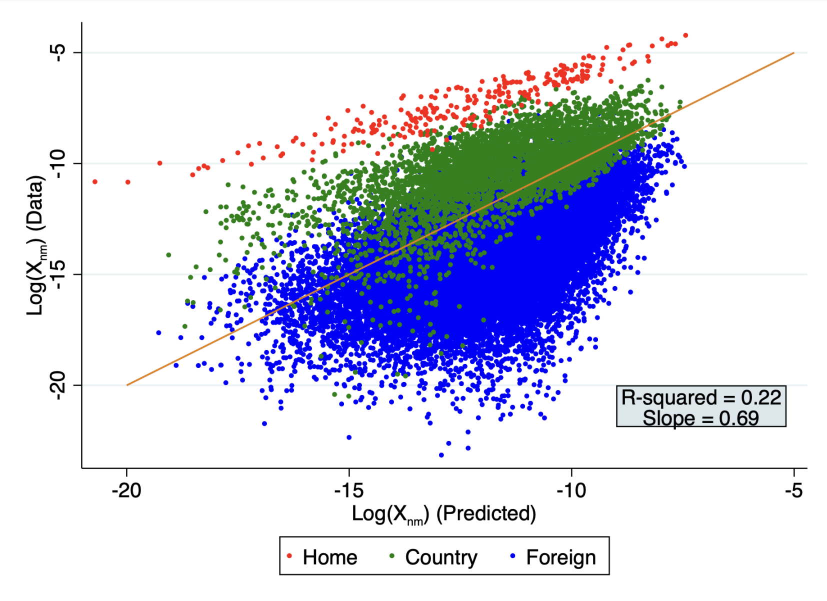

ship for a region pair with the ratio of the actual trade probability and the trade probability predicted by the independence benchmark:

$$S_{nm} = \frac{X_{nm}}{X_n^O X_m^D} \tag{3.2}$$

We refer to this measure as a normalized market share.[2] This measure corrects for the mechanical effect of size on trade and it has a very simple interpretation: if $S_{nm} = 2$ (0.5), shipments from origin n to destination m are twice (half) as large as one would be able to predict knowing only the sizes of the regions. Thus, S_{nm} is a size-free measure of how strong a trade

[2]The reason is that S_{nm} has two alternative interpretations that suggest this name. First, S_{nm} is the share of origin n in destination market m, i.e., X_{nm}/X_m^D; normalized by the share of origin n in the entire European market, i.e., X_n^O. Second, S_{nm} is the share of destination m in origin market n, i.e., X_{nm}/X_n^O; normalized by the share of destination m in the entire European market, i.e., X_m^D. The World Bank uses a related measure for country trade named Trade Intensity Index (https://wits.worldbank.org/wits/wits/witshelp/Content/Utilities/e1.trade_indicators.htm.). This index normalizes probabilities by international trade instead of total trade, i.e., it does not include home trade.

relationship is.[3]

Figure 3.5 plots histograms of (log) normalized market shares for home, country and foreign trade. The average values for the different types of trade are 469.5, 11.22 and 0.44, respectively. The distributions of normalized market shares for these types of trade have little overlap. The ranking home > country > foreign trade is not surprising. But the magnitude of the differences is (at least to us!). More so, since we are using data on trade in goods and not trade in services.

Figure 3.5: Home, country and foreign normalized market shares

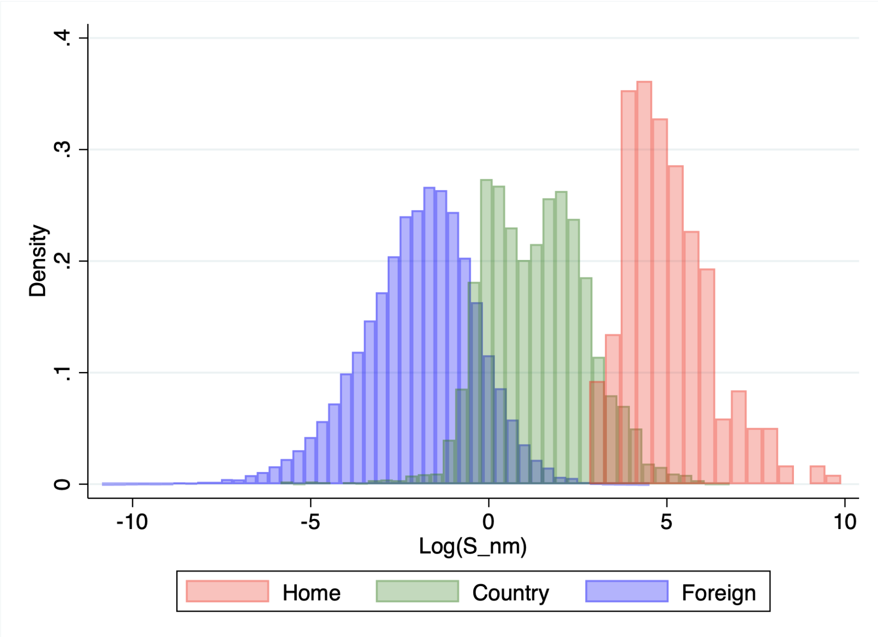

Finally, and just to whet the appetite for what is coming next, Figure 3.6 plots the (log) normalized market shares against the (log) of actual distances. It is apparent that the strength of trade relationships declines with distance.

[3]If we go back to the example of Catalonia and La Rioja, we find that normalized market shares for Catalonia are 2.83 (sales/exports) and 3.91 (purchases/imports) and for La Rioja 16.17 and 15.19. Catalonia and the Basque Country trade between three and four times more than one would predict given their sizes, but La Rioja and the Basque Country trade between fifteen and sixteen times more! Thus, it is La Rioja that has a stronger trade relationship with the Basque Country. One reason for this is that La Rioja is much closer geographically to the Basque Country than Catalonia is.

This surely helps explain part of the home and country biases in trade. But Figure 3.6 also shows that distance cannot be the single explanation for these biases. Within any given distance interval, we can observe the ranking of home > country > foreign trade. What else is going on? We turn next to a systematic examination of the data using the standard gravity framework.

Figure 3.6: (Log) normalized market shares and (log) distance

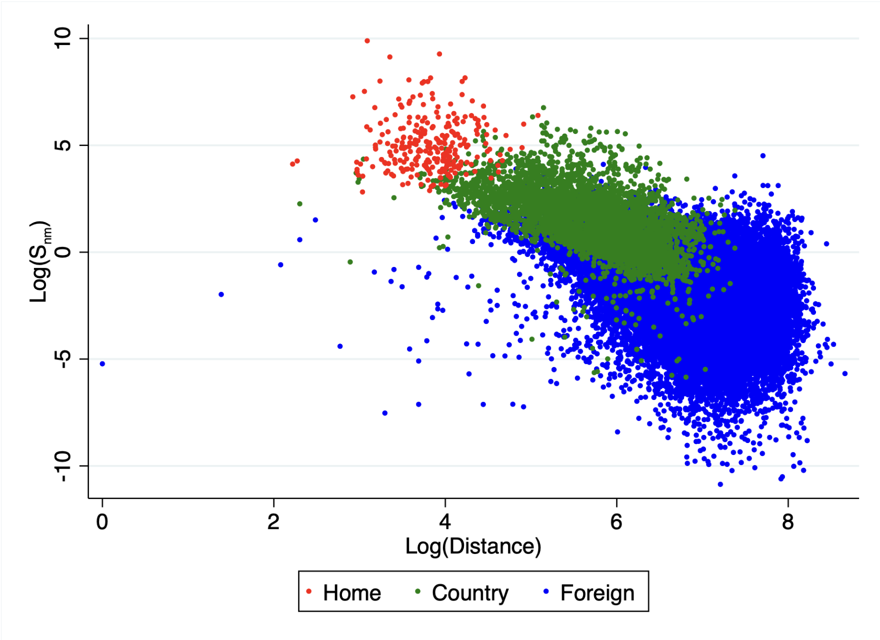

3.3 A gravity look at the data

Figure 3.7 shows the matrix of (log) normalized market shares. The goal of this section is to provide a parsimonious description of this matrix. To do this, we use the gravity framework to guide our search for patterns. The bottom line is simple again: using distance and borders we can explain about two thirds of the variation in (log) normalized market shares. To reach this conclusion, we explore a battery of increasingly flexible specifications for distance and border effects.

Figure 3.7: Bilateral matrix of (log) normalized market shares for European regions

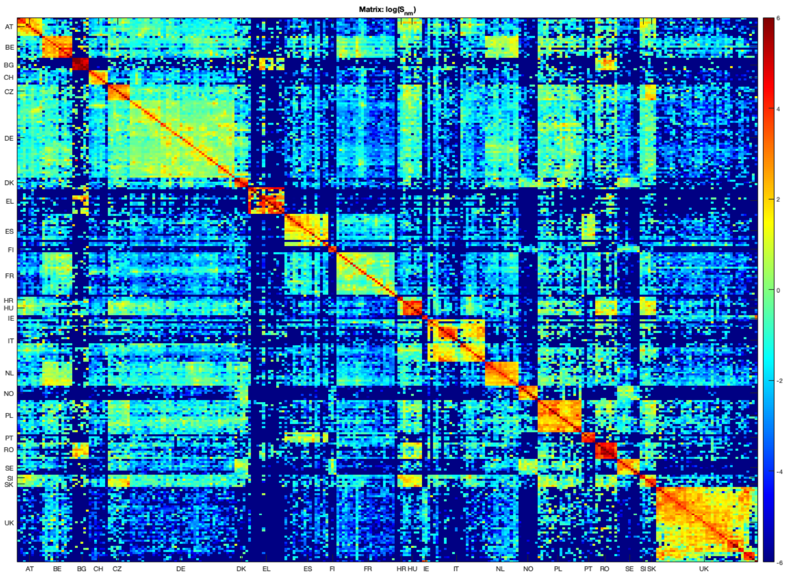

3.3.1 The gravity framework

The gravity framework provides a specific mathematical structure that adjusts trade probabilities to take into account distance, borders and other variables. Let M_{nm} be a measure of the cost of shipping goods from origin n to destination m. We refer to M_{nm} as bilateral market access. Gravity models postulate a bilateral market access function of this form:

$$M_{nm} = \exp\left\{\sum_i \theta^i Z^i_{nm}\right\} \tag{3.3}$$

where $\{Z^i_{nm}\}$ is a set of bilateral variables that jointly determine market access and $\{\theta^i\}$ is a set of theoretical coefficients. The set of bilateral variables typically contains a distance variable and a border dummy measuring whether the regions are in the same country or not. In many cases, other variables that might affect the costs of shipping goods are added such as dummies measuring whether the regions have a common language or currency.

The gravity framework consists of the following mathematical model:

$$X_{nm} = \frac{M_{nm}}{M_n^O M_m^D} X_n^O X_m^D \tag{3.4}$$

which, alternatively, can be expressed in terms of normalized market shares as follows:

$$S_{nm} = \frac{M_{nm}}{M_n^O M_m^D} \tag{3.5}$$

where M_n^O and M_m^D is a set of numbers that satisfy the following restrictions:

$$1 = \sum_m X_m^D \frac{M_{nm}}{M_n^O M_m^D} \tag{3.6}$$

$$1 = \sum_n X_n^O \frac{M_{nm}}{M_n^O M_m^D} \tag{3.7}$$

We refer to M_n^O and M_m^D as origin and destination measures of average market access.[4] Equations (3.6) and (3.7) are not additional theoretical restrictions, but instead consistency requirements that ensure that probabilities add, i.e., $1 = \sum_m X_m^D S_{nm}$ and $1 = \sum_n X_n^O S_{nm}$.

It is well known that there is a large set of theoretical models that are consistent with the formulation of the gravity framework in Equations (3.5), (3.6) and (3.7) (See Head and Mayer (2014)). These models predict that the trade relationship of a region pair is strong if its bilateral market access is large relative to the average market access of origin and destination regions.

[4]The literature often refers to these terms as multilateral resistance terms or price levels, but labeling them as origin and destination measures of market access seems more transparent.

3.3.2 An important example

We explore next a parsimonious version of the gravity model that offers a number of interesting insights and, as we shall show soon enough, it explains a substantial fraction of the variation in the matrix of (log) normalized market shares. In particular, let us assume the following bilateral market access function:

$$M_{nm} = \exp\{\sigma D_{nm} + \beta B_{nm}\} \tag{3.8}$$

where $\sigma, \beta \leq 0$. The variable $D_{nm} \geq 0$ is the (log) average kilometers travelled between regions n and m. The variable B_{nm} is a dummy variable that takes value 0 if regions n and m belong to the same country, and takes value 1 otherwise. The coefficients σ and β measure the (negative) effect of distance and borders on bilateral market access, respectively.

Figure 3.8 shows three theoretical matrices of (log) normalized market shares produced with this model. In all of them, we set $\sigma = 0$ so that:

$$M_{nm} = \begin{cases} 1 & \text{if } B_{nm} = 0 \\ e^{\beta} & \text{if } B_{nm} = 1 \end{cases} \tag{3.9}$$

From left to right, these matrices assume that $\beta = 0$, $\beta = -1.2$ and $\beta = -2.4$, respectively. Thus, we start from the independence benchmark with all (log) normalized market shares equal to zero on the left, and then increase the border effect in two steps as we move right. As the border effect becomes stronger, bilateral market access for region pairs in different countries shrinks. As a result, average market access for all origin and destination regions also shrinks. Crucially, this shrinkage is larger for regions within small countries than for regions within large ones.[5] The reason, of course, is that the costs of trade have increased more for the former than for the latter.

These observations lead to two important theoretical predictions. The first one is that, as the border effect becomes stronger, country/home trade

[5] By the size of the country, we mean the sum of the sizes of its regions.

Figure 3.8: Borders and trade

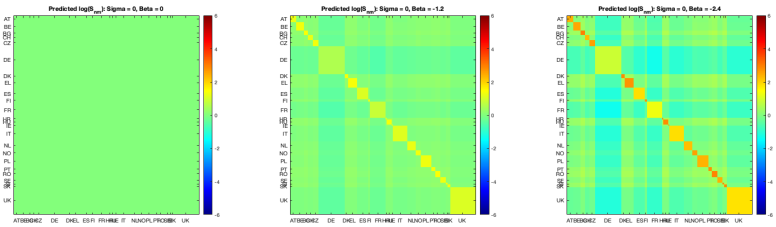

grows and foreign trade shrinks. This generates squares centered along the diagonal with high-trade entries inside them and low-trade entries outside. The second theoretical prediction is that, as the border effect becomes stronger, regions in small countries experiment more growth of country/home trade and less shrinkage of foreign trade. This small-country effect (which is due exclusively to the differential change in average market access) creates a specific source of heterogeneity and it has a very simple intuition. If you have above-average trade relationships with many/large regions (i.e. large country), not only each of these relationships cannot be too much above average but also the remaining relationships must be well below average. If you have above-average trade relationships with few/small regions (i.e. small country), these relationships can be well above average and yet the remaining relationships do not have to be much below average.

Figure 3.9 plots actual (log) normalized market shares against country size, using different colors for home, country and foreign trade. Not surprisingly, we see again that home/country trade is larger than foreign trade, which is consistent with the first theoretical prediction. More interesting is that regions in small countries have larger (log) normalized market shares than regions in large countries. This can be seen when we compare (log) normalized market shares within each type of trade. Clearly, the small-country effect is present in the European regional trade data.

Figure 3.10 shows three additional theoretical matrices of (log) normalized

Figure 3.9: (Log) normalized market shares and country size

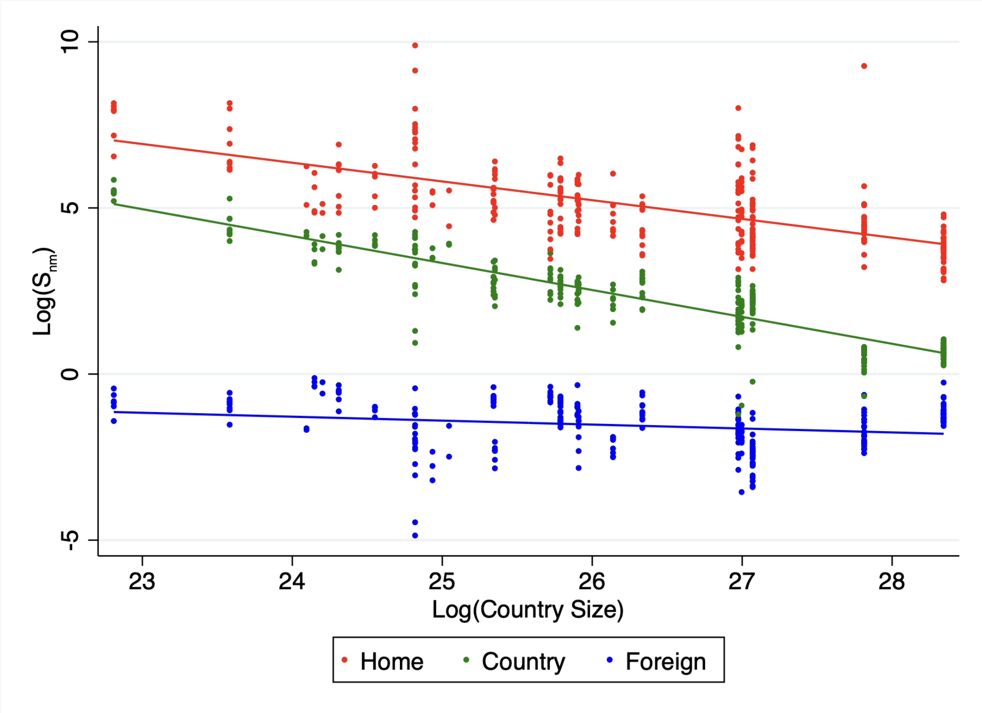

market shares produced with the model. In all of them, we set $\beta = 0$ so that:

$$M_{nm} = e^{\sigma D_{nm}} \tag{3.10}$$

From left to right, these matrices assume that $\sigma = 0$, $\sigma = -0.6$ and $\sigma = -1.2$, respectively. Thus, we start with the independence benchmark again, and then increase the cost of distance in two steps as we move right. As the distance effect becomes stronger, bilateral market access for all region pairs shrink. This shrinkage is larger for region pairs that are far away from each other. As bilateral market access shrinks, average market access for all origin and destination regions also shrink. Now, this shrinkage is larger for regions that are remote within Europe than for regions that are central. The reason, again, is that the costs of trade have increased more for the former than for the latter.

These observations lead to two theoretical predictions. The first one

Figure 3.10: Distance and trade

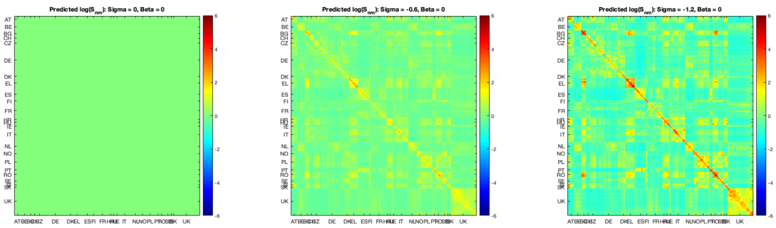

is again that, as the distance effect becomes stronger, country/home trade grows and foreign trade shrinks. The reason is that regions in different countries are far away from regions in the same country (recall Figure 3.3). This generates again squares centered along the diagonal with high-trade entries inside them and low-trade entries outside. An interesting novelty is that now trade is not homogeneous inside these squares. In particular, there is more trade in the diagonal than in the rest of these squares since regions are closer to themselves than to other regions within the same country. The second theoretical prediction prediction is that remote regions experiment more growth in country/home trade and more shrinkage of foreign trade. This remoteness effect creates a second specific source of heterogeneity, which is also quite intuitive.

Figure 3.11 plots actual (log) normalized market shares against an index of remoteness.[6] A quick look at the figure shows that (log) normalized market shares for home and country trade do indeed grow with remoteness, while (log) normalized market shares for foreign trade shrink. The remoteness effect is also present in European regional trade data.

Armed with these intuitions, we search next for the combination of σ and β that provides the best fit of this model to the data. To do this, we define a two-dimensional grid over σ and β. For each point in the grid, we compute: (i) a complete set of bilateral market access measures $\{M_{nm}\}$; (ii) a complete set of origin/destination average market access measures $\left\{M_n^O\right\}$

[6]This index is the average distance to all other regions in Europe.

Figure 3.11: (Log) normalized market shares and remoteness

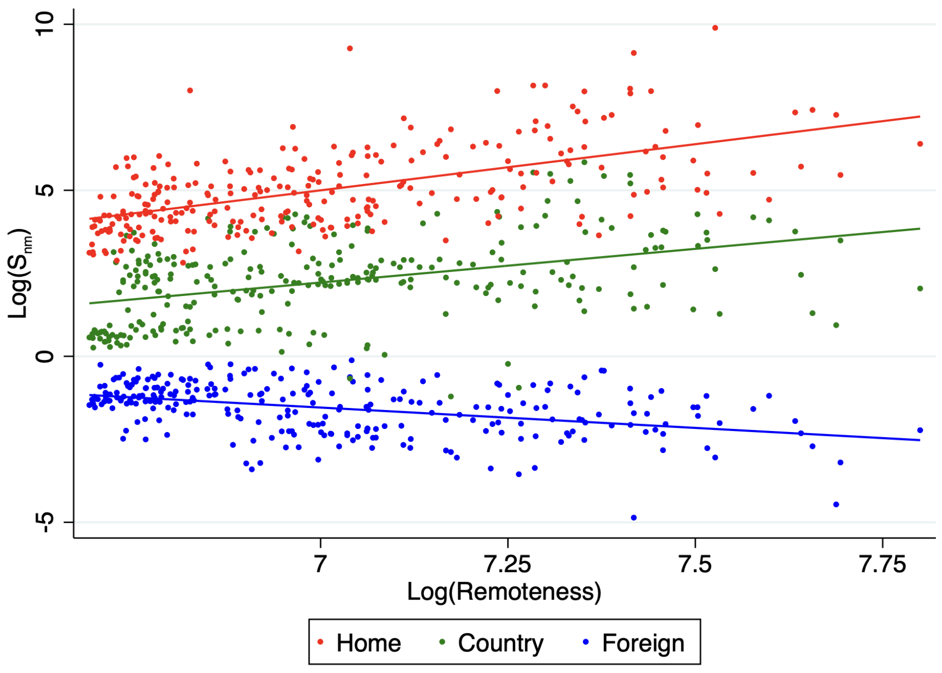

and $\left\{M_m^D\right\}$; and (iii) the matrix of predicted (log) normalized market shares. We then choose the values of σ and β that minimize the distance between the matrices of actual and predicted (log) normalized market shares.[7] This procedure leads us to choose $\sigma = -1.3$ and $\beta = -2.4$. Figure 3.12 shows how sensitive is the fit of the model to changes in parameter values.

Figure 3.13 plots the actual matrix of (log) normalized market shares in the left panel and the matrix of predicted (log) normalized market shares in the right panel. Even though there are differences across the two matrices, it seems that the parsimonious model discussed here captures some of the most important patterns in the data. To reinforce this message, Figure 3.14 plots the entries of these matrices against each other.

[7]To minimize the distance we use as a criterion the Frobenious norm.

Figure 3.12: Sensitivity analysis

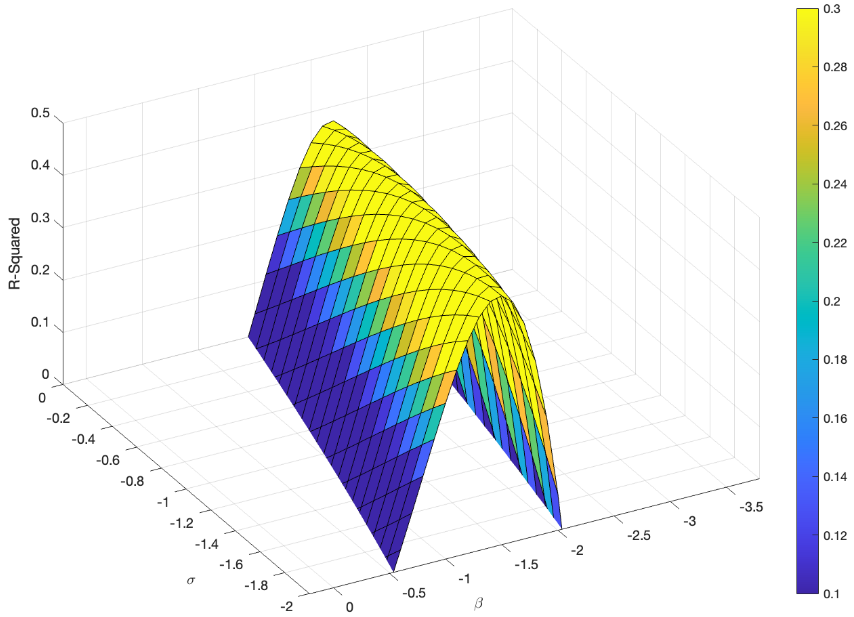

Figure 3.13: Actual vs predicted matrices of (log) normalized market shares

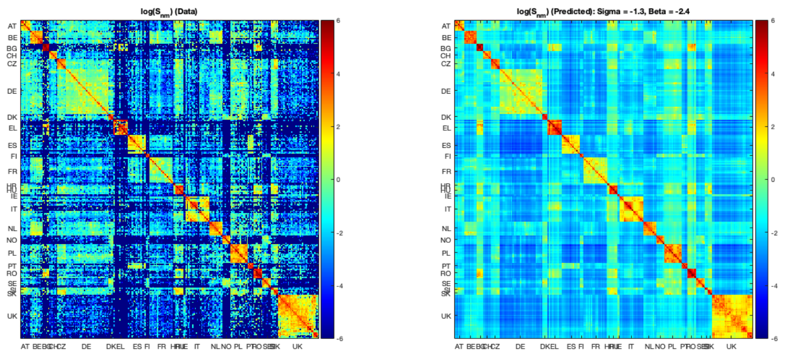

3.3.3 Fixed-effects regressions

Next, we estimate the following fixed-effects regression:

$$\ln S_{nm} = \phi_n^O + \phi_m^D + \sum_i \theta^i Z_{nm} + u_{nm} \tag{3.11}$$

Figure 3.14: Actual vs. predicted (log) normalized market shares

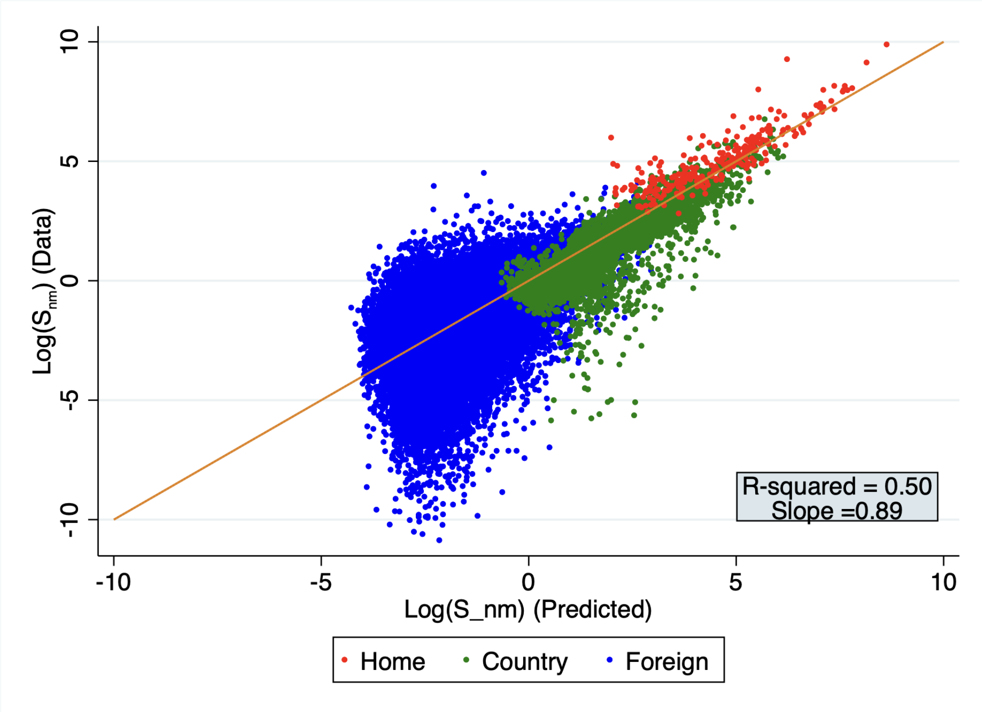

where ϕ_n^O and ϕ_m^D are region fixed effects and u_{nm} is an error term that is assumed to be orthogonal to the regressors. The idea behind this regression is to allow the data to choose the parameters $\{\theta^i\}$ that give the model the best chance to explain the data. The estimates of the fixed effects are then interpreted as our estimates of $\ln M_n^O$ and $\ln M_m^D$.[8]

Table 1 shows the results of estimating regression (3.11) for six different gravity models. Column (1) shows the parsimonious model that we used in the previous subsection. In particular, there is a border dummy B_{nm} and a measure of distance D_{nm} which is the (log) average kilometers travelled from n to m. This specification therefore assumes a constant elasticity of trade to distance.

Column (1) shows that the parsimonious model explains almost two-

[8]Recovering origin and destination market access measures from a fixed-effects regression is much more difficult when the dependent variable is $\ln X_{nm}$. See Fally (2015) for a discussion of this problem.

Table 3.1: Gravity: Fixed Effects Regressions

	(1) Log(S_nm)	(2) Log(S_nm)	(3) Log(S_nm)	(4) Log(S_nm)	(5) Log(S_nm)	(6) Log(S_nm)
Border dummy	-2.384***	-2.340***				
	(0.260)	(0.243)				
Border / common language / common currency dummy			-1.530***	-1.491***		
			(0.189)	(0.185)		
Border / common language / different currency dummy			-1.799***	-1.742***		
			(0.228)	(0.221)		
Border / different language / common currency dummy			-2.267***	-2.242***		
			(0.183)	(0.171)		
Border / different language / different currency dummy			-2.777***	-2.744***		
			(0.221)	(0.208)		
Border dummies for each country pair	No	No	No	No	Yes	Yes
Distance (constant-elasticity)	-1.190***		-1.071***		-1.006***	
	(0.0668)		(0.0607)		(0.0712)	
Distance (variable-elasticity)	No	Yes	No	Yes	No	Yes
Origin FE	Yes	Yes	Yes	Yes	Yes	Yes
Dest FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	46505	46505	46505	46505	46505	46505
R^2	0.610	0.611	0.623	0.624	0.666	0.668

Standard errors in parentheses. Distance is measured as the (log) of average kilometers travelled.
* $p < .1$, ** $p < .05$, *** $p < .01$

thirds of the variation in trade probabilities. This is especially remarkable given that we have eliminated the effects of size using (log) trade normalized market shares instead of (log) trade probabilities.[9] Border and distance effects are significative, economically large and not far away from those that we found in the calibration exercise above. The estimated coefficient for the border dummy means that, controlling for distance, a national border reduces bilateral trade to $\exp\{-2.384\} \times 100 = 9.21$ percent of the independence benchmark. The estimated coefficient for distance implies that, controlling for borders, a one percent increase in distance traveled reduces bilateral trade by 1.19 percent with respect to the independence benchmark. Clearly, borders and distances can predict deviations from the independence benchmark.

In Column (2) we use a more general distance function that allows for the

[9]We have estimated all the regressions in Table 1 using $\ln X_{nm}$ as the dependent variable instead of $\ln S_{nm}$. All the coefficients of bilateral variables remain unchanged up to the third decimal. Since the size correction is picked up by the fixed effects, now to be interpreted as $\ln\left(\frac{X_n^O}{M_n^O}\right)$ and $\ln\left(\frac{X_m^D}{M_m^D}\right)$, the R-squared of the regressions is a bit inflated. Going from Column (1) to (6) the R-squared starts at 0.681 and grows up to 0.729.

elasticity of trade to distance to vary across distance brackets. The results are very similar. The R-squared and the border coefficient are essentially the same. Figure 3.15 plots the effect of distance on trade using the estimates of the regressions in columns (1) and (2). The constant-elasticity specification is always above the variable-elasticity one, indicating that the former might be overestimating the effects of distance on trade. But the difference does not seem to be large.

Figure 3.15: Constant vs. variable elasticity distance functions

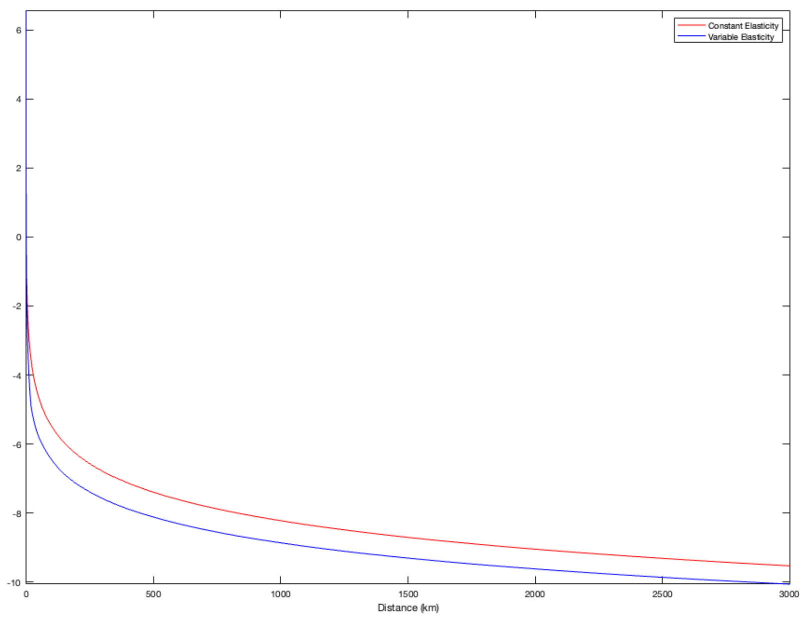

Columns (3) and (4) allow for some heterogeneity in the border effect. In particular, the border effect is allowed to depend on whether the regions involved have a common language and currency. The idea is that sharing a language and/or a currency facilitates trade and reduces the border effect. Using this flexible specification of the border effect raises the R-squared of the regression only marginally. Interestingly, we see that the distance effect is a bit smaller now since the estimated elasticity of trade to distance is -1.071. Again, there is not much difference between the constant- and variable-elasticity specifications for the distance effect.

The results in Columns (3) and (4) indicate that indeed the border effect depends on whether the region pair shares a language and/or a currency. At one extreme, a national border separating a region pair that shares both language and currency reduces bilateral trade to $\exp\{-1.491\} \times 100 = 22.52$ percent of the independence benchmark. At the other extreme, a national border separating a region pair that shares neither language nor currency reduces bilateral trade to $\exp\{-2.744\} \times 100 = 6.43$ percent of the independence benchmark. The estimated coefficients suggest that not sharing a language is more deleterious to trade than not sharing a currency, even though both variables seem to matter.

Columns (5) and (6) estimate different border effect for each country pair. That is, we allow the French-Spanish border to have different effects than the Finish-Spanish or the Irish-British borders. Since there are 24 countries in our sample, we are estimating 276 different border effects. This is the most flexible specification of the border effect so far. Yet, we find that the R-squared of the regression increases only marginally. The distance effect is reduced even further as the estimated elasticity of trade to distance is now -1.006. We confirm again that using the constant- or the variable-elasticity specifications of distance does not make much of a difference. The estimates of border effects for each country pair show substantial heterogeneity. Figure 3.16 and Table 2 show this.

We also estimate the regressions in Table 1 using a Poisson Pseudo-Maximum Likelihood (PPML) estimator (See Silva and Tenreyro (2006)). The results are shown in the Table A.3.2 in the Appendix. The estimates are quite similar to those obtained with OLS and reported in Table 2. The main difference is that, in Columns (3)-(4) not sharing a currency now is more important than not sharing a language.

Tables 3.2 and 3.3 show the results for our baseline fixed-effects regressions for each industry individually. Our estimation shows some heterogeneity

Figure 3.16: Histogram of country pair dummies

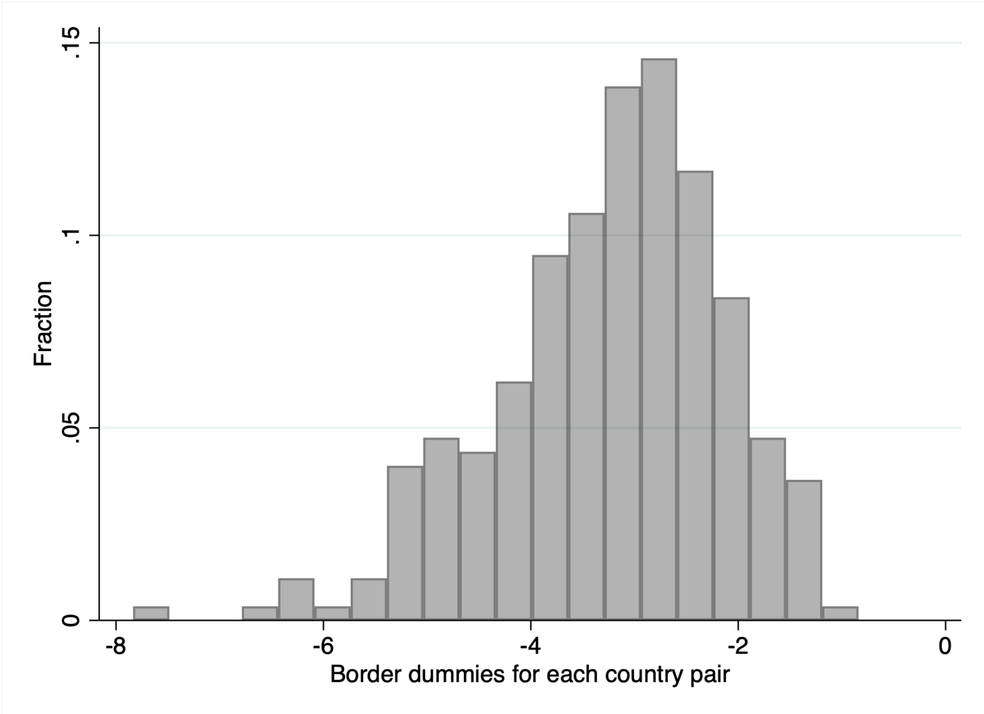

across industries. The first observation is that this model retains a high explanatory power for all industries, with the R-squared ranging between 0.554 and 0.798. The second observation is that the border effect is also substantial for all industries. It ranges from -0.728 (Coke and Petroleum) to -2.426 (Food, Beverage and Tobacco). For most industries (8 out of 12) it is between -1.4 and -1.8, slightly smaller than the average coefficient we obtained in Table 1. The third and final observation is that the distance coefficient varies substantially across industries, ranging from -0.494 to -1.884. For most industries this coefficient is close to -1, which is close to the average coefficient that we obtained in Table 1.

Table 3.2: Gravity: Fixed Effects Regressions

	(1) Agri	(2) Mining	(3) FBT	(4) Textiles	(5) Wood	(6) Coke Pet
Border Effect	-1.698***	-1.191***	-2.426***	-0.991***	-1.656***	-0.728***
	(0.191)	(0.209)	(0.142)	(0.147)	(0.104)	(0.182)
Distance (constant-elasticity)	-1.174***	-1.884***	-1.006***	-0.494***	-1.065***	-1.458***
	(0.0932)	(0.193)	(0.0650)	(0.0938)	(0.0480)	(0.169)
Origin FE	Yes	Yes	Yes	Yes	Yes	Yes
Dest FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	20226	10072	27764	11428	21348	6870
R^2	0.672	0.798	0.699	0.554	0.660	0.718

Standard errors in parentheses. Distance is measured as the (log) of average kilometers travelled.
* $p < .1$, ** $p < .05$, *** $p < .01$

Table 3.3: Gravity: Fixed Effects Regressions

	(1) Chem	(2) Non-Metal	(3) Metal	(4) Machinery	(5) Vehicles	(6) Other
Border Effect	-1.619***	-1.860***	-1.610***	-1.810***	-1.674***	-1.422***
	(0.144)	(0.131)	(0.123)	(0.146)	(0.151)	(0.147)
Distance (constant-elasticity)	-1.005***	-1.388***	-0.914***	-0.640***	-0.570***	-0.603***
	(0.0581)	(0.0998)	(0.0603)	(0.0820)	(0.0753)	(0.0678)
Origin FE	Yes	Yes	Yes	Yes	Yes	Yes
Dest FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	24073	16764	22527	22368	20014	16100
R^2	0.633	0.766	0.623	0.586	0.566	0.565

Standard errors in parentheses. Distance is measured as the (log) of average kilometers travelled.
* $p < .1$, ** $p < .05$, *** $p < .01$

3.4 The home bias in trade

In our previous exploration of the data, we have treated all trade flows within the same country in the same way. However, we have shown in Section 2 that home trade is orders of magnitude larger than country trade, accounting for 40% of intra-European flows in our data. We now explore how large is this difference by adding a home-bias dummy to our gravity estimation. Table 3.4 shows the same fixed-effects regressions that we saw in Table 3.1, including this additional variable. There are three key takeaways. First, the coefficient on the home-bias dummy is large and significant. Across all columns this coefficient is positive and comparable in size to the border effect. Focusing on our extended model in columns (3)-(4), the average market share of a region with itself ranges from $\exp\{1.013\} \times 100 = 275$ and $\exp\{2.233\} \times 100 = 932$ percent larger than the average market share between two different regions in the same country, controlling for distance. Second, the R-squared of the regressions does not change much after we introduce the home-bias dummy. This reflects the fact that home trade has a very small number of observations in the overall trade matrix. Failing to fit those is not severely penalized in the tests we performed above. Third, our estimates of the border effect and distance does not change much as a result of adding the home-bias dummy.[10]

What explains this strong home bias in trade? To make progress in answering this question, we perform two exercises. First, we explore which regional characteristics are correlated with a large home bias. Second, we separate statistical and political regions and show that it is the latter and not the former that exhibit a large home bias in trade.

Table 3.4: Gravity: Fixed Effects Regressions

	(1) Log(S_nm)	(2) Log(S_nm)	(3) Log(S_nm)	(4) Log(S_nm)	(5) Log(S_nm)	(6) Log(S_nm)
Border dummy	-2.380***	-2.321***				
	(0.261)	(0.241)				
Border / common language / common currency dummy			-1.499***	-1.466***		
			(0.182)	(0.179)		
Border / common language / different currency dummy			-1.763***	-1.726***		
			(0.228)	(0.218)		
Border / different language / common currency dummy			-2.265***	-2.217***		
			(0.176)	(0.165)		
Border / different language / different currency dummy			-2.782***	-2.729***		
			(0.222)	(0.208)		
Border dummies for each country pair	No	No	No	No	Yes	Yes
Home Bias	1.013***	2.079***	1.271***	2.166***	1.424***	2.233***
	(0.259)	(0.409)	(0.218)	(0.352)	(0.184)	(0.289)
Distance (constant-elasticity)	-1.150***		-1.016***		-0.903***	
	(0.0689)		(0.0604)		(0.0670)	
Distance (variable-elasticity)	No	Yes	No	Yes	No	Yes
Origin FE	Yes	Yes	Yes	Yes	Yes	Yes
Dest FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	46505	46505	46505	46505	46505	46505
R^2	0.611	0.613	0.625	0.627	0.669	0.671

Standard errors in parentheses. Distance is measured as the (log) of average kilometers travelled.
* $p < .1$, ** $p < .05$, *** $p < .01$

[10] As a robustness check, we report the results using PPML in Tables A.3.2 and A.3.3 in the Appendix.

Figure 3.17: Normalised Market share: Home

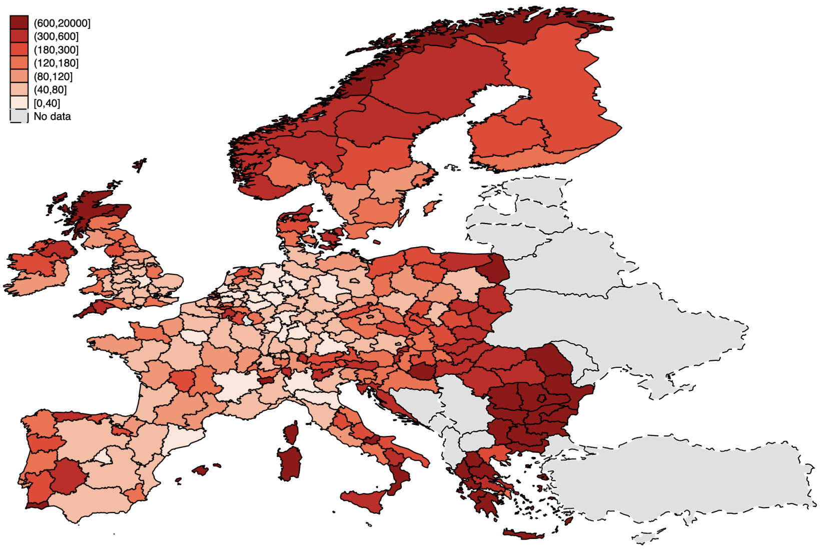

3.4.1 Correlates of the home bias in trade

Figure 3.17 shows the spatial distribution of the home market share in Europe. The first striking pattern is how heterogeneous these shares are across regions. They range from a low of 40 to a high of about 20,000. The map also shows that geography plays an important role. Regions in the periphery of Europe, like Greek and Bulgarian regions in the South and Norwegian and Swedish regions in the North tend to have higher home trade. Island and mountainous regions also have higher home trade. Interestingly, within-country geography also plays a role: regions in the periphery of a country display higher home trade than more central regions. For instance, regions in the south of Italy and Portugal, in the west of Spain and in the north of the UK and Denmark have higher home trade than the rest of the country.

Interestingly, we see that home trade tends be lower in more densely populated regions of Europe. We see this pattern at the European level in the so-called Blue Banana.[11] We also see this pattern within some countries

[11]The Blue Banana is a corridor of highly urbanized land spreading over Western and Central Europe. It stretches approximately from North West England through the English Midlands across Greater London to the European Metropolis of Lille, the Benelux

that are outside the Blue Banana. For instance, Madrid and Catalonia have the lowest home trade in Spain, while Warsaw and Athens have the lowest home trade in Poland and Greece.

Table 3.5 shows regressions of home trade on a number of regional characteristics, and country fixed effects. Column (1) reports the results using the following geographical variables: distance, remoteness plus island and mountain region dummies. All these variables are significant, except for distance. This formally confirms that remote regions, island regions and mountainous regions have higher home trade. These simple geographical variables explain 41 percent of the variation in the home market share.

Column (2) adds economic variables: presence of ports, motorway density, population, share of employment in manufacturing and in the public sector, the share of population with at least secondary education and the share of foreign-born population. The introduction of economic variables reduces the coefficients of the geographic variables. All economic variables are significant except for presence of ports. Motorway density reduces the home market share, showing that infrastructure helps overcome geographical obstacles. As we observed in the map, the most populated regions also have lower home market shares. Economic structure also matters, regions with high manufacturing shares, larger governments, more educated populations and more migrants have lower home market shares. Adding all these economic variables raises the R-squared from 41 to 80 percent.

Column (3) adds country fixed effects. The R-squared increases to almost 90 percent, indicating that some of the variation in home trade has a country component. Some variables seem to be correlated with this country component since they now lose their significance and the magnitude of their coefficients is reduced: the share of employment in the public sector, the share of population with at least secondary eduction and the share of

states with the Dutch Randstad and Brussels and along the German Rhineland, Southern Germany, Alsace-Moselle in France in the west and Switzerland (Basel and Zürich) to Northern Italy (Milan and Turin) in the south.

Table 3.5: Home Bias: Determinants

	(1) Home	(2) Home	(3) Home
Distance	-0.0171	0.229**	-0.0266
	(0.145)	(0.0904)	(0.187)
Log(European Remoteness)	2.345***	1.353***	1.551***
	(0.265)	(0.194)	(0.466)
Island Region	1.872***	0.915**	0.988***
	(0.509)	(0.364)	(0.328)
Mountain Region	0.304**	0.154**	0.193**
	(0.118)	(0.0722)	(0.0831)
Major Port Region		-0.197	-0.127
		(0.129)	(0.107)
Motorway Density		-6.379***	-6.510***
		(1.179)	(1.454)
Log(Population)		-0.819***	-0.758***
		(0.0488)	(0.0590)
Share of Emp. (Manuf.)		-10.48***	-10.01***
		(1.174)	(1.905)
Share of Emp. (Public)		-16.84***	-0.410
		(1.634)	(3.917)
Sh. Secondary or tertiary educ		1.511***	-1.399
		(0.398)	(0.903)
Share Migrant Pop.		-2.287***	-0.386
		(0.500)	(0.702)
Country FE	No	No	Yes
Observations	269	265	265
R^2	0.410	0.799	0.890

Standard errors in parentheses. Distance is measured as the (log) of average kilometers travelled.
* $p < .1$, ** $p < .05$, *** $p < .01$

foreign-born population. However most of our variables remain significant and their coefficients are stable. This means that the country component does not explain all the variation in home market shares. To confirm this, we estimate a regression that includes only country fixed effects and find that explains 56 percent which is well below the 90 percent obtained in Column (3).

Finally, we explore industry heterogeneity in home bias correlates. We estimate the regression in Column (3) for each of our 12 industries and report the results in Tables A.3.4 and A.3.5 in the Appendix. For most of the industries the results align with the average findings reported here. The exceptions are Agriculture, Mining and Coke/Petroleum, for which remoteness does not play a role.

3.4.2 Government structure and home trade

To learn more about the source of this home bias, we exploit a peculiarity of the data collection and harmonization process of our dataset. Since our shipment data is collected and provided by Eurostat, our units of observation are NUTS2 regions. In some countries these NUTS2 regions are only statistical regions created for the purpose of sharing data with Eurostat. In other countries, however, they coincide with political divisions with different levels of self-government. This provides us a unique opportunity to see the extent to which regional governments are behind this home bias in trade. In particular, we want to compare region pairs separated by statistical and political borders.

We work with two geographical classifications that partition our set of 24 countries into regions. The finer one is the NUTS2 classification that we have been using up to this point which includes 269 regions. The coarser one is the NUTS1 classification that includes 101 regions. We group countries in the following way:

Group 1: Countries with no political borders. These countries have regional governments neither at the NUTS1 nor at the NUTS2 level. Therefore all internal borders in these countries are statistical. These countries include Bulgaria, Estonia, Latvia, Lithuania, Portugal and Slovenia.

Group 2: Countries with political borders at a lower level of aggregation than NUTS2. These countries have regional governments, but every NUTS2 region contains more than one. One example is Switzerland that is divided into 26 Cantons, that have their own government. But Eurostat collects data by aggregating these Cantons into 7 NUTS2 regions. These countries include

Croatia, Czech Republic, Finland, Hungary, Ireland, Norway, Romania, Slovakia, Sweden and Switzerland.

Group 3: Countries with political borders that coincide with NUTS2 regions. These countries are Austria, Denmark, France, Greece, Italy, Netherlands, Poland and Spain.

Group 4: Countries with political borders that coincide with data at a higher level of aggregation than NUTS2. These countries include Belgium, Germany and the United Kingdom. In the case of Belgium and Germany, both NUTS2 and NUTS1 regions correspond to political borders (provinces and regions in Belgium, government regions and L*ä*nders in Germany). In the case of the United Kingdom, political regions either coincide with the NUTS1 classification (Northern Ireland, Scotland and Wales) or contain several NUTS1 regions (England). All NUTS2 regions are statistical.

We exploit this heterogeneity in statistical and political borders. To this end, we define two dummies $HB1_{nm}$ and $HB2_{nm}$. $HB1_{nm}$ takes value 1 if n and m are in the same NUTS1 region and 0 zero otherwise. $HB2_{nm}$ takes value 1 if n and m are in the same NUTS2 region and 0 zero otherwise. Then we estimate the following regression for each country:

$$lnS_{nm} = \phi_n^O + \phi_m^D + \sigma D_{nm} + \lambda_1 HB1_{nm} + \lambda_2 HB2_{nm} + u_{nm} \tag{3.12}$$

We use this regression to assess differences between political and statistical borders. The two home-bias dummies allows us to distinguish among three types of trade: (i) trade within NUTS2 regions (HB1 = HB2 = 1), (ii) trade across NUTS2 regions but within the same NUTS1 region (HB1 = 1, HB2 = 0), and (iii) trade across NUTS1 regions (HB1 = HB2 = 0). Therefore we interpret λ_1 as the difference between average trade of type (ii) and type (iii). We also interpret $\lambda_1 + \lambda_2$ as the difference between average trade of type (i) and type (iii). With this in mind, if it is political borders that cause the observed home bias, we would expect that:

1. For countries in group 1, we expect $\lambda_1 = 0$ and $\lambda_2 = 0$. The reason is that none of the three types of trade crosses a political border.

2. For countries in group 2, we expect $\lambda_1 = 0$, and $\lambda_2 > 0$. The reason is that only a fraction of the trade within a NUTS2 region (type (i)) does not cross a border but the totality of trade that goes across a NUTS1 or NUTS2 region (type (ii) and (iii)) crosses a border.

3. For countries in group 3, we also expect $\lambda_1 = 0$ and $\lambda_2 > 0$. The reason is that none of the trade within a NUTS2 region crosses a border while all of the trade that goes across a NUTS1 or NUTS2 region does.

4. For countries in group 4 the expected coefficients depend on the specific country. For Belgium and Germany, we expect $\lambda_1 > 0$ and $\lambda_2 > 0$. The reason is that trade across NUTS1 regions crosses two borders, while trade across NUTS2 regions but within a NUTS1 region crosses one border, and trade within NUTS2 regions crosses none. For the United Kingdom we expect $\lambda_1 > 0$ and $\lambda_2 = 0$. The reason is that only trade between NUTS1 regions crosses a border.[12]

Figure 3.18 shows our estimates of λ_1 and λ_2 for each country. There are four panels in the figure, each one showing one group of countries. Our hypotheses are confirmed by the data, with few exceptions.

Panel 1 presents the results for countries in group 1. Unfortunately we have to drop Estonia, Latvia, Lithuania and Slovenia since they only have one or two regions and therefore we cannot estimate the coefficients. We cannot estimate λ_1 for Portugal neither since the NUTS1 level includes the entire country. Thus we only have three estimates. Consistent with our expectations these estimates are small and not significative.

[12]For the United Kingdom, when we refer to NUTS1 we are really referring to the four nations: England, Northern Ireland, Scotland and Wales.

Panel 2 shows the results for countries in group 2. We are forced to drop Croatia and Ireland since they only have two regions and therefore we cannot estimate the coefficients. As in the case of Portugal, we cannot estimate λ_1 for Czech Republic, Finland, Norway, Slovakia and Switzerland. Consistent with our expectations we cannot reject that λ_1 is zero for the three countries where we can estimate this coefficient. The point estimates for λ_2 are positive but the standard errors are large. Only for two countries we can reject the hypothesis that λ_2 is zero.

Panel 3 shows the results for countries in group 3. As in the case of Portugal we cannot estimate λ_1 for Denmark. We cannot reject λ_1 is different than zero with the exception of Greece. The point estimates for λ_2 are positive for all countries except for Denmark. We can reject λ_2 is different from zero in all cases except for Denmark and Spain.

Finally, Panel 4 shows the results for countries in group 4. As expected, we find that λ_1 is positive and significant for all countries. Again confirming our hypotheses we find that λ_2 is zero for the United Kingdom and positive for Belgium and Germany.

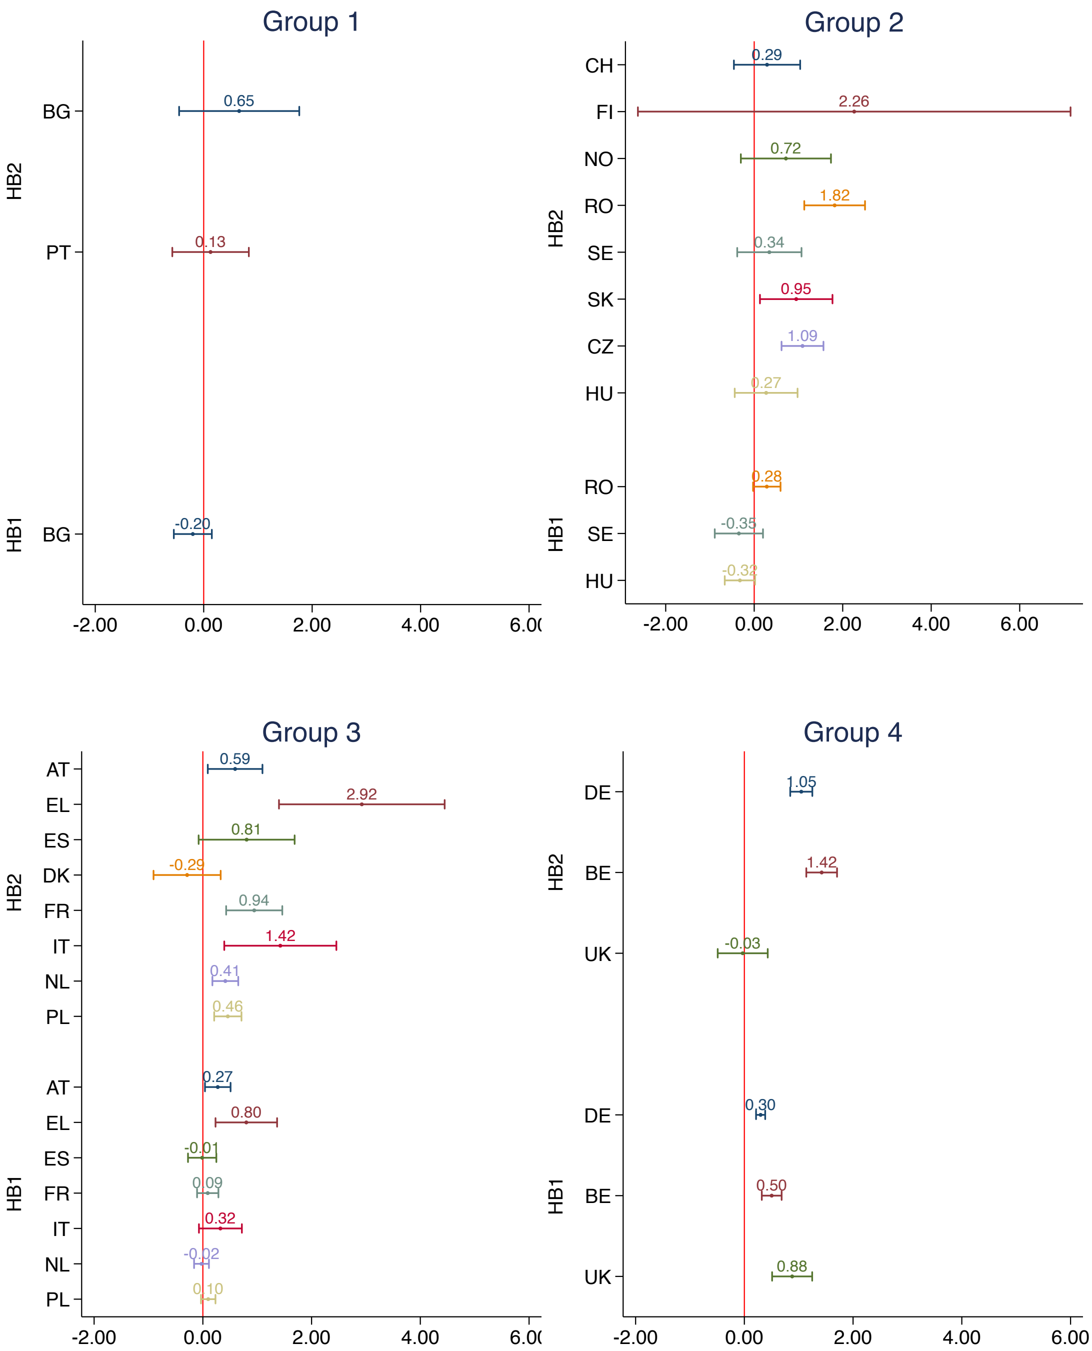

Figure 3.18: Home bias: Statistical and political borders

Notes: Figure shows the coefficients on the Home bias dummy at the NUTS1 level, λ_1 (HB1) and at the NUTS2 level λ_2 (HB2) from estimating regression 3.4.2 in each country in our sample. The four groups of countries are defined in the main text.

Taken together, these results indicate that political borders are an obstacle for trade while statistical borders are not. Thus it seems that political borders are an important factor behind the observed home bias in trade.

Finally, we perform the same exercise for each industry using the countries in group 3. We focus on this group because they provide the cleanest comparison between political and statistical borders. Recall that for this group political borders coincide with NUTS2 borders while NUTS1 border are purely statistical aggregates. Figure 3.19 shows the average estimate for each industry. The result confirm our aggregate findings. For essentially all industries we cannot reject that λ_1 is different from zero. For all industries we find that the point estimate of λ_2 is positive and significantly different from zero. We have performed this exercise for all country groups and the results are consistent with the average estimates reported above.

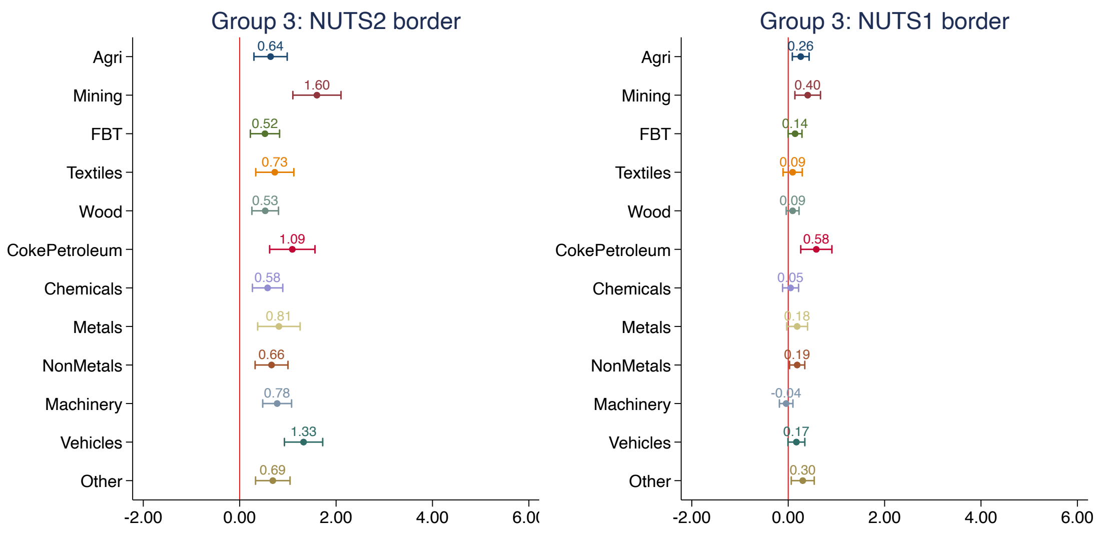

Figure 3.19: Home bias: Statistical and political borders

Notes: Figure shows the coefficients on the Home bias dummy at the NUTS1 level, λ_1 (HB1) and at the NUTS2 level λ_2 (HB2) from estimating regression 3.4.2 in each industry for the countries in group 3 in our sample. The group of countries is defined in the main text.

3.5 Concluding remarks

This paper has provided an integrated view of intranational and international trade in Europe using the new dataset we constructed in Santamaría et al. (2020). The picture that emerges is clear: (i) European regional trade has a strong home and country bias, (ii) geographic distance and national borders are important determinants of regional trade, but they cannot explain the strong home bias and (iii) this home bias is quite heterogeneous across regions and seems to be caused by political borders at the regional level.

Our findings open up several interesting questions. Why is it that political borders and geographical distance still remain such a strong impediment to trade in the context of Europe? How does the behaviour of governments shape regional trade flows, contributing to the large home bias in trade? Which factors explain the heterogeneous home bias and border effects that we see across countries? Providing a sound answer to these questions will have huge policy implications.

The key tool that we have used to explore trade interactions is the matrix of bilateral trade. This matrix provides a snapshot of all trade flows within and across European regions and countries. Unfortunately, many important economic indicators such as migration flows, foreign direct investment or bank lending relationships, are not yet available at the region-pair level in such a unified way. Has Europe achieved a higher degree of integration in these areas? It would also be useful to construct similar matrices for other social and cultural interactions such as travel and tourism, cultural exchanges, sports competitions, joint research projects, and so on. These matrices would help us form an accurate picture of how European citizens interact with each other.

Bibliography

Ahlfeldt, G. M., Redding, S. J., Sturm, D. M., and Wolf, N. (2015). The economics of density: Evidence from the berlin wall. *Econometrica*, 83(6):2127–2189.

Alchian, A. A. and Allen, W. R. (1964). *Exchange and Production; Theory in Use*. Wadsworth Publishing Company.

Allen, T. and Arkolakis, C. (2014). Trade and the topography of the spatial economy. *The Quarterly Journal of Economics*, 129(3):1085–1140.

Anderson, J. E. and Van Wincoop, E. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, 93(1):170–192.

Bailey, M., Gupta, A., Hillenbrand, S., Kuchler, T., Richmond, R., and Stroebel, J. (2020). International trade and social connectedness. *Journal of International Economics*, page 103418.

Bakker, J. D. (2020). Trade and agglomeration: Theory and evidence from france. *Available at SSRN 3757053*.

Baldwin, R., Francois, J., and Portes, R. (1997). The costs and benefits of eastern enlargement: the impact on the eu and central europe. *Economic Policy*, 12(24):125–176.

Becker, S. O. and Ichino, A. (2002). Estimation of average treatment effects based on propensity scores. *The Stata Journal*, 2(4):358–377.

Beestermöller, M. and Rauch, F. (2018). A dissection of trading capital: Trade in the aftermath of the Fall of the Iron Curtain. *The Journal of Economic History*, 78(2):358–393.

Berry, C. R. and Glaeser, E. L. (2005). The divergence of human capital levels across cities. (11617).

Bleakley, H. and Lin, J. (2012). Portage and Path Dependence *. *The Quarterly Journal of Economics*, 127(2):587–644.

Brooks, L., Gendron-Carrier, N., and Rua, G. (2018). The Local Impact of Containerization. (2018-045).

Brülhart, M., Carrere, C., and Trionfetti, F. (2012). How wages and employment adjust to trade liberalization: Quasi-experimental evidence from austria. *Journal of International Economics*, 86(1):68–81.

Brülhart, M., Carrère, C., and Robert-Nicoud, F. (2018). Trade and towns: Heterogeneous adjustment to a border shock. *Journal of Urban Economics*, 105(C):162–175.

Caliendo, L., Dvorkin, M., and Parro, F. (2019). Trade and labor market dynamics: General equilibrium analysis of the china trade shock. *Econometrica*, 87(3):741–835.

Caliendo, L., Parro, F., Opromolla, L. D., and Sforza, A. (2021). Goods and factor market integration: A quantitative assessment of the eu enlargement. *Journal of Political Economy*, 0(0):000–000.

Carlino, G. A. and Saiz, A. (2008). Beautiful city: Leisure amenities and urban growth.

Cervellati, M., Lazzaroni, S., Prarolo, G., and Vanin, P. (2019). Political geography and pre-industrial development: A theory and evidence for Europe 1000-1850. London, CEPR.

Chen, N. (2004). Intra-national versus international trade in the European Union: Why do national borders matter? *Journal of International Economics*, 63(1):93–118.

Chen, N. and Novy, D. (2012). On the measurement of trade costs: Direct vs. indirect approaches to quantifying standards and technical regulations. *World Trade Review*, 11:401.

Combes, P.-P., Lafourcade, M., and Mayer, T. (2005). The trade-creating effects of business and social networks: Evidence from France. *Journal of International Economics*, 66(1):1–29.

Coughlin, C. C. and Novy, D. (2012). Is the International Border Effect Larger than the Domestic Border Effect? Evidence from U.S. Trade. CEP discussion papers, Centre for Economic Performance, LSE.

Coughlin, C. C. and Novy, D. (2021). Estimating Border Effects: The Impact Of Spatial Aggregation. *International Economic Review*, 62(4):1453–1487.

Coşar, A. K. and Fajgelbaum, P. D. (2016). Internal geography, international trade, and regional specialization. *American Economic Journal: Microeconomics*, 8(1):24–56.

Crump, R. K., Hotz, V. J., Imbens, G. W., and Mitnik, O. A. (2009). Dealing with Limited Overlap in Estimation of Average Treatment Effects. *Biometrika*, 96(1):187–199.

Dauth, W., Findeisen, S., and Suedekum, J. (2014). The rise of the east and the far east: German labor markets and trade integration. *Journal of the European Economic Association*, 12(6):1643–1675.

Dawid, A. P. (1979). Conditional independence in statistical theory. *Journal of the Royal Statistical Society: Series B (Methodological)*, 41(1):1–15.

Desmet, K. and Rossi-Hansberg, E. (2014). Spatial development. *American Economic Review*, 104(4):1211–43.

Diamond, R. (2016). The determinants and welfare implications of us workers' diverging location choices by skill: 1980-2000. *American Economic Review*, 106(3):479–524.

Dijkstra, L. and Poelman, H. (2014). Regional working paper 2014. *A harmonised definition of cities and rural areas: the new degree of urbanisation. European Comission Directorate-General for Regional and Urban Policy: Working Paper*.

Dix-Carneiro, R. and Kovak, B. K. (2017). Trade liberalization and regional dynamics. *American Economic Review*, 107(10):2908–46.

Ducruet, C., Juhász, R., Nagy, D. K., and Steinwender, C. (2020). All aboard: The effects of port development.

Dustmann, C. and Frattini, T. (2012). Immigration: The European Experience. (2012001).

Ehrlich, M. v. and Overman, H. G. (2020). Place-based policies and spatial disparities across european cities. *Journal of Economic Perspectives*, 34(3):128–49.

Fajgelbaum, P. and Redding, S. J. (2018). Trade, structural transformation and development: Evidence from argentina 1869-1914. Working Paper 20217, National Bureau of Economic Research.

Feenstra, R. C. (2002). Border effects and the gravity equation: Consistent methods for estimation. *Scottish Journal of Political Economy*, 49(5):491–506.

Feyrer, J. (2019). Trade and income—exploiting time series in geography. *American Economic Journal: Applied Economics*, 11(4):1–35.

Fukao, K. and Okubo, T. (2004). Why has the border effect in the japanese market declined? the role of business networks in East Asia. Technical report, Graduate Institute of International Studies.

Gallego, N. and Llano, C. (2015). Thick and thin borders in the European Union: How deep internal integration is within countries, and how shallow between them. *The World Economy*, 38(12):1850–1879.

Gancia, G., Ponzetto, G. A., and Ventura, J. (2020). Globalization and political structure. *CEPR Discussion Paper No. DP11159*.

Ganong, P. and Shoag, D. (2017). Why has regional income convergence in the u.s. declined? *Journal of Urban Economics*, 102:76–90.

Garmendia, A., Llano, C., Minondo, A., and Requena, F. (2012). Networks and the disappearance of the intranational home bias. *Economics Letters*, 116(2):178 – 182.

Giannone, E. (2017). Skill-Biased Technical Change and Regional Convergence. Technical report.

Gil-Pareja, S., Llorca-Vivero, R., Martínez-Serrano, J. A., and Oliver-Alonso, J. (2005). The border effect in Spain. *World Economy*, 28(11):1617–1631.

Glaeser, E. L., Kolko, J., and Saiz, A. (2001). Consumer city. *Journal of Economic Geography*, 1(1):27–50.

Grossman, G. M. and Helpman, E. (2001). *Special Interest Politics*. MIT press.

Head, K. and Mayer, T. (2009). Illusory border effects: distance mismeasurement inflates estimates of home bias in trade. In *In: The Gravity Model in International Trade: Advances and Applications. Editors: Bergeijk and Brakman*. Citeseer.

Head, K. and Mayer, T. (2014). Gravity equations: Workhorse, toolkit, and cookbook. In *Handbook of International Economics*, volume 4, pages 131–195. Elsevier.

Head, K. and Mayer, T. (2021). The united states of europe: A gravity model evaluation of the four freedoms. *Journal of Economic Perspectives*, 35(2):23–48.

Helpman, E., Melitz, M., and Rubinstein, Y. (2008). Estimating trade flows: Trading partners and trading volumes. *The Quarterly Journal of Economics*, 123(2):441–487.

Henderson, J. V. and Wang, H. G. (2005). Aspects of the rural-urban transformation of countries. *Journal of Economic Geography*, 5(1):23–42.

Henderson, J. V. and Wang, H. G. (2007). Urbanization and city growth: The role of institutions. *Regional Science and Urban Economics*, 37(3):283–313.

Henrekson, M., Torstensson, J., and Torstensson, R. (1997). Growth effects of european integration. *European Economic Review*, 41(8):1537–1557.

Herz, B. and Varela-Irimia, X.-L. (2020). Border Effects in European Public Procurement. *Journal of Economic Geography*.

Hillberry, R. (1999). Explaining the 'border effect': what can we learn from disaggregated commodity flow data. *Indiana University Graduate Student Economics Working Paper Series*.

Hillberry, R. and Hummels, D. (2008). Trade responses to geographic frictions: A decomposition using micro-data. *European Economic Review*, 52(3):527–550.

Hummels, D. and Skiba, A. (2004). Shipping the good apples out? an empirical confirmation of the alchian-allen conjecture. *Journal of political Economy*, 112(6):1384–1402.

Imbens, G. W. and Rubin, D. B. (2015). *Causal inference in Statistics, Social, and Biomedical Sciences*. Cambridge University Press.

Kennan, J. (2017). Open borders in the european union and beyond: Migration flows and labor market implications. (23048).

Krugman, P. (1991). Increasing returns and economic geography. *Journal of political economy*, 99(3):483–499.

Krugman, P. and Elizondo, R. L. (1996). Trade policy and the third world metropolis. *Journal of Development Economics*, 49(1):137–150. Increasing Returns, Monopolistic Competition and Economic Development.

Marin, A. G., Potlogea, A. V., Voigtländer, N., and Yang, Y. (2020). Cities, Productivity, and Trade. NBER Working Papers 28309, National Bureau of Economic Research, Inc.

McCallum, J. (1995). National borders matter: Canada-US regional trade patterns. *The American Economic Review*, 85(3):615–623.

Monte, F., Redding, S. J., and Rossi-Hansberg, E. (2018). Commuting, migration, and local employment elasticities. *American Economic Review*, 108(12):3855–90.

Moretti, E. (2012). *The new geography of jobs*. Houghton Mifflin Harcourt.

Mori, T., Smith, T. E., and Hsu, W.-T. (2020). Common power laws for cities and spatial fractal structures. *Proceedings of the National Academy of Sciences*, 117(12):6469–6475.

Mori, T. and Wrona, J. (2021). Centrality bias in inter-city trade. *KIER Discussion Paper*, 1056:1–52.

Morten, M. and Oliveira, J. (2018). The effects of roads on trade and migration : Evidence from a planned capital city .

Nagy, D. (2020). Trade and urbanization: Evidence from hungary. *American Economic Journal: Microeconomics (Revise and resubmit)*.

Nitsch, V. (2000). National borders and international trade: evidence from the European Union. *Canadian Journal of Economics/Revue canadienne d'économique*, 33(4):1091–1105.

Nitsch, V. and Wolf, N. (2013). Tear down this wall: on the persistence of borders in trade. *Canadian Journal of Economics/Revue canadienne d'économique*, 46(1):154–179.

Pascali, L. (2017). The wind of change: Maritime technology, trade, and economic development. *American Economic Review*, 107(9):2821–54.

Rappaport, J. (2007). Moving to nice weather. *Regional Science and Urban Economics*, 37(3):375–398.

Rappaport, J. (2008). Consumption amenities and city population density. *Regional Science and Urban Economics*, 38(6):533–552.

Redding, S. J. and Rossi-Hansberg, E. (2017). Quantitative spatial economics. *Annual Review of Economics*, 9(1):21–58.

Redding, S. J. and Sturm, D. M. (2008). The costs of remoteness: Evidence from german division and reunification. *American Economic Review*, 98(5):1766–97.

Rodríguez-Pose, A. (2018). The revenge of the places that don't matter (and what to do about it). *Cambridge journal of regions, economy and society*, 11(1):189–209.

Santamaría, M. A., Ventura, J., and Yeşilbayraktar, U. (2020). Borders within europe. (28301).

Schulze, M.-S. and Wolf, N. (2009). On the origins of border effects: insights from the Habsburg Empire. *Journal of Economic Geography*, 9(1):117–136.

Shapiro, J. M. (2006). Smart cities: quality of life, productivity, and the growth effects of human capital. *The review of economics and statistics*, 88(2):324–335.

Silva, J. S. and Tenreyro, S. (2006). The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658.

Simonovska, I. and Waugh, M. E. (2014). The elasticity of trade: Estimates and evidence. *Journal of international Economics*, 92(1):34–50.

Spolaore, E. and Alesina, A. (2003). *The Size of Nations*. Mit Press Cambridge, MA.

Tinbergen, J. . (1962). *Economic geography and international inequality*. New York, NY: Twentieth Century Fund.

Tombe, T. and Zhu, X. (2019). Trade, Migration, and Productivity: A Quantitative Analysis of China. *American Economic Review*, 109(5):1843–1872.

Turrini, A. and van Ypersele, T. (2010). Traders, courts, and the border effect puzzle. *Regional Science and Urban Economics*, 40(2-3):81–91.

Weber, S., Péclat, M., and Warren, A. (2021). Travel distance and travel time using Stata: New features and major improvements in georoute. IRENE Working Papers 21-04, IRENE Institute of Economic Research.

Wei, S.-J. (1996). Intra-national versus international trade: how stubborn are nations in global integration? Technical report, National Bureau of Economic Research.

Wolf, H. C. (2000). Intranational home bias in trade. *Review of Economics and Statistics*, 82(4):555–563.

Appendix A

APPENDIX

A.1 Appendix: Chapter 1

A.1.1 Data and Empirical Results

Additional Tables

Table A.1.1: Changes in Population Growth: Cities

	Dep: Annualized Pop. Growth	Dep: Annualized Pop. Growth
Acc x Hub	-0.0856	-0.0243
	(0.0934)	(0.0612)
Border(0-50) x Acc x Hub	0.865***	0.816***
	(0.252)	(0.245)
Border(50-100) x Acc x Hub	0.517***	0.467***
	(0.153)	(0.141)
Border(100-150) x Acc x Hub	0.457	0.408
	(0.479)	(0.476)
Border(0-50) x Hub		-0.464**
		(0.196)
Border(50-100) x Hub		-0.217***
		(0.0786)
Border(100-150) x Hub		-0.397
		(0.462)
Year FE	Yes	Yes
Location FE	Yes	No
Observations	604	604
R^2	0.710	0.045

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.1.2: Poland (NUTS3): Changes in Employment

	Dep: Employment	Dep: Urban Employment
Border(0to50km) x Acc	0.0101**	0.0151***
	(0.00296)	(0.00376)
Border(50to100km) x Acc	0.00975**	0.0165***
	(0.00292)	(0.00386)
Border(100to150km) x Acc	-0.00264	0.00549
	(0.00453)	(0.00422)
Year FE	Yes	Yes
Location FE	Yes	Yes
Observations	1320	1320
R^2	0.765	0.636

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A.1.3: Robustness: Main Results

	α	Average change in real income	Change in urban population	Change in urbanization
Baseline Specification	0.133	.214%	.393 %	.084%
10% higher taste heterogeneity parameter θ	0.138	.214%	.371%	.081%
10% lower taste heterogeneity parameter θ	0.127	.214%	.413%	.089%
10% higher shipping cost parameter ψ	0.133	.214%	.387%	.083%
10% lower shipping cost parameter ψ.	0.132	.214%	.397%	.085
10% higher transport cost parameter ϕ	0.135	.200%	.385%	.083
10% lower transport cost parameter ϕ	0.130	.231%	.401%	.086
10% higher elasticity of substitution σ	0.141	.093%	.164%	.035
10% lower elasticity of substitution σ	0.112	.502%	.998%	.210

List of Hubs

City Code	City Name	Country
AT003C1	Linz	Austria
AT002C1	Graz	Austria
AT001C1	Wien	Austria
CZ001C1	Praha	Czech Republic
CZ003C1	Ostrava	Czech Republic
CZ002C1	Brno	Czech Republic
CZ004C1	Plzen	Czech Republic

DE002C1	Hamburg	Germany
DE013C1	Hannover	Germany
DE012C1	Bremen	Germany
DE011C1	Düsseldorf	Germany
DE006C1	Essen	Germany
DE036C1	Mönchengladbach	Germany
DE034C1	Bonn	Germany
DE004C1	Köln	Germany
DE010C1	Dortmund	Germany
DE005C1	Frankfurt am Main	Germany
DE020C1	Wiesbaden	Germany
DE007C1	Stuttgart	Germany
DE035C1	Karlsruhe	Germany
DE502C1	Mannheim	Germany
DE027C1	Freiburg im Breisgau	Germany
DE003C1	München	Germany
DE014C1	Nürnberg	Germany
DE033C1	Augsburg	Germany
DE001C1	Berlin	Germany
DE505C1	Chemnitz	Germany
DE009C1	Dresden	Germany
DE008C1	Leipzig	Germany
DE018C1	Halle (Saale)	Germany
DE019C1	Magdeburg	Germany
HU001C1	Budapest	Hungary
HU005C1	Debrecen	Hungary
HU002C1	Miskolc	Hungary
IT004C1	Torino	Italy
IT006C1	Genova	Italy
IT002K1	Milano	Italy
IT029C1	Brescia	Italy

IT012C1	Verona	Italy
IT011C1	Venezia	Italy
IT009C1	Bologna	Italy
IT007C1	Firenze	Italy
IT001C1	Roma	Italy
IT003K1	Napoli	Italy
IT031C1	Foggia	Italy
IT008C1	Bari	Italy
IT022C1	Taranto	Italy
IT005C1	Palermo	Italy
IT501C1	Messina	Italy
IT010C1	Catania	Italy
IT027C1	Cagliari	Italy
PL004C1	Wrocław	Poland
PL008C1	Bydgoszcz	Poland
PL013C1	Toruń	Poland
PL009C1	Lublin	Poland
PL002C1	Łódź	Poland
PL003C1	Kraków	Poland
PL025C1	Radom	Poland
PL001C1	Warszawa	Poland
PL015C1	Rzeszów	Poland
PL011C1	Białystok	Poland
PL006C1	Gdańsk	Poland
PL501C1	Gdynia	Poland
PL506C1	Bielsko-Biała	Poland
PL024C1	Częstochowa	Poland
PL001K1	Katowice	Poland
PL012C1	Kielce	Poland
PL014C1	Olsztyn	Poland
PL005C1	Poznań	Poland

PL007C1	Szczecin	Poland
SI001C1	Ljubljana	Slovenia
SK001C1	Bratislava	Slovakia
SK002C1	Kosice	Slovakia

Additional Figures

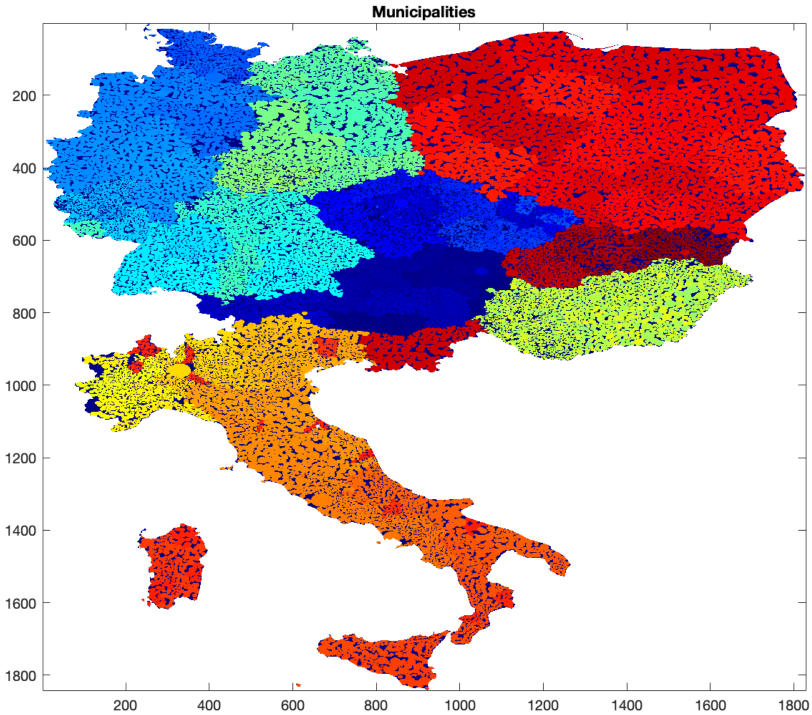

Figure A.1.1: Model Geography: Municipalities

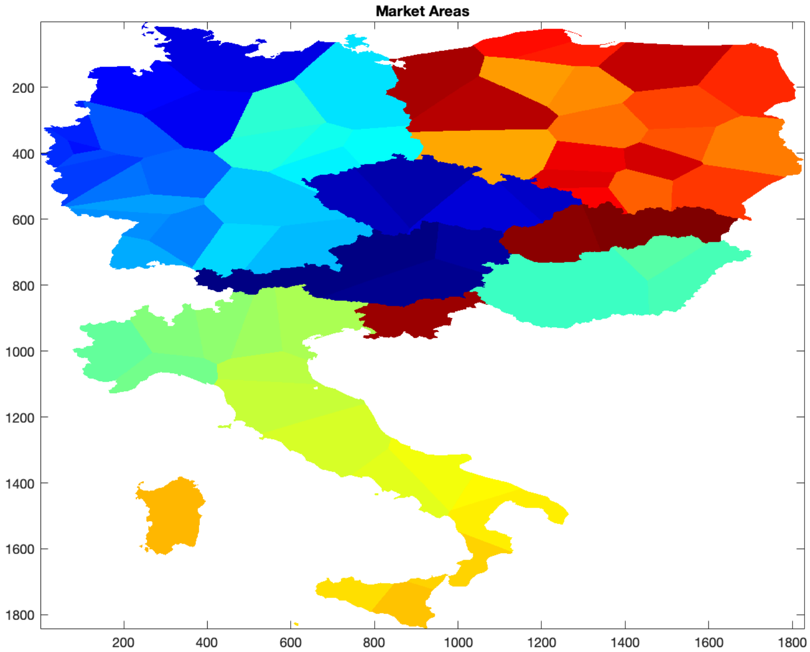

Figure A.1.2: Model Geography: Market Areas

A.1.2 Deriving Equilibrium Conditions

In this section I derive equations (7) to (11), which characterize the spatial equilibrium of the model. Due to the CES assumption in consumer demand for tradable goods, price index at hub h is given by

$$P_h = \left[\sum_h p_h^{1-\sigma} L_h \tau_{oh}^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \tag{B1}$$

Each worker produces a uniquely differentiated good and is endowed with one unit of labor. Producing one unit of each good requires one unit of labor. Since marginal cost of production does not increase with the quantity produced and workers do not value leisure, workers utility is strictly increasing in their output. The intuition is simple: Given an elasticity of substitution greater than 1, workers revenue is increasing in their quantity produced. As a consequence of this, each worker will produce the maximum quantity i.e $x^j = 1\ \forall j$. As neither the demand nor the supply of a product depends on the workers identity. Thus, in equilibrium all workers trading via the same hub h set the same price for their product i.e $p_h^j = p_h\ \forall j \in h$. Given this, one can rewrite the goods market clearing condition provided in equation (6) as

$$p_m^{\sigma} = \sum_o \tau_{ho}^{1-\sigma}.P_o^{\sigma-1} p_o L_o$$

Given the Gumbel distribution of idiosyncratic amenities, the share of population living at residential location r is given by

$$\frac{L(r)}{\bar{L}} = \frac{[e^{a(r)+\max_h \zeta(\mu_h,r)^{-1}\frac{p_h}{P_h}}]^{\theta-1}}{\sum_s [e^{a(s)+\max_o \zeta(\mu_o,s)^{-1}\frac{p_o}{P_o}}]^{\theta-1}} \tag{B3}$$

Thus the fraction of population living at residential location r is equal to the conventional logit probability, or the fraction of non-idiosyncratic utility at location r. Naturally, all workers that reside at location r choose the same trading hub denoted by $\mu(r)$. Denoting real income by $\omega_h = \frac{p_h}{P_h}$, I can restate this equation as

$$\log L(r) = v + \theta^{-1}[a(r) + \zeta(\mu(r), r)^{-1}\omega_{\mu(r)}] \tag{B4}$$

Since by definition $\mu(r) = \text{argmax}_h \zeta(\mu_h, r)^{-1}\omega_{\mu_h}$, or the hub that provides the best economic incentives at location r, it must hold that

$$\zeta(\mu(r), r)^{-1}\omega_{\mu(r)} \geq \zeta(\mu_h, r)^{-1}\omega_{\mu_h} \qquad \forall h. \tag{B5}$$

Using $P_h = \frac{p_h}{\omega_h}$, one can rewrite equation (B1) and (B2) as

$$p_h^{1-\sigma}\omega_h^{\sigma-1} = \sum_o \tau_{oh}^{1-\sigma} p_o^{1-\sigma} L_o \tag{B1'}$$

$$p_h^\sigma = \sum_o \tau_{ho}^{1-\sigma} p_o^\sigma \omega_o^{1-\sigma} L_o \tag{B2'}$$

Under the assumption of symmetric shipping costs, these two equations can be reduced to one by using the guess and verify trick used in Allen and Arkolakis (2014). In particular, guess that price and real income are related via

$$p_h = \omega_h^k$$

where k is a constant. Under this assumption, equations (B1') and (B2') become

$$\omega_h^{k(1-\sigma)+\sigma-1} = \sum_o \tau_{ho}^{1-\sigma} L_o \omega_h^{k(1-\sigma)}$$

$$\omega_h^{k\sigma} = \sum_o \tau_{ho}^{1-\sigma} L_o \omega_h^{k(1-\sigma)}$$

and reduce down to a single equation for $k = \frac{\sigma-1}{2\sigma-1}$

$$\omega_h^{\frac{\sigma(\sigma-1}{2\sigma-1}} = \sum_o \tau_{ho}^{1-\sigma} L_o \omega_h^{-\frac{(\sigma-1)^2}{2\sigma-1}} \tag{B6}$$

Here I follow Nagy (2020) and define the right hand side of equation (B6) as the market access of hub h or MA_h and get

$$\omega_h = MA_h^{\frac{2\sigma-1}{\sigma(\sigma-1)}} \tag{B8}$$

Using this definition, one can substitute for ω_h in equations (B4), (B5) and (B6) to get equations (7), (8) and (9).

A.2 Appendix: Chapter 2

A.2.1 Additional Figures

Figure A.2.1: Correlation with aggregate international trade data

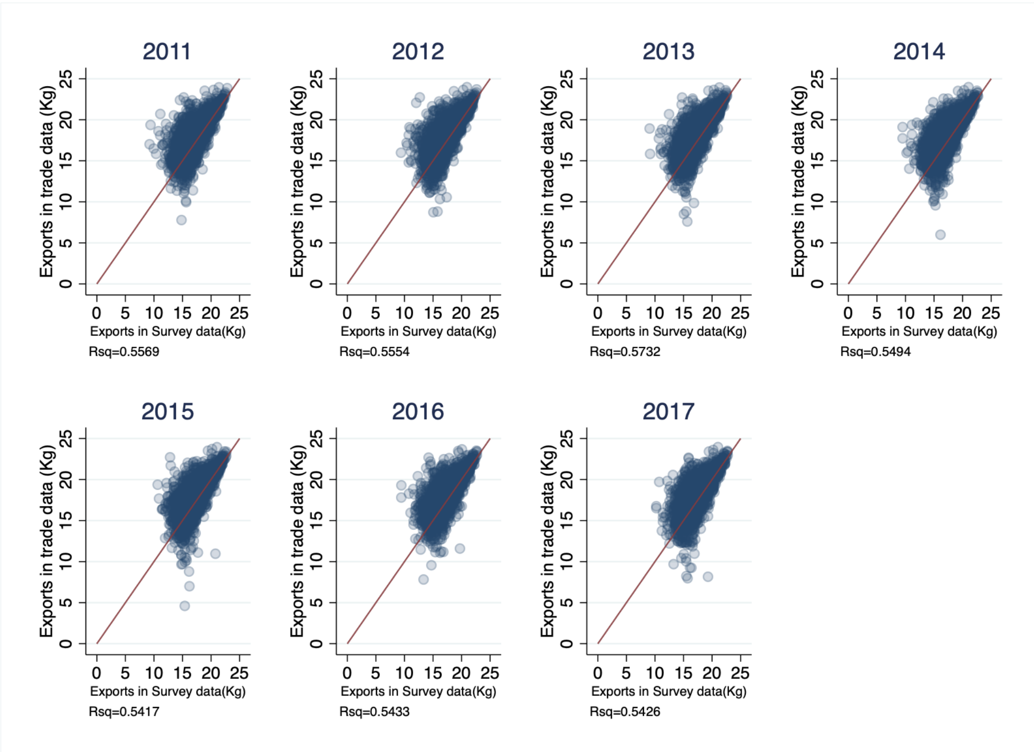

Notes: The figures show the correlation between exports and shipments in the ERFT survey in kilograms in each year. The Y-axis represents (log) bilateral trade (kg) by country-pair-industry-year using international trade data from Eurostat. The X-axis shows bilateral shipments by road (kg) aggregated by country-pair-industry-year obtained from the ERFT survey.

Figure A.2.2: Correlation with aggregate international trade data

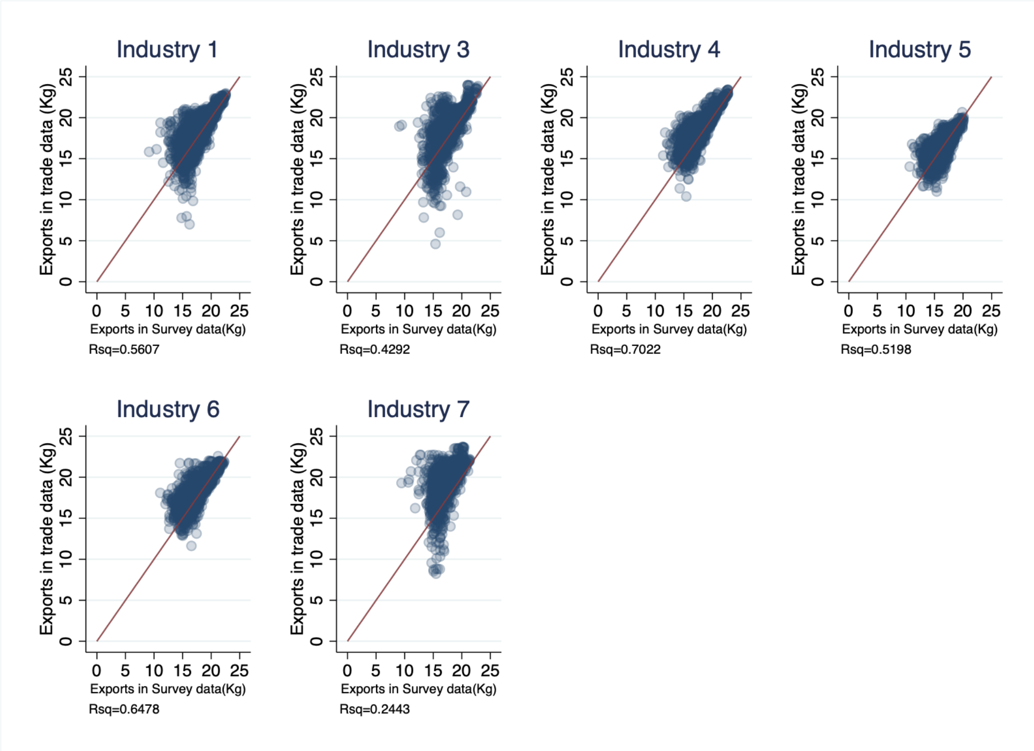

Notes: The figures show the correlation between exports and shipments in the ERFT survey in kilograms in each industry. The Y-axis represents (log) bilateral trade (kg) by country-pair-industry-year using international trade data from Eurostat. The X-axis shows bilateral shipments by road (kg) aggregated by country-pair-industry-year obtained from the ERFT survey.

Figure A.2.3: Correlation with aggregate international trade data

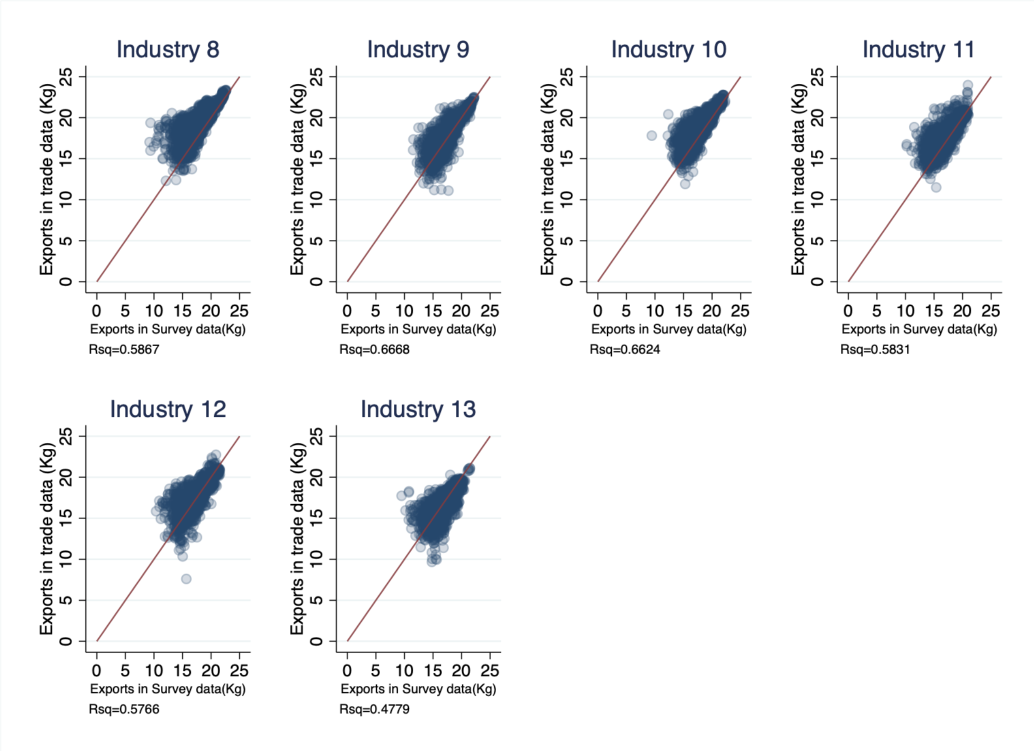

Notes: The figures show the correlation between exports and shipments in the ERFT survey in kilograms in each industry. The Y-axis represents (log) bilateral trade (kg) by country-pair-industry-year using international trade data from Eurostat. The X-axis shows bilateral shipments by road (kg) aggregated by country-pair-industry-year obtained from the ERFT survey.

Figure A.2.4: Out-of-sample Estimates

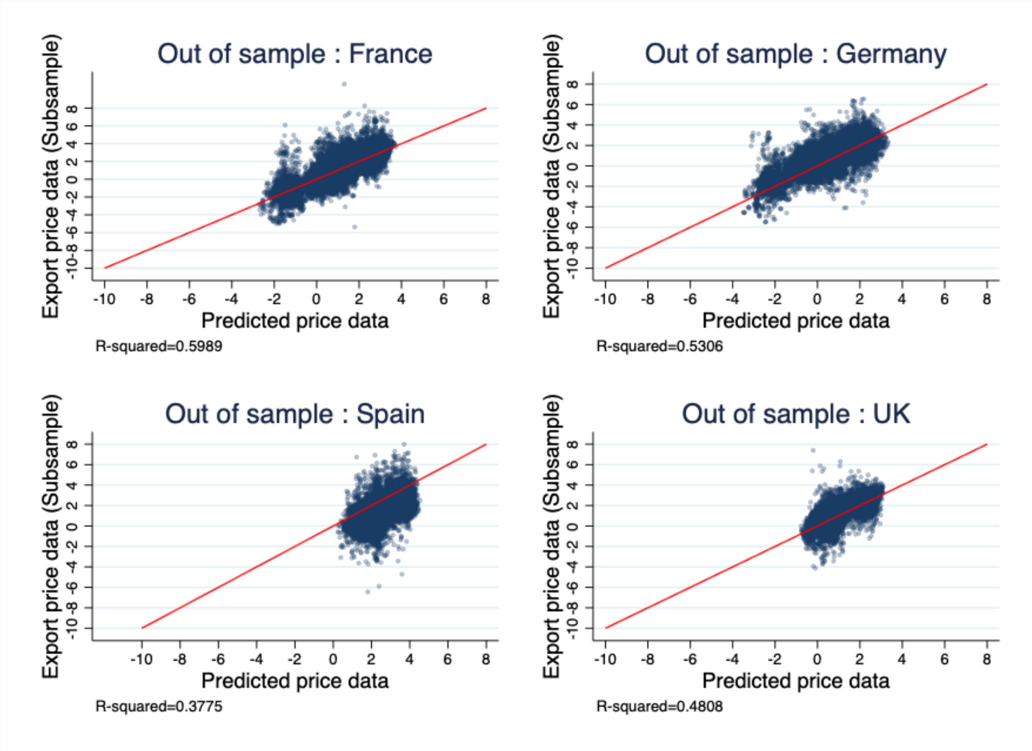

Notes: These figures show the out-of sample check to confirm the performance of the price imputation methodology. Each figure reports the (log) price per kg of exports of France, Germany, Spain and UK to all the countries in our sample by industry and year. The X-axis reports the estimated (log) price per kg of shipment in our regional trade dataset aggregated at the country-pair-industry-year level, predicted when we drop France, Germany, Spain and UK respectively.

Figure A.2.5: Country-to-Country Estimates

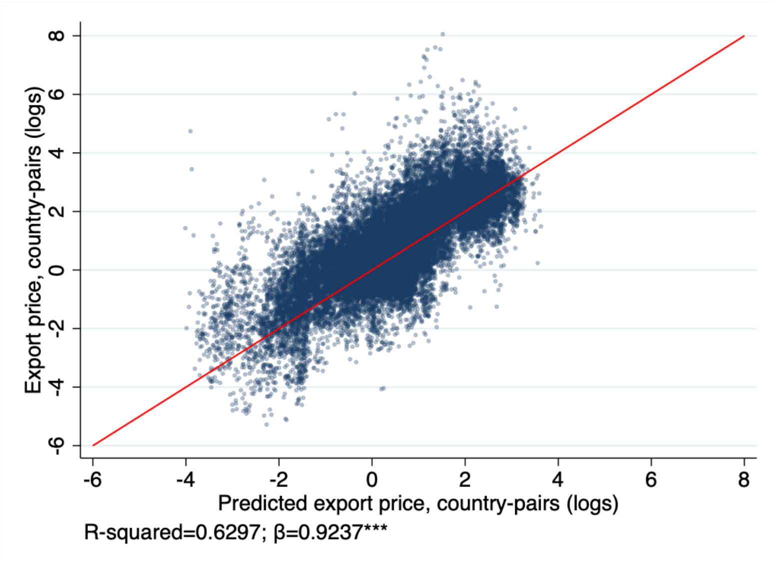

Notes: This figure shows the correlation between the price per kg of exports in international trade data and the imputed prices in our sample. The Y-axis reports the (log) price per kg of exports by country-pair, industry and year. The X-axis reports the estimated (log) price per kg of shipment in our regional trade dataset aggregated at the country-pair-industry-year level. In this figure we use all countries except France, Germany, Spain and United Kingdom.

Figure A.2.6: Correlation with aggregate international trade data

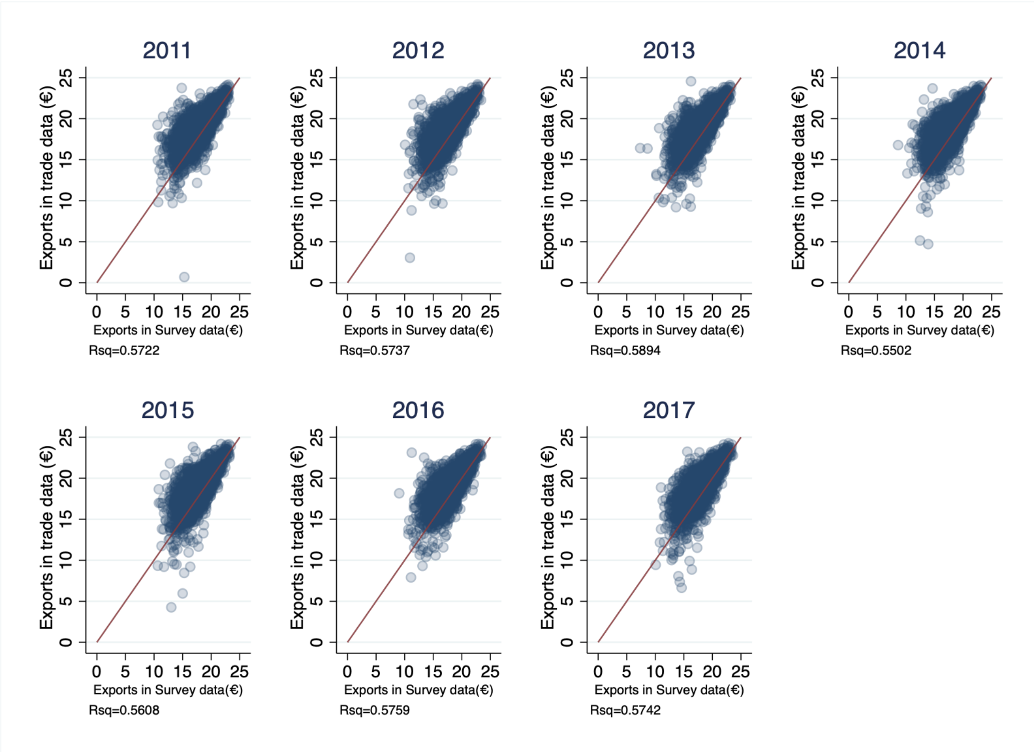

Notes: The figures show the correlation between exports and shipments in the ERFT survey in euros in each year. The Y-axis represents (log) bilateral trade (euros) by country-pair-industry-year using international trade data from Eurostat. The X-axis shows bilateral shipments by road (euros) aggregated by country-pair-industry-year obtained from the ERFT survey.

Figure A.2.7: Correlation with aggregate international trade data

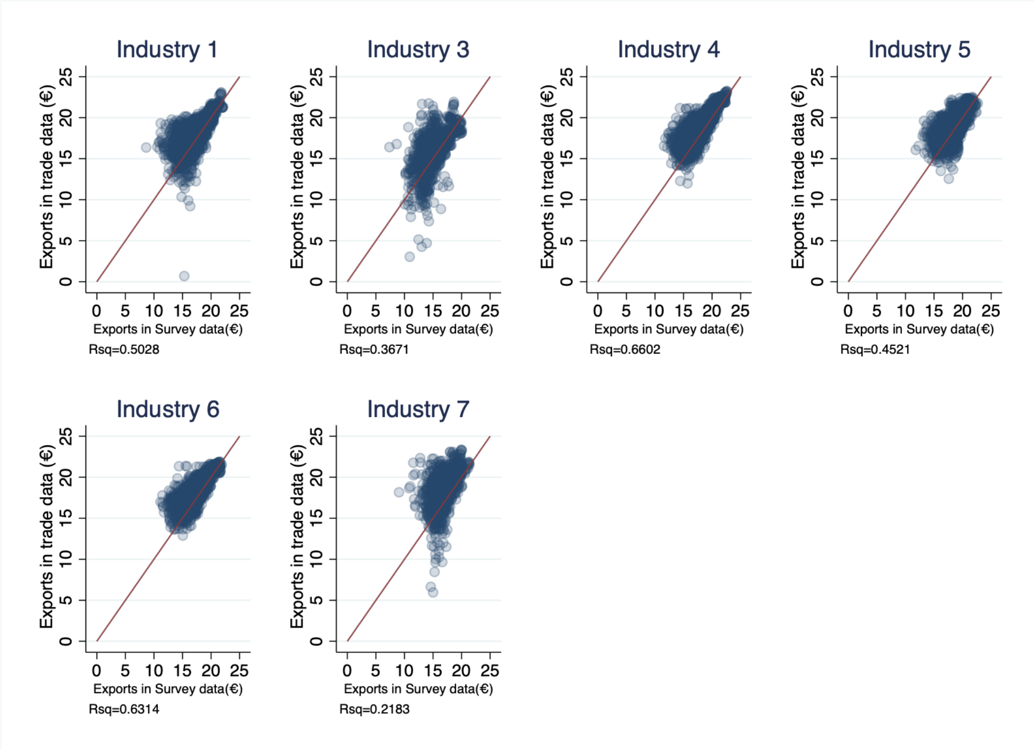

Notes: The figures show the correlation between exports and shipments in the ERFT survey in euros in each industry. The Y-axis represents (log) bilateral trade (euros) by country-pair-industry-year using international trade data from Eurostat. The X-axis shows bilateral shipments by road (euros) aggregated by country-pair-industry-year obtained from the ERFT survey.

Figure A.2.8: Correlation with aggregate international trade data

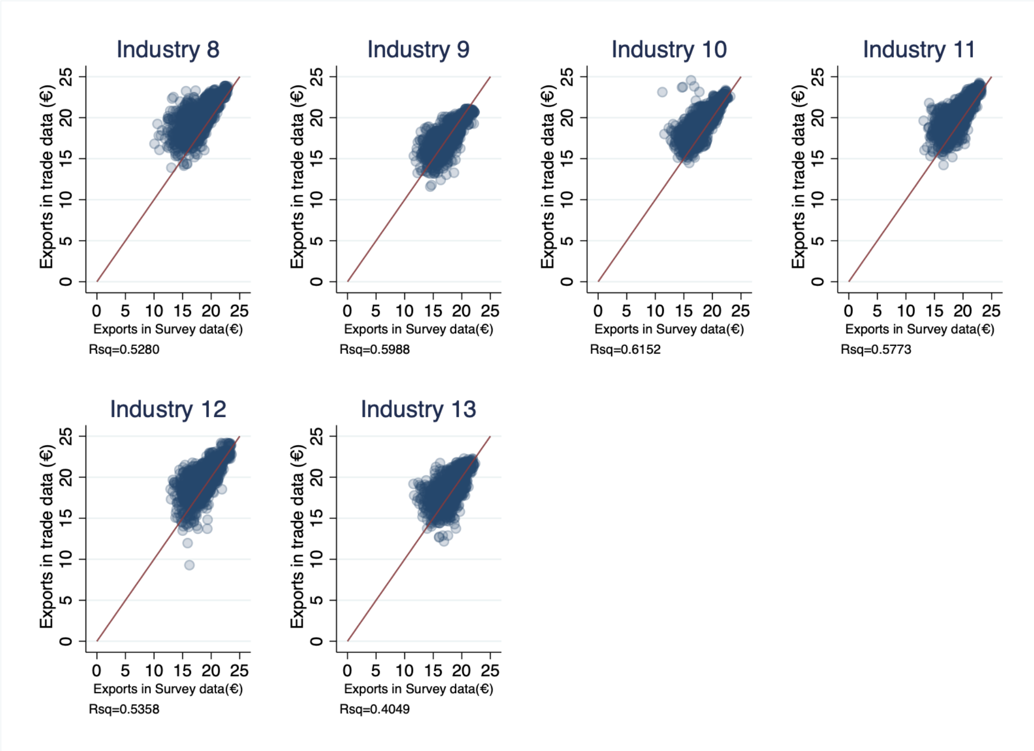

Notes: The figures show the correlation between exports and shipments in the ERFT survey in euros in each industry. The Y-axis represents (log) bilateral trade (euros) by country-pair-industry-year using international trade data from Eurostat. The X-axis shows bilateral shipments by road (euros) aggregated by country-pair-industry-year obtained from the ERFT survey.

Figure A.2.9: Composition of regions in block 1

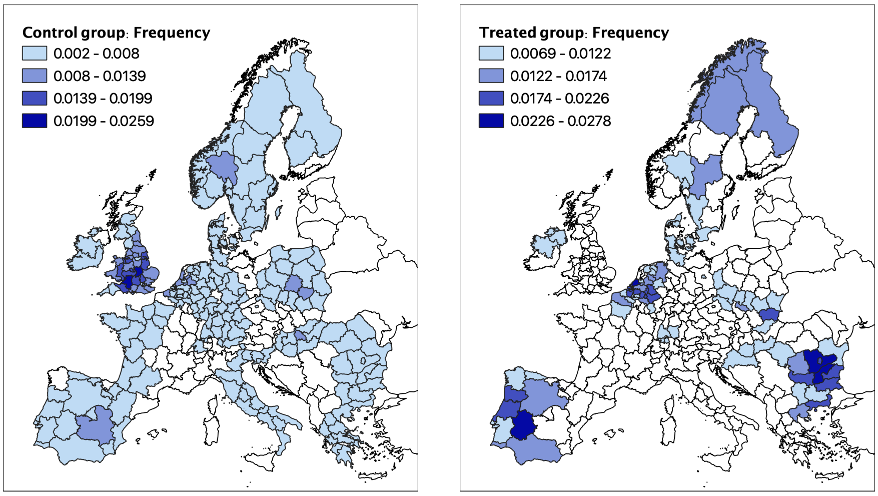

A) Control group **B) Treated group**

Notes: This figure shows the regions that are part of the block. The shading represents the frequency with which each region appears (as a part of a pair) in the control group (first panel) and treated group (second panel) in the block.

Figure A.2.10: Composition of regions in block 2

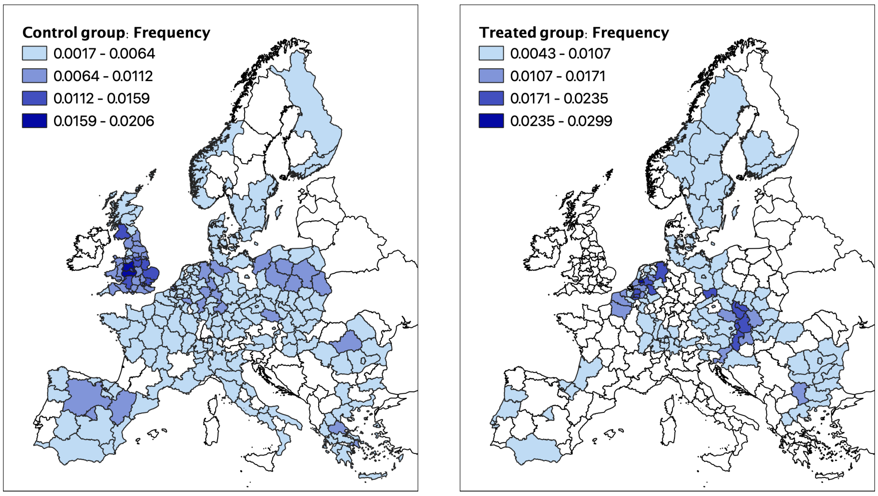

A) Control group **B) Treated group**

Notes: This figure shows the regions that are part of the block. The shading represents the frequency with which each region appears (as a part of a pair) in the control group (first panel) and treated group (second panel) in the block.

Figure A.2.11: Composition of regions in block 3

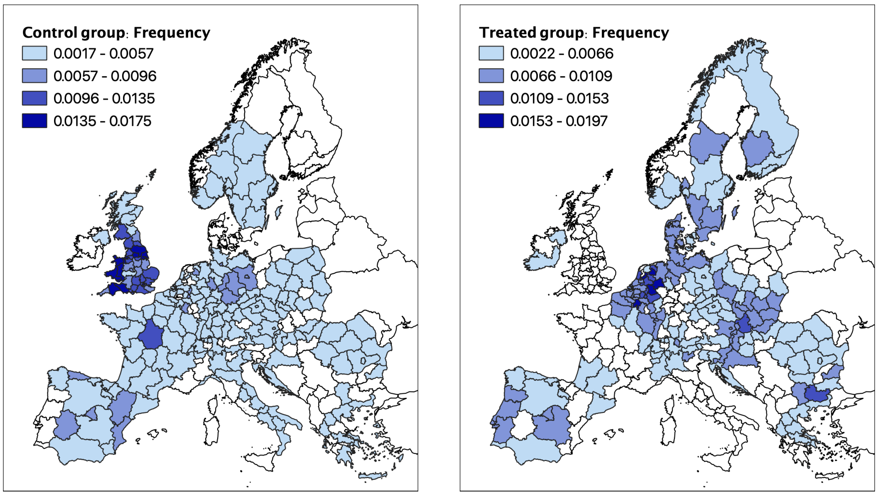

A) Control group **B) Treated group**

Notes: This figure shows the regions that are part of the block. The shading represents the frequency with which each region appears (as a part of a pair) in the control group (first panel) and treated group (second panel) in the block.

Figure A.2.12: Composition of regions in block 5

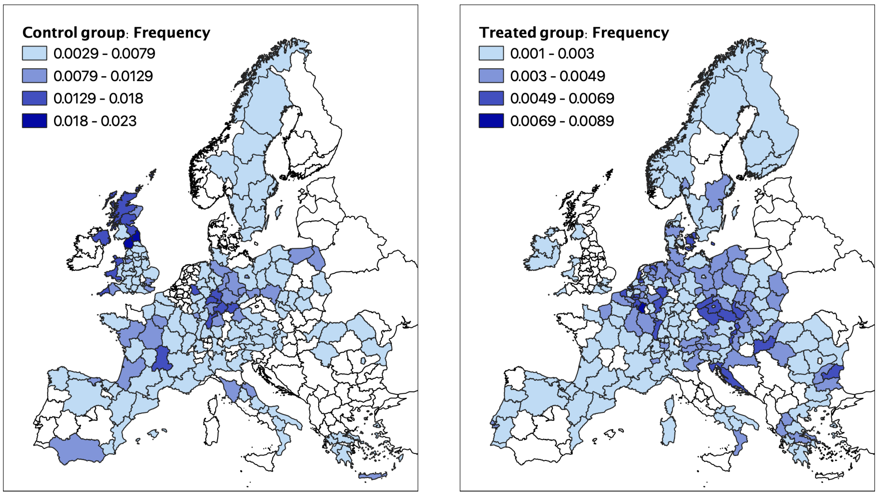

A) Control group **B) Treated group**

Notes: This figure shows the regions that are part of the block. The shading represents the frequency with which each region appears (as a part of a pair) in the control group (first panel) and treated group (second panel) in the block.

Figure A.2.13: Composition of regions in block 6

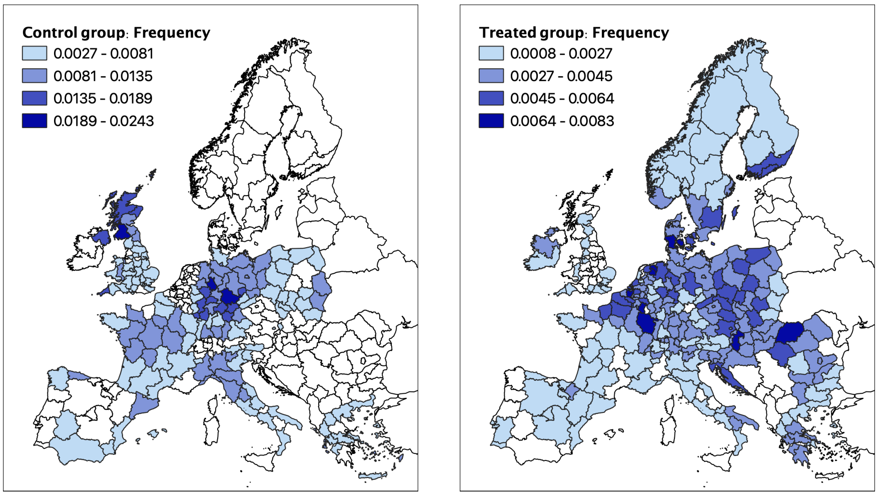

A) Control group **B) Treated group**

Notes: This figure shows the regions that are part of the block. The shading represents the frequency with which each region appears (as a part of a pair) in the control group (first panel) and treated group (second panel) in the block.

Figure A.2.14: Composition of regions in block 7

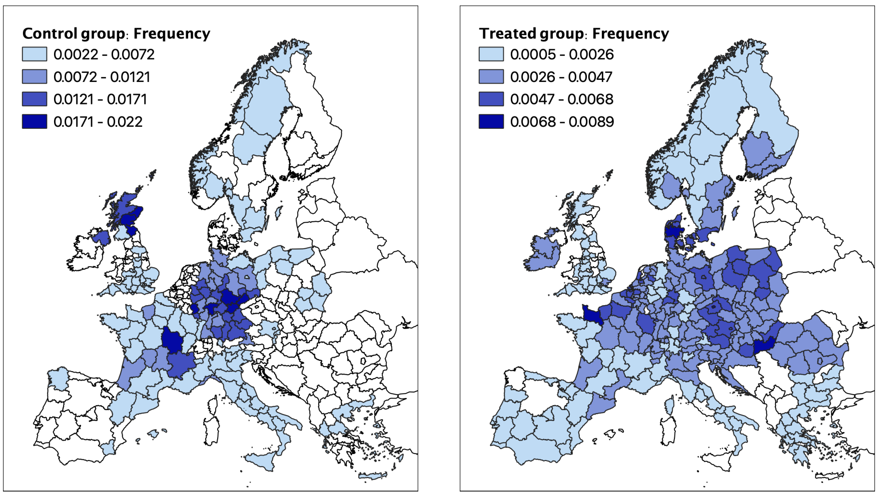

A) Control group **B) Treated group**

Notes: This figure shows the regions that are part of the block. The shading represents the frequency with which each region appears (as a part of a pair) in the control group (first panel) and treated group (second panel) in the block.

Figure A.2.15: Composition of regions in block 8

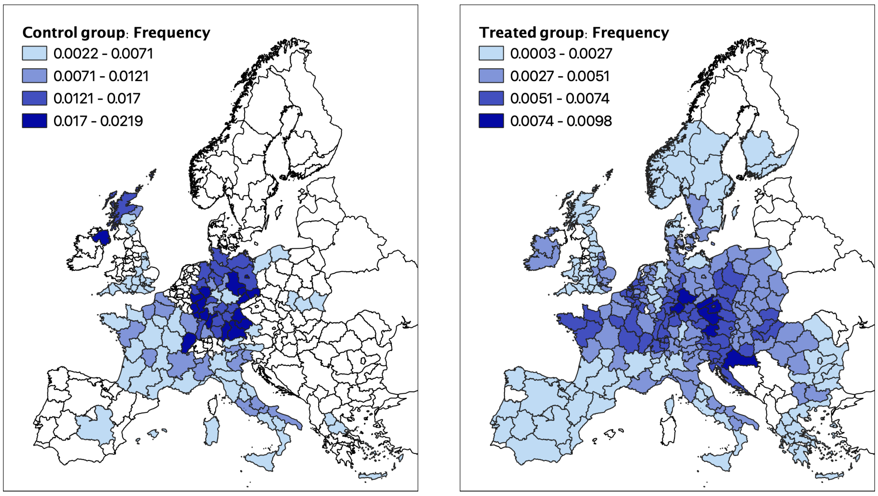

A) Control group **B) Treated group**

Notes: This figure shows the regions that are part of the block. The shading represents the frequency with which each region appears (as a part of a pair) in the control group (first panel) and treated group (second panel) in the block.

Figure A.2.16: Composition of regions in block 9

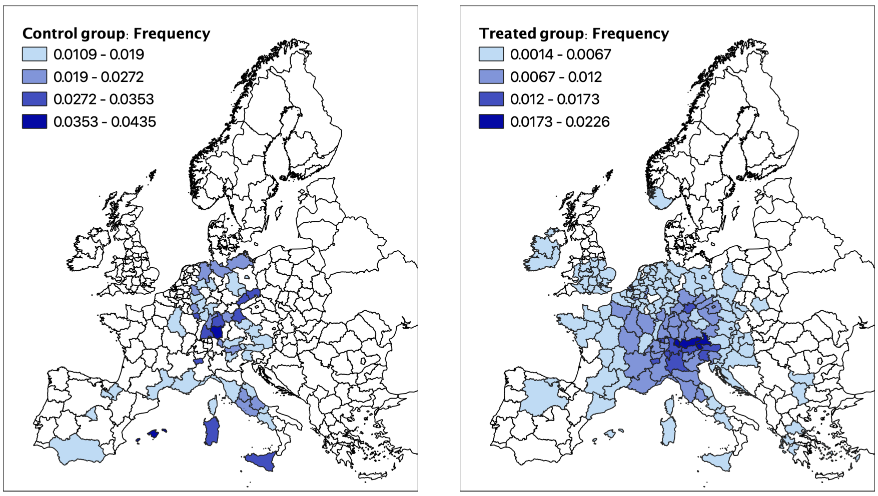

A) Control group **B) Treated group**

Notes: This figure shows the regions that are part of the block. The shading represents the frequency with which each region appears (as a part of a pair) in the control group (first panel) and treated group (second panel) in the block.

Figure A.2.17: Border effect - Industry level

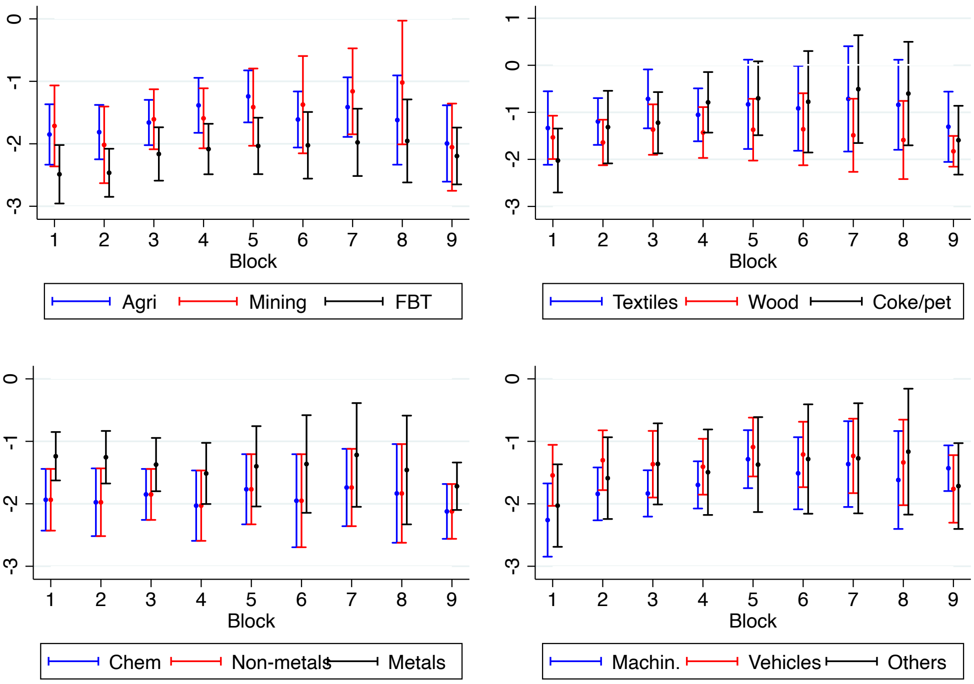

Notes: These figures show the coefficient of the dummy Border estimated with specification (2.11) in each block and industry (dot). The confidence interval for the coefficient is represented by the vertical lines.

Figure A.2.18: Participation rates across industries

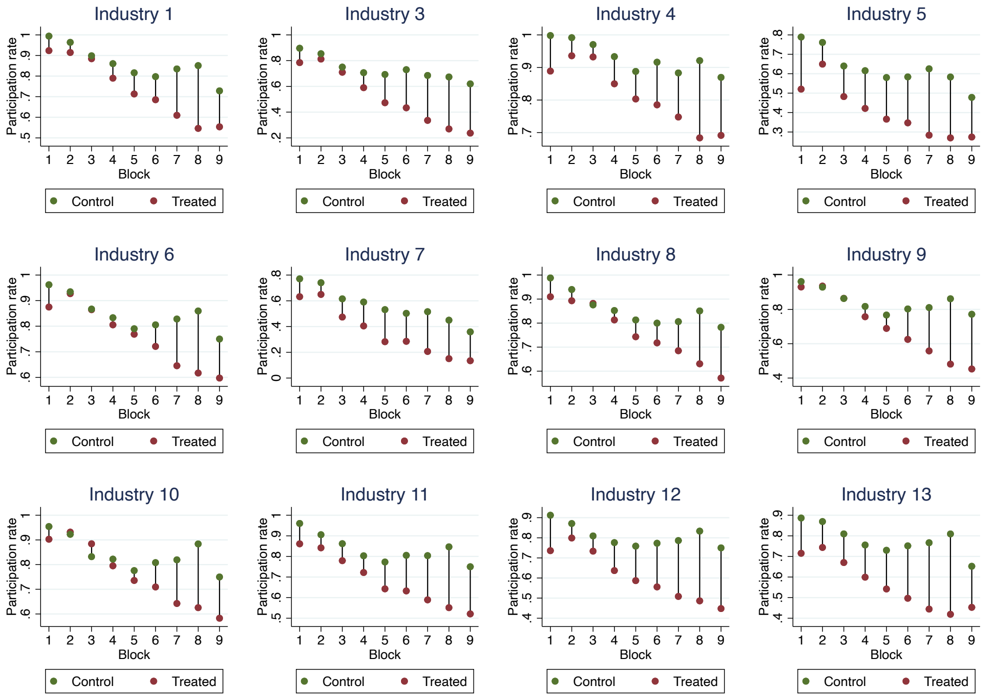

Notes: These figures show the participation rate (share of region pairs that display positive trade) in the control group, red circles, and in the treated group, green circles.

A.2.2 Additional Tables

Table A.2.1: Industries in ERFT survey

Industry	Label	Sample
1	Products of agriculture, hunting, and forestry; fish and other fishing products	1
2	Coal and lignite; crude petroleum and natural gas	0
3	Metal ores and other mining and quarrying products; peat; uranium and thorium ores	1
4	Food products, beverages and tobacco	1
5	Textiles and textile products; leather and leather products	1
6	Wood and products of wood and cork (except furniture);pulp, paper and paper; printed matter and recorded media	1
7	Coke and refined petroleum products	1
8	Chemicals, chemical products, and man-made fibers; rubber and plasticproducts; nuclear fuel	1
9	Other non-metallic mineral products	1
10	Basic metals; fabricated metal products, except machinery and equipment	1
11	Machinery and equipment n.e.c.; communication equipment; medical, precision and optical instruments; watches and clocks	1
12	Transport equipment	1
13	Furniture; other manufactured goods n.e.c.	1
14	Secondary raw materials; municipal wastes and other wastes	0
15	Mail, parcels	0
16	Equipment and material utilized inthe transport of goods	0
17	Goods moved in the course of household and office removals	0
18	Grouped goods	0
19	Unidentifiable goods: goods which for any reason cannot be identified and therefore cannot be assigned to groups 01-16	0
20	Other goods n.e.c.	0

Table A.2.2: Sample of Regions

Country	**Region**	**Label**
AT	AT11	Burgenland (AT)
AT	AT12	Niederösterreich
AT	AT13	Wien
AT	AT21	Kärnten
AT	AT22	Steiermark
AT	AT31	Oberösterreich
AT	AT32	Salzburg
AT	AT33	Tirol
AT	AT34	Vorarlberg
BE	BE10	Bruxelles-Capitale/Brussels Hoofdstedelijk Gewest
BE	BE21	Prov. Antwerpen
BE	BE22	Prov. Limburg (BE)
BE	BE23	Prov. Oost-Vlaanderen
BE	BE24	Prov. Vlaams-Brabant
BE	BE25	Prov. West-Vlaanderen
BE	BE31	Prov. Brabant Wallon
BE	BE32	Prov. Hainaut
BE	BE33	Prov. Liège
BE	BE34	Prov. Luxembourg (BE)
BE	BE35	Prov. Namur
BG	BG31	Severozapaden
BG	BG32	Severen tsentralen
BG	BG33	Severoiztochen
BG	BG34	Yugoiztochen
BG	BG41	Yugozapaden
BG	BG42	Yuzhen tsentralen
CZ	CZ01	Praha
CZ	CZ02	Střední Čechy

Continued on next page

Table A.2.2 – *Continued from previous page*

Country	Region	Label
CZ	CZ03	Jihozápad
CZ	CZ04	Severozápad
CZ	CZ05	Severovýchod
CZ	CZ06	Jihovýchod
CZ	CZ07	Střední Morava
CZ	CZ08	Moravskoslezsko
DE	DE11	Stuttgart
DE	DE12	Karlsruhe
DE	DE13	Freiburg
DE	DE14	Tübingen
DE	DE21	Oberbayern
DE	DE22	Niederbayern
DE	DE23	Oberpfalz
DE	DE24	Oberfranken
DE	DE25	Mittelfranken
DE	DE26	Unterfranken
DE	DE27	Schwaben
DE	DE30	Berlin
DE	DE40	Brandenburg
DE	DE50	Bremen
DE	DE60	Hamburg
DE	DE71	Darmstadt
DE	DE72	Gießen
DE	DE73	Kassel
DE	DE80	Mecklenburg-Vorpommern
DE	DE91	Braunschweig
DE	DE92	Hannover
DE	DE93	Lüneburg

Continued on next page

Table A.2.2 – *Continued from previous page*

Country	Region	Label
DE	DE94	Weser-Ems
DE	DEA1	Düsseldorf
DE	DEA2	Köln
DE	DEA3	Münster
DE	DEA4	Detmold
DE	DEA5	Arnsberg
DE	DEB1	Koblenz
DE	DEB2	Trier
DE	DEB3	Rheinhessen-Pfalz
DE	DEC0	Saarland
DE	DED2	Dresden
DE	DED4	Chemnitz
DE	DED5	Leipzig
DE	DEE0	Sachsen-Anhalt
DE	DEF0	Schleswig-Holstein
DE	DEG0	Thüringen
DK	DK01	Hovedstaden
DK	DK02	Sjælland
DK	DK03	Syddanmark
DK	DK04	Midtjylland
DK	DK05	Nordjylland
EL	EL30	Attiki
EL	EL41	Voreio Aigaio
EL	EL42	Notio Aigaio
EL	EL43	Kriti
EL	EL51	Anatoliki Makedonia, Thraki
EL	EL52	Kentriki Makedonia
EL	EL53	Dytiki Makedonia

Continued on next page

Table A.2.2 – *Continued from previous page*

Country	Region	Label
EL	EL54	Thessalia
EL	EL61	Ipeiros
EL	EL62	Ionia Nisia
EL	EL63	Dytiki Ellada
EL	EL64	Sterea Ellada
EL	EL65	Peloponnisos
ES	ES11	Galicia
ES	ES12	Principado de Asturias
ES	ES13	Cantabria
ES	ES21	País Vasco
ES	ES22	Comunidad Foral de Navarra
ES	ES23	La Rioja
ES	ES24	Aragón
ES	ES30	Comunidad de Madrid
ES	ES41	Castilla y León
ES	ES42	Castilla-La Mancha
ES	ES43	Extremadura
ES	ES51	Cataluña
ES	ES52	Comunidad Valenciana
ES	ES53	Illes Balears
ES	ES61	Andalucía
ES	ES62	Región de Murcia
FI	FI19	Länsi-Suomi
FI	FI18	Helsinki-Uusimaa+Etelä-Suomi
FI	FI1D	Pohjois- ja Itä-Suomi
FR	FR10	Île de France
FR	FR21	Champagne-Ardenne
FR	FR22	Picardie

Continued on next page

Table A.2.2 – *Continued from previous page*

Country	**Region**	**Label**
FR	FR23	Haute-Normandie
FR	FR24	Centre
FR	FR25	Basse-Normandie
FR	FR26	Bourgogne
FR	FR30	Nord - Pas-de-Calais
FR	FR41	Lorraine
FR	FR42	Alsace
FR	FR43	Franche-Comté
FR	FR51	Pays de la Loire
FR	FR52	Bretagne
FR	FR53	Poitou-Charentes
FR	FR61	Aquitaine
FR	FR62	Midi-Pyrénées
FR	FR63	Limousin
FR	FR71	Rhône-Alpes
FR	FR72	Auvergne
FR	FR81	Languedoc-Roussillon
FR	FR82	Provence-Alpes-Côte d'Azur
FR	FR83	Corse
HR	HR03	Jadranska Hrvatska
HR	HR04	Kontinentalna Hrvatska
HU	HU10	Közép-Magyarország
HU	HU21	Közép-Dunántúl
HU	HU22	Nyugat-Dunántúl
HU	HU23	Dél-Dunántúl
HU	HU31	Észak-Magyarország
HU	HU32	Észak-Alföld
HU	HU33	Dél-Alföld

Continued on next page

Table A.2.2 – *Continued from previous page*

Country	**Region**	**Label**
IE	IE01	Border, Midland and Western
IE	IE02	Southern and Eastern
IT	ITC1	Piemonte
IT	ITC2	Valle d'Aosta/Vallée d'Aoste
IT	ITC3	Liguria
IT	ITC4	Lombardia
IT	ITF1	Abruzzo
IT	ITF2	Molise
IT	ITF3	Campania
IT	ITF4	Puglia
IT	ITF5	Basilicata
IT	ITF6	Calabria
IT	ITG1	Sicilia
IT	ITG2	Sardegna
IT	ITH1	Provincia Autonoma di Bolzano/Bozen
IT	ITH2	Provincia Autonoma di Trento
IT	ITH3	Veneto
IT	ITH4	Friuli-Venezia Giulia
IT	ITH5	Emilia-Romagna
IT	ITI1	Toscana
IT	ITI2	Umbria
IT	ITI3	Marche
IT	ITI4	Lazio
NL	NL11	Groningen
NL	NL12	Friesland (NL)
NL	NL13	Drenthe
NL	NL21	Overijssel
NL	NL22	Gelderland

Continued on next page

Table A.2.2 – *Continued from previous page*

Country	Region	Label
NL	NL23	Flevoland
NL	NL31	Utrecht
NL	NL32	Noord-Holland
NL	NL33	Zuid-Holland
NL	NL34	Zeeland
NL	NL41	Noord-Brabant
NL	NL42	Limburg (NL)
PL	PL11	Łódzkie
PL	PL12	Mazowieckie
PL	PL21	Małopolskie
PL	PL22	Śląskie
PL	PL31	Lubelskie
PL	PL32	Podkarpackie
PL	PL33	Świętokrzyskie
PL	PL34	Podlaskie
PL	PL41	Wielkopolskie
PL	PL42	Zachodniopomorskie
PL	PL43	Lubuskie
PL	PL51	Dolnośląskie
PL	PL52	Opolskie
PL	PL61	Kujawsko-Pomorskie
PL	PL62	Warmińsko-Mazurskie
PL	PL63	Pomorskie
PT	PT11	Norte
PT	PT15	Algarve
PT	PT16	Centro (PT)
PT	PT17	Lisboa
PT	PT18	Alentejo

Continued on next page

Table A.2.2 – *Continued from previous page*

Country	**Region**	**Label**
RO	RO11	Nord-Vest
RO	RO12	Centru
RO	RO21	Nord-Est
RO	RO22	Sud-Est
RO	RO31	Sud - Muntenia
RO	RO32	Bucureşti - Ilfov
RO	RO41	Sud-Vest Oltenia
RO	RO42	Vest
SE	SE11	Stockholm
SE	SE12	Östra Mellansverige
SE	SE21	Småland med öarna
SE	SE22	Sydsverige
SE	SE23	Västsverige
SE	SE31	Norra Mellansverige
SE	SE32	Mellersta Norrland
SE	SE33	Övre Norrland
SI	SI03	Vzhodna Slovenija
SI	SI04	Zahodna Slovenija
SK	SK01	Bratislavský kraj
SK	SK02	Západné Slovensko
SK	SK03	Stredné Slovensko
SK	SK04	Východné Slovensko
UK	UKC1	Tees Valley and Durham
UK	UKC2	Northumberland and Tyne and Wear
UK	UKD1	Cumbria
UK	UKD3	Greater Manchester
UK	UKD4	Lancashire
UK	UKD6	Cheshire

Continued on next page

Table A.2.2 – *Continued from previous page*

Country	Region	Label
UK	UKD7	Merseyside
UK	UKE1	East Yorkshire and Northern Lincolnshire
UK	UKE2	North Yorkshire
UK	UKE3	South Yorkshire
UK	UKE4	West Yorkshire
UK	UKF1	Derbyshire and Nottinghamshire
UK	UKF2	Leicestershire, Rutland and Northamptonshire
UK	UKF3	Lincolnshire
UK	UKG1	Herefordshire, Worcestershire and Warwickshire
UK	UKG2	Shropshire and Staffordshire
UK	UKG3	West Midlands
UK	UKH1	East Anglia
UK	UKH2	Bedfordshire and Hertfordshire
UK	UKH3	Essex
UK	UKI1	Inner London
UK	UKI2	Outer London
UK	UKJ1	Berkshire, Buckinghamshire and Oxfordshire
UK	UKJ2	Surrey, East and West Sussex
UK	UKJ3	Hampshire and Isle of Wight
UK	UKJ4	Kent
UK	UKK1	Gloucestershire, Wiltshire and Bristol/Bath area
UK	UKK2	Dorset and Somerset
UK	UKK3	Cornwall and Isles of Scilly
UK	UKK4	Devon
UK	UKL1	West Wales and The Valleys
UK	UKL2	East Wales
UK	UKM2	Eastern Scotland
UK	UKM3	South Western Scotland

Continued on next page

Table A.2.2 – *Continued from previous page*

Country	**Region**	**Label**
UK	UKM5	North Eastern Scotland
UK	UKM6	Highlands and Islands
UK	UKN0	Northern Ireland
CH	CH01	Lake Geneva Region
CH	CH02	Espace Mittelland
CH	CH03	Northwestern Switzerland
CH	CH04	Zurich
CH	CH05	Eastern Switzerland
CH	CH06	Central Switzerland
CH	CH07	Ticino
NO	NO01	Oslo og Akershus
NO	NO02	Hedmark og Oppland
NO	NO03	Sør-Østlandet
NO	NO04	Agder og Rogaland
NO	NO05	Vestlandet
NO	NO06	Trøndelag
NO	NO07	Nord-Norge

Table A.2.3: Price regressions

	DEP.VAR: Log Price	DEP.VAR: Log Price
	(1)	(2)
$\log(\text{Dist})_{nm}$		0.451
		(0.012)
Constant	10.550	0.728
	(1.184)	(1.194)
Industry-Year FE	Yes	Yes
Origin Variables	Yes	Yes
Destination Variables	Yes	Yes
Obs.	48995	48995
R-squared	0.525	0.539

Notes: First column displays the results including only origin and destination level variables. The second column reports the results when adding the bilateral distance between origin and destination as a determinant of export prices.

Table A.2.4: Average border effect - Complete table

Dep. Var: $ln(S_{n,m})$	Block 1	Block 2	Block 3	Block 4	Block 5	Block 6	Block 7	Block 8	Block 9
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Border	-1.786	-1.721	-1.699	-1.768	-1.686	-1.796	-1.687	-1.754	-1.858
	(0.182)	(0.178)	(0.175)	(0.175)	(0.238)	(0.289)	(0.268)	(0.290)	(0.201)
Distance	-0.899	-1.378	-1.643	-0.618	-1.949	-0.532	-1.105	-1.066	-1.118
	(0.440)	(0.276)	(0.377)	(0.315)	(0.828)	(0.696)	(0.497)	(0.372)	(1.873)
Insularity	1.120	-0.861	-0.157	-0.491	-1.777	-0.913	-1.596	-1.554	-1.024
	(0.754)	(0.376)	(0.430)	(0.412)	(0.534)	(0.418)	(0.351)	(0.319)	(0.862)
Mountain Ranges	0.014	-0.137	-0.180	-0.134	-0.322	-0.088	-0.229	-0.257	-0.095
	(0.074)	(0.071)	(0.080)	(0.082)	(0.175)	(0.102)	(0.089)	(0.097)	(0.243)
River Basin	0.220	0.141	0.132	0.477	0.155	0.514	0.413	0.348	0.594
	(0.182)	(0.123)	(0.168)	(0.166)	(0.203)	(0.181)	(0.192)	(0.174)	(0.458)
Remoteness	2.236	3.236	3.339	1.335	3.412	0.803	2.086	2.167	1.356
	(0.625)	(0.783)	(0.595)	(0.606)	(1.557)	(1.219)	(0.889)	(0.833)	(2.833)
Number of Borders	7.058	6.695	7.041	10.779	11.294	11.833	9.234	8.091	0.420
	(1.756)	(1.970)	(2.034)	(1.730)	(2.064)	(2.783)	(2.792)	(3.063)	(2.944)
Constant	-52.432	-53.962	-55.214	-70.492	-79.496	-74.367	-63.052	-56.456	-4.131
	(11.534)	(12.696)	(12.979)	(10.102)	(12.606)	(15.500)	(15.239)	(16.468)	(20.347)
N	645	813	1024	1364	968	1267	2011	2948	637
R^2	.572	.533	.501	.47	.375	.388	.31	.285	.299

Notes: Standard errors clustered at the country-pair level, are in parentheses. Dependent variable is the (log) normalized market share of n in m. *Border* is a dummy for international border. *Number of Borders* is the average of the share of international borders that are faced by n and m. *Distance* is (log) bilateral distance between origin and destination in kilometers, *Insularity* takes value 1 if one of the regions is an island. *Mountain Ranges* is the highest difference in elevation between two regions in metres (difference between highest point and lowest point, in logs), *River Basin* takes value 1 if the region pair shares a river basin and *Remoteness* is the log of the average remoteness of the origin and the destination regions.

Table A.2.5: Average border effect - No number of borders

Dep. Var: $ln(S_{n,m})$	Block 1	Block 2	Block 3	Block 4	Block 5	Block 6	Block 7	Block 8	Block 9
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Border	-1.478	-1.382	-1.336	-1.166	-1.092	-1.162	-1.215	-1.350	-1.841
	(0.193)	(0.208)	(0.174)	(0.202)	(0.210)	(0.216)	(0.189)	(0.199)	(0.209)
Distance	-0.807	-1.508	-1.917	-0.660	-2.161	-0.843	-1.289	-1.190	-1.098
	(0.458)	(0.327)	(0.439)	(0.349)	(0.898)	(0.771)	(0.523)	(0.361)	(1.843)
Insularity	1.298	-0.833	-0.259	-0.653	-2.202	-1.419	-1.969	-1.836	-1.020
	(0.520)	(0.379)	(0.380)	(0.417)	(0.588)	(0.490)	(0.398)	(0.303)	(0.853)
Mountain Ranges	-0.002	-0.110	-0.137	-0.079	-0.269	-0.079	-0.211	-0.242	-0.090
	(0.091)	(0.089)	(0.091)	(0.106)	(0.193)	(0.106)	(0.093)	(0.098)	(0.232)
River Basin	0.471	0.300	0.212	0.732	0.315	0.635	0.470	0.409	0.606
	(0.219)	(0.171)	(0.212)	(0.232)	(0.257)	(0.210)	(0.198)	(0.194)	(0.431)
Remoteness	2.795	3.886	4.087	2.213	4.674	2.384	3.327	3.228	1.368
	(0.735)	(1.005)	(0.797)	(0.776)	(1.691)	(1.211)	(0.913)	(0.720)	(2.838)
Constant	-13.122	-16.590	-15.672	-10.163	-17.628	-10.528	-13.645	-13.189	-1.784
	(3.456)	(5.477)	(3.780)	(3.555)	(5.732)	(3.785)	(3.384)	(3.466)	(8.180)
N	645	813	1024	1364	968	1267	2011	2948	637
R^2	.499	.473	.454	.384	.302	.314	.276	.262	.299

Notes: Standard errors clustered at the country-pair level, are in parentheses. Dependent variable is the (log) normalized market share of n in m. *Border* is a dummy for international border. *Distance* is (log) bilateral distance between origin and destination in kilometers, *Insularity* takes value 1 if one of the regions is an island. *Mountain Ranges* is the highest difference in elevation between two regions in metres (difference between highest point and lowest point, in logs), *River Basin* takes value 1 if the region pair shares a river basin and *Remoteness* is the log of the average remoteness of the origin and the destination regions.

Table A.2.6: Summary statistics of covariates by block: Conditional on Border in 1910=0

	(1)	(2)	(3)	(4)	(5)	(6)
	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd	mean/sd
Distance	271.968	371.468	455.596	535.897	657.689	665.057
	97.15	113.65	131.67	170.34	217.44	274.37
Insularity	0.072	0.098	0.086	0.040	0.050	0.056
	0.26	0.30	0.28	0.20	0.22	0.23
Mountain Ranges	5.532	5.873	6.114	6.157	6.344	6.552
	0.98	0.90	0.81	0.78	0.81	0.80
River Basin	0.203	0.170	0.174	0.229	0.275	0.689
	0.40	0.38	0.38	0.42	0.45	0.46
Remoteness	1067.843	1028.199	1003.654	992.166	1001.298	1036.640
	235.34	206.65	183.78	164.26	142.07	144.14
Estimated propensity score	0.170	0.311	0.439	0.559	0.685	0.814
	0.04	0.04	0.04	0.04	0.04	0.04
N	775	552	466	375	298	161

Notes: This table reports the mean and standard deviation of each geographical covariate and the propensity score in each block. *Distance* is bilateral distance between origin and destination in kilometers, *Insularity* takes value 1 if one of the regions is an island. *Mountain Ranges* is the highest difference in elevation between two regions in metres (difference between highest point and lowest point), *River Basin* takes value 1 if the region pair shares a river basin and *Remoteness* is the average remoteness of the origin and the destination regions.

Table A.2.7: Balancing test of covariates by block: Conditional on Border in 1910=0

	(1)	(2)	(3)	(4)	(5)	(6)
Distance	0.0219	-0.0256	0.0362	-0.0197	-0.0510	-0.141
	(0.0388)	(0.0339)	(0.0321)	(0.0382)	(0.0513)	(0.0779)
Insularity	-0.0862	-0.114	0.000715	0.0277	0.211	0.333
	(0.0240)	(0.0273)	(0.0263)	(0.0202)	(0.0272)	(0.0410)
Mountain Ranges	-0.353	-0.227	-0.0558	0.162	0.701	0.692
	(0.0910)	(0.0836)	(0.0755)	(0.0806)	(0.103)	(0.160)
River Basin	-0.00557	0.0327	-0.0474	0.0168	0.0455	-0.0274
	(0.0376)	(0.0351)	(0.0355)	(0.0435)	(0.0609)	(0.0982)
Remoteness	0.0349	-0.0208	-0.00301	0.00436	0.0235	-0.104
	(0.0188)	(0.0172)	(0.0154)	(0.0150)	(0.0170)	(0.0252)
N	775	552	466	375	298	161

Notes: This table reports the difference in means between treated and control region pairs for each geographical covariate by block (defined as control minus treated). Standard errors in parenthesis. *Distance* is bilateral distance between origin and destination in kilometers, *Insularity* takes value 1 if one of the regions is an island. *Mountain Ranges* is the highest difference in elevation between two regions in metres (difference between highest point and lowest point), *River Basin* takes value 1 if the region pair shares a river basin and *Remoteness* is the average remoteness of the origin and the destination regions.

A.2.3 Construction of European regional trade dataset

in this section we explain the methodology we follow to construct the matrix of regional trade flows in Europe. First, we explain the data sets used for the price imputation procedure. Second, we provide additional details about how we clean and use the European Road Freight dataset (ERFT).

Regional price data

The subsample of region to country level trade data is collected individually for our subset of four countries:

France The French Douane administration provides international trade data for the different Regions and Departements in France. The data is available quarterly for the years 2011 and 2014 at the industry level (4 digits of disaggregation of CPA4) for the different origin/destination countries. The trade flows are collected in value and weight, both imports and exports.[1] We use a 2 digits industrial disaggregation (22 industries).

Germany The German agency of statistics, Destatis, provides Foreign trade data for the 16 German states (Bundeslander). The data is available monthly for the years 2008 to November 2016 at the industry level (1, 2 or 3 digits of aggregation) for the different origin/destination countries. The trade flows are collected in value and weight (Tons). For this paper we use annual data for the years 2011 to 2014, at a 2 digits level of disaggregation (30 industries).[2]

Spain The Spanish secretary of commerce provides Foreign trade data for the 17 Spanish regions (Comunidades Autonomas). The data is available monthly for the years 1995 to 2015 at different industry levels for the different origin/destination countries. The trade flows are collected in value, not weight. For this paper we use annual data for the years 2011 to 2014, at a 2 digits level of disaggregation (22 industries).[3]

United Kingdom The UK Customs department provides Foreign trade data for the 12 regions in the UK. The data is available monthly for the years 2009 to 2016 at different industry levels (several digits available) for the different origin/destination countries. The trade flows are collected in

[1] The data can be accessed at `http://lekiosque.finances.gouv.fr/portail_default.asp`.

[2] The data can be accessed at `https://www-genesis.destatis.de/genesis/online/data`.

[3] The data can be accessed at `http://datacomex.comercio.es/principal_comex_es.aspx`.

value and weight. For this paper we use annual data for the years 2011 to 2014, at a 2 digit level of disaggregation.[4]

Table A.2.8: Foreign Trade Sample

Country	**Unit**	**Freq**	**Year**	**Industries**	**Unit**
Spain	NUTS2	Monthly	2011-2014	22, 99	€, kg
Germany	NUTS1	Monthly	2011-2014	30, 211	€, kg
France	NUTS3	Trimester	2011-2014	22 ,>200	€, kg
UK	NUTS1	Quarterly	2011-2014	67	£, kg

We aggregate each dataset to a 20 industry NST 2007 classification (European classification system for transport statistics), which is the classification used in the European Road Freight Transport Survey. This subsample of 58 regions allows us to observe 2,688 region to country trade flows (region-country pairs) each year.

Variables for price imputation and robustness checks

We put together an extensive database of economic and geographic characteristics at the regional and country to use as determinants of price levels across regions. Our preferred specification is to pool all time periods and industries in the following regression:

$$\ln P_{nm}^{it} = \eta_n^t X_n^{t-1} + \pi_m^t Z_m^{t-1} + \beta d_{nm} + \phi^{it} + e_{nm}^{it},$$

where P_{nm}^{it} is the unit price of exports of industry i shipped from origin n to destination m in year t. The price of exports is calculated as the ratio between the value of exports and the weight of exports for each industry, origin, destination and year. Table A.2.9 reports the complete list of variables that we include as controls.[5] In addition, we also compute the geodesic distance

[4] The data can be access at: Statistical department of the United Kingdom government.

[5] EuroRegional Map: https://eurogeographics.org/products-and-services/euroregionalmap/

between the centroid of the origin and the destination region, and we use it as a proxy for bilateral distance d_{nm}.

Table A.2.9: Explanatory Variables for Price regressions

Label	Included	Level	Source
log(Pop Dens)	or/dest	NUTS2, year	Eurostat
log(GDP pc)	or/dest	NUTS2, year	Eurostat
log(Life Exp.)	or/dest	NUTS2, year	Eurostat
log(Total Emp.)	or/dest	NUTS2, year	Eurostat
Manuf. Sh. of Emp.	or/dest	NUTS2, year	Eurostat
Low Tech. Sh. of Emp.	or/dest	NUTS2, year	Eurostat
Edu (None) Sh.	or/dest	NUTS2, year 2011	2011 census
Edu (ISEC3) Sh.	or/dest	NUTS2, year 2011	2011 census
Edu (ISEC6) Sh.	or/dest	NUTS2, year 2011	2011 census
Ind Agri. Sh.	or/dest	NUTS2, year 2011	2011 census
Ind Manu. Sh.	or/dest	NUTS2, year 2011	2011 census
Ind. Prof/Science Sh.	or/dest	NUTS2, year 2011	2011 census
Ind. Fin. Sh.	or/dest	NUTS2, year 2011	2011 census
Ind. Pub. Sh	or/dest	NUTS2, year 2011	2011 census
Birth (Other EU) Sh.	or/dest	NUTS2, year 2011	2011 census
Birth (Non-EU) Sh.	or/dest	NUTS2, year 2011	2011 census
log(Heating h)	or/dest	NUTS2, year	Eurostat
log(av sun h)	or/dest	NUTS2	PVGIS 5 solar irradiation
log(max_sun h)	or/dest	NUTS2	PVGIS 5 solar irradiation
log(distRiver)	or/dest	NUTS2	EuroRegional map
log(distCoast)	or/dest	NUTS2	EuroRegional map

To test the accuracy of our predicted prices, we collect data of country to country trade flows at the year-industry level from Eurostat dataset COMEXT. Comext is Eurostat's reference database for detailed statistics on international trade in goods. It provides information about the value and quantity of the trade transaction, allowing us to compute the price per kilo of exports. We download the data for the years in our sample, 2011-2017, from the website: `http://epp.eurostat.ec.europa.eu/newxtweb/`.

European Road Freight Transport survey

The European Road Freight Transport survey microdata is a database collected by Eurostat in order to understand the magnitude of the shipment of goods across Europe. The ERFT survey covers 27 EU countries (except Malta) and EFTA countries (except Iceland). Each member state collects statistics on the carriage of goods by road by means of any road freight vehicle from a representative sample of road vehicles collected from the national vehicle registry. In case such a registry is not available, the sample will be selected either from the registry of licensed road haulage operators or the registry of persons licensed to operate such vehicles. In particular, Eurostat provides three interlinked datasets that contain the micro data at the vehicle, journey and goods level.

The Vehicle dataset (Dataset A1) records characteristics of each individual road vehicle and besides identifying each respondent vehicle contains information such as the age, axle configuration, unladen weight, total permissible weight and total kilometers performed during the survey.

The Journey dataset (Dataset A2) contains information about specific journeys performed by a vehicle identified in the A1 dataset. Each journey is assigned a journey identifier and can be linked to the corresponding vehicle in the A1 dataset that performs it. Journey related variables include gross weight of goods transported, place of loading and unloading (reported at a NUTS 2 level of disaggregation), actual distance travelled, tonne-km effected, degree of loading in terms of total volume and countries crossed in transit during each journey. Notably, survey distinguishes different journey types based on their laden/unladen status and the number of distinct transport operations involved. As a result four main journey types are identified: Laden-Involving one single transport operation, laden-Involving multiple transport operations, laden-collection/distribution and unladen. Journeys that involve 5 or more distinct locations are considered to be of collection/distribution type.

The goods dataset (Dataset A3) each journey is broken down to represent specific shipments of goods between two geographical units. Each goods' transfer between any two geographical units is identified and linked to the specific journey it is part of. Journeys that involve either multiple destinations for loading/unloading and/or different types of goods are further broken down in the goods dataset (Dataset A3). Each observation in Dataset A3 represents a flow of one type of good between two specific geographical units.

Region border changes Throughout the paper we use the classification NUTS2013 for most regions for consistency. In cases for which there was a change, a region split in more regions, from NUTS2010 to NUTS2013 we use the aggregated NUTS2010. This is the case for regions FI1B and FI1C (NUTS2013) in Finland, which we aggregate for all years in our data and corresponds to FI18 (NUTS2010). For London area regions UKI3, UKI4, UKI5, UKI6 and UKI7 (NUTS2013) we use the aggregated UKI1 (UKI3 + UKI4) and UKI2 (UKI5 + UKI6 + UKI7) NUTS2010 regions.

Cleaning data To create our matrix of weights of goods shipped between each region-pair we merge the good-level dataset (A3) for the years 2011 to 2017. We drop "unladen" journeys. We also drop "distribution" journeys, since these are journeys that involve five or more stops in distinct locations considered to be of collection/distribution nature. These are more likely associated with distribution or logistics than with trade. We then normalise the region identifiers to the 2013 NUTS version, since there are some regions that change name between 2011 and 2017. Finally, we apply the weights provided by Eurostat to each shipment to account for under-sampling of some journeys.

We then aggregate the value traded across all industries by each region pair by adding up the value traded in all industries for each region-pair in each year. Finally, to construct our region-pair level dataset we take the average of the value traded by each region pair (n,m) across all years 2011-2017.

A.2.4 Additional data sources

Construction of geographical variables

1. *Distance*: We construct bilateral distance by calculating the length of the curve linking the central point of the origin region (centroid) and the central point of the destination region, in kilometers. We use a curve since we take into account the curvature of earth's surface. We compute the centroid as the center point of the polygon of the area of the region, using the software ArcGIS.

2. *Insularity*: Dummy variable taking value one if there is the need to cross a sea to reach from one region to the other, and zero otherwise.

3. *Mountain ranges*: Largest altitude difference between two regions, computed as the difference between the highest altitude point and the lowest altitude point along the straight line that joins the centre the origin region (centroid) and the centre of the destination region. To compute this maximum difference in altitude we use a topographic layer of Europe. We compute the straight line segment that links each possible region-pair (centroid to centroid). We then compute the altitude at different intervals along the line (computed using the cells of the altitude raster) and keep the highest and the lowest points. Finally, we take the difference between the highest and the lowest point.

4. *River basin*. Dummy variable taking value 1 if both regions belong to the same river basin. We consider the largest rivers in Europe. A map of the areas covered by each river basin is shown in figure A.2.19. We consider the major European rivers: Danube, Douro/Duero, Elbe, Ebro, Glomma, Garonne, Gota Alv, Loire, Meuse/Maas, Maritsa, Oder, Ouse, Po, Rhein, Rhone, Seine, Severn, Tejo/Tajo, Thames, Tiber, Trent, Weser, Vistula.

5. *Remoteness*. We calculate the remoteness of a region as the sum of the

bilateral distance from that region to every other region in the sample. Then, we calculate the remoteness of a pair as the average remoteness of both regions.

6. *Number of borders* We sum the number of borders of the origin region and the number of border of the destination region, and we take the log of the sum. We compute the number of borders of the origin as the number of regions in the sample minus one (the border of the origin with itself) minus the number of regions in the country that the origin region belongs to (regions with which the origin region does not have a border). We do the same for the destination.

Figure A.2.19: Regions that share a river basin

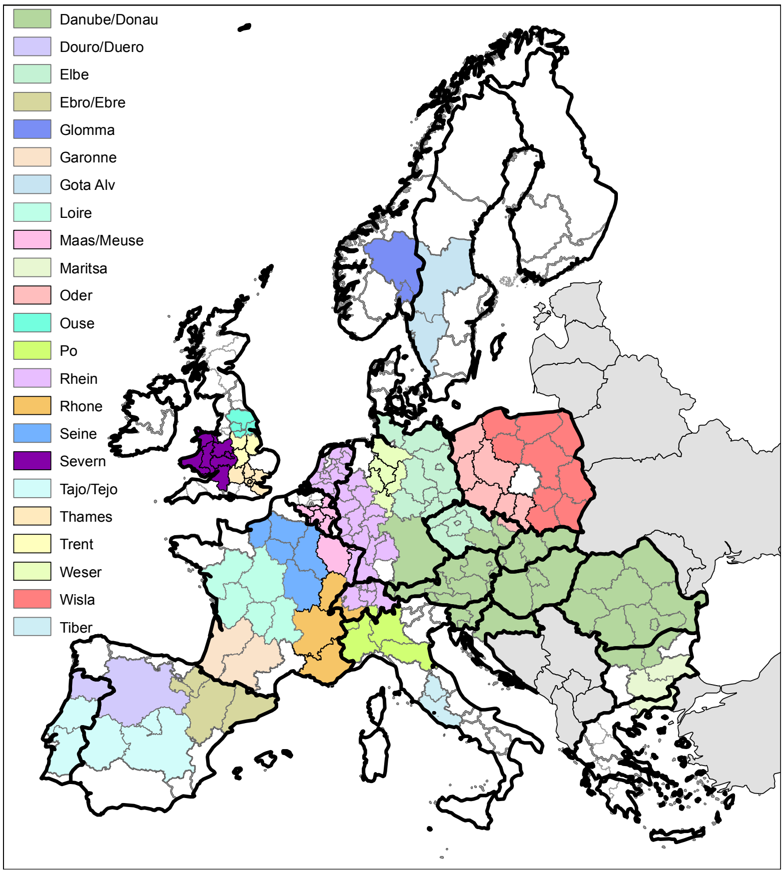

Notes: This figure shows the different river basins that we consider, represented by different colors.

Collection of historical borders

We thank Matteo Cervellati, Sara Lazzaroni, Giovanni Prarolo and Paolo Vanin for kindly sharing their digitised data of historical borders in Europe from their paper Cervellati et al. (2019). We use the shapefile provided by the authors to identify borders in 1910 between our 269 regions.

A.3 Appendix: Chapter 3

A.3.1 Additional Figures and Tables

Table A.3.1: Border effects for country pairs

Country	**Border (Mean)**	**Border (SD)**	**Highest (1)**	**Highest (2)**	**Highest (3)**	**Lowest (1)**	**Lowest (2)**	**Lowest (3)**
AT	-2.93	0.88	-5.24 (FI)	-4.68 (IE)	-3.84 (NO)	-1.45 (DE)	-1.55 (SI)	-1.96 (SK)
BE	-2.71	1.03	-5.30 (FI)	-4.57 (IE)	-4.32 (NO)	-1.35 (FR)	-1.53 (NL)	-1.68 (CZ)
BG	-3.18	1.25	-7.84 (NO)	-4.57 (HR)	-3.93 (PT)	-1.21 (IE)	-1.67 (UK)	-2.28 (EL)
CH	-3.50	0.77	-5.18 (IE)	-4.90 (SI)	-4.88 (HR)	-2.14 (DE)	-2.77 (SK)	-2.77 (BE)
CZ	-2.58	0.76	-4.07 (IE)	-3.82 (FI)	-3.65 (HR)	-0.84 (SK)	-1.48 (DE)	-1.68 (BE)
DE	-2.53	1.07	-5.17 (FI)	-4.99 (IE)	-3.92 (NO)	-1.43 (SK)	-1.45 (AT)	-1.48 (CZ)
DK	-3.18	0.68	-4.61 (FI)	-4.44 (UK)	-4.42 (IE)	-2.28 (BG)	-2.29 (PL)	-2.33 (SE)
EL	-2.90	1.15	-5.19 (HR)	-4.05 (SE)	-3.83 (ES)	-1.80 (NO)	-2.28 (BG)	-2.36 (UK)
ES	-3.17	1.10	-5.60 (IE)	-5.18 (FI)	-4.82 (NO)	-1.43 (PT)	-1.70 (FR)	-2.01 (BE)
FI	-4.47	1.39	-6.50 (PT)	-6.39 (IE)	-6.39 (UK)	-2.60 (BG)	-2.99 (SE)	-3.82 (CZ)
FR	-3.10	1.12	-5.36 (IE)	-5.06 (FI)	-4.89 (NO)	-1.35 (BE)	-1.70 (ES)	-1.93 (SI)
HR	-4.11	0.91	-5.87 (IE)	-5.59 (PT)	-5.27 (FI)	-1.91 (SI)	-3.02 (AT)	-3.13 (HU)
HU	-2.94	0.97	-5.09 (IE)	-4.69 (NO)	-4.29 (FI)	-1.50 (SI)	-1.59 (DE)	-1.72 (SK)
IE	-4.57	1.47	-6.39 (FI)	-6.19 (PT)	-5.87 (HR)	-1.21 (BG)	-3.31 (UK)	-4.02 (NL)
IT	-2.98	0.84	-4.93 (IE)	-4.68 (FI)	-4.48 (NO)	-1.88 (SI)	-2.04 (SK)	-2.13 (PL)
NL	-2.84	0.76	-4.89 (FI)	-4.02 (IE)	-3.93 (NO)	-1.53 (BE)	-1.72 (DE)	-2.08 (PL)
NO	-4.07	1.17	-7.84 (BG)	-4.97 (IE)	-4.89 (FR)	-1.80 (EL)	-2.13 (SE)	-2.72 (DK)
PL	-2.73	0.69	-4.13 (HR)	-4.07 (IE)	-3.94 (FI)	-1.58 (DE)	-1.76 (BE)	-2.08 (NL)
PT	-3.38	1.24	-6.50 (FI)	-6.19 (IE)	-5.59 (HR)	-1.43 (ES)	-2.11 (SK)	-2.20 (FR)
RO	-3.25	0.66	-4.70 (FI)	-4.32 (HR)	-4.27 (NO)	-2.28 (BE)	-2.43 (HU)	-2.48 (UK)
SE	-3.65	0.79	-5.39 (IE)	-4.97 (UK)	-4.64 (FR)	-2.13 (NO)	-2.33 (DK)	-2.86 (PL)
SI	-2.86	1.11	-5.01 (IE)	-4.90 (CH)	-4.30 (NO)	-1.31 (SK)	-1.50 (HU)	-1.54 (DE)
SK	-2.60	1.00	-5.64 (IE)	-3.91 (FI)	-3.51 (HR)	-0.84 (CZ)	-1.31 (SI)	-1.43 (DE)
UK	-3.37	1.03	-6.39 (FI)	-4.97 (SE)	-4.78 (NO)	-1.67 (BG)	-2.29 (PL)	-2.36 (EL)

Table A.3.2: Gravity: PPML Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
	S_{nm}	S_{nm}	S_{nm}	S_{nm}	S_{nm}	S_{nm}
Border Effect	-1.808***	-2.002***				
	(0.123)	(0.108)				
Border / common language / common currency dummy			-1.724***	-1.725***		
			(0.214)	(0.182)		
Border / common language / different currency dummy			-1.855***	-1.833***		
			(0.146)	(0.151)		
Border / different language / common currency dummy			-1.719***	-1.995***		
			(0.148)	(0.147)		
Border / different language / different currency dummy			-1.848***	-2.096***		
			(0.145)	(0.127)		
Distance (constant-elasticity)	-1.412***		-1.410***		-1.473***	
	(0.0644)		(0.0655)		(0.0708)	
Origin FE	Yes	Yes	Yes	Yes	Yes	Yes
Dest FE	Yes	Yes	Yes	Yes	Yes	Yes
Distance (variable-elasticity)	No	Yes	No	Yes	No	Yes
Border dummies for each country pair	No	No	No	No	Yes	Yes
Observations	46505	46505	46505	46505	46505	46505
R^2	0.975	0.977	0.975	0.977	0.975	0.977

Standard errors in parentheses. Distance is measured as the (log) of average kilometers travelled.
* $p < .1$, ** $p < .05$, *** $p < .01$

Table A.3.3: Gravity: PPML Regressions

	(1)	(2)	(3)	(4)	(5)	(6)
	S_{nm}	S_{nm}	S_{nm}	S_{nm}	S_{nm}	S_{nm}
Border Effect	-2.187***	-1.871***				
	(0.119)	(0.120)				
Border / common language / common currency dummy			-1.713***	-1.505***		
			(0.142)	(0.128)		
Border / common language / different currency dummy			-1.840***	-1.704***		
			(0.145)	(0.141)		
Border / different language / common currency dummy			-2.136***	-1.829***		
			(0.141)	(0.142)		
Border / different language / different currency dummy			-2.317***	-2.002***		
			(0.135)	(0.134)		
Home Bias	1.475***	2.128***	1.508***	2.143***	1.486***	2.122***
	(0.414)	(0.522)	(0.418)	(0.526)	(0.480)	(0.555)
Distance (constant-elasticity)	-0.783***		-0.762***		-0.776***	
	(0.141)		(0.145)		(0.183)	
Origin FE	Yes	Yes	Yes	Yes	Yes	Yes
Dest FE	Yes	Yes	Yes	Yes	Yes	Yes
Distance (variable-elasticity)	No	Yes	No	Yes	No	Yes
Border dummies for each country pair	No	No	No	No	Yes	Yes
Observations	46505	46505	46505	46505	46505	46505
R^2	0.988	0.991	0.988	0.991	0.988	0.991

Standard errors in parentheses. Distance is measured as the (log) of average kilometers travelled.

* $p < .1$, ** $p < .05$, *** $p < .01$

Table A.3.4: Home Bias: Determinants - by Industry

	Agri	Mining	FBT	Textiles	Wood	Coke/Pet
	(1)	(2)	(3)	(4)	(5)	(6)
	Home	Home	Home	Home	Home	Home
Distance	-0.271	-0.187	0.199	-0.762*	-0.175	-0.00967
	(0.196)	(0.233)	(0.203)	(0.448)	(0.251)	(0.212)
Log(European Remoteness)	0.184	0.489	1.471***	2.331**	1.351**	0.551
	(0.454)	(0.509)	(0.468)	(0.920)	(0.579)	(0.499)
Island Region	1.784***	0.581	1.388***	1.869*	2.310***	0.554
	(0.289)	(0.510)	(0.353)	(1.023)	(0.440)	(0.363)
Mountain Region	0.336***	0.0968	0.242**	0.300	0.214*	0.0826
	(0.103)	(0.100)	(0.0964)	(0.187)	(0.111)	(0.101)
Major Port Region	-0.251**	0.0190	-0.0792	0.0156	-0.0859	-0.250**
	(0.0974)	(0.132)	(0.105)	(0.252)	(0.129)	(0.111)
Motorway Density	-2.032	1.010	-4.170***	-11.90***	-5.245***	-4.121**
	(1.608)	(1.494)	(1.420)	(3.126)	(1.842)	(1.617)
Log(Population)	-0.606***	-0.786***	-0.553***	-0.889***	-0.549***	-0.752***
	(0.0662)	(0.0814)	(0.0750)	(0.130)	(0.0736)	(0.0659)
Share of Emp. (Manuf.)	-4.761**	-7.716***	-3.778**	-11.24***	-10.27***	-4.877**
	(1.942)	(2.152)	(1.881)	(3.769)	(2.387)	(2.067)
Share of Emp. (Public)	4.018	-2.725	-5.189	-6.696	1.462	2.999
	(5.205)	(5.219)	(5.714)	(9.221)	(6.176)	(4.713)
Sh. Secondary or tertiary educ	-2.094**	-1.605	-0.259	-2.029	-3.105***	-0.580
	(1.011)	(1.029)	(0.994)	(1.985)	(1.126)	(1.106)
Share Migrant Pop.	-0.639	-2.321**	-0.187	-3.550**	-0.118	0.0618
	(0.824)	(0.952)	(0.781)	(1.698)	(0.919)	(0.960)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	265	265	265	254	265	265
R^2	0.838	0.827	0.843	0.625	0.866	0.817

Standard errors in parentheses. Distance is measured as the (log) of average kilometers travelled.
* $p < .1$, ** $p < .05$, *** $p < .01$

Table A.3.5: Home Bias: Determinants - by Industry (cont.)

	Chem	Non-Metal	Metal	Mach.	Vehicles	Other
	(1)	(2)	(3)	(4)	(5)	(6)
	Home	Home	Home	Home	Home	Home
Distance	0.478*	0.0678	0.509**	-0.287	-0.300	-0.0701
	(0.264)	(0.214)	(0.243)	(0.201)	(0.361)	(0.317)
Log(European Remoteness)	2.130***	1.176**	1.904***	2.029***	3.167***	2.991***
	(0.554)	(0.464)	(0.522)	(0.508)	(0.754)	(0.859)
Island Region	2.548***	0.264	2.379***	1.238***	0.946*	2.164***
	(0.665)	(0.377)	(0.753)	(0.264)	(0.560)	(0.691)
Mountain Region	0.226	0.129	0.0782	0.260***	0.180	0.260**
	(0.140)	(0.0928)	(0.141)	(0.0918)	(0.139)	(0.109)
Major Port Region	-0.197	-0.0357	-0.0174	-0.157	0.0162	-0.0142
	(0.150)	(0.105)	(0.134)	(0.147)	(0.210)	(0.216)
Motorway Density	-5.406***	1.304	-5.518***	-8.651***	-10.75***	-7.042***
	(2.009)	(1.664)	(1.849)	(1.732)	(3.232)	(2.164)
Log(Population)	-0.901***	-0.838***	-0.791***	-0.785***	-0.887***	-0.599***
	(0.102)	(0.0715)	(0.0990)	(0.0689)	(0.111)	(0.111)
Share of Emp. (Manuf.)	-6.680***	-5.895***	-11.43***	-11.76***	-13.94***	-5.466**
	(2.253)	(1.923)	(2.171)	(2.127)	(2.820)	(2.679)
Share of Emp. (Public)	-9.292*	-2.286	-4.960	-6.169	0.836	-9.645
	(5.606)	(4.283)	(5.250)	(4.832)	(6.817)	(6.522)
Sh. Secondary or tertiary educ	0.467	-1.795*	-0.675	-1.950**	-1.659	-1.912
	(1.064)	(0.966)	(1.021)	(0.927)	(1.364)	(1.364)
Share Migrant Pop.	2.693**	-1.220	2.207**	-2.229**	-0.767	-0.312
	(1.283)	(0.963)	(1.066)	(0.947)	(1.608)	(1.055)
Country FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	264	265	264	263	261	258
R^2	0.836	0.852	0.821	0.854	0.809	0.787

Standard errors in parentheses. Distance is measured as the (log) of average kilometers travelled.

* $p < .1$, ** $p < .05$, *** $p < .01$