Price Dispersion and Market Segmentation: Evidence from Bottled Water[*]

Joris Hoste[†]
University of Cambridge

September 12, 2024

[PRELIMINARY AND INCOMPLETE, PLEASE DO NOT CIRCULATE]

Abstract

It is well known that substantial spatial price dispersion persists between countries of the European Single Market (ESM). Whether this spatial price dispersion simply reflects destination-specific marginal costs or is driven by non-trade-policy barriers that fragment European markets remains an open question. This paper leverages technological and institutional features of the European bottled water industry to estimate the level of non-trade-policy barriers and their effect on equilibrium price dispersion and consumer welfare. Using a partial equilibrium model of the bottled water industry, I estimate non-trade-policy barriers between European countries of the (ESM) that are equivalent to a 20% import tariff. Relative to an integrated benchmark economy, non-trade-policy barriers increase equilibrium international spatial price dispersion by 5% and reduce consumer welfare equivalent to a 10% tax on bottled water consumption.

JEL codes: D43, D61, F12, F15, L13 and R2
Keywords: Geographic market integration, trade frictions and misallocation

---

[*] I thank Filip Abraham, Meredith Crowley, Nestor Duch-Brown, Christos Genakos, Joep Konings, Thierry Mayer, Hylke Vandenbussche, Gonzague Vanoorenberghe, Jakob Vanschoonbeek and Frank Verboven for helpful discussions. I would like to thank the seminar participants at the Cambridge Trade Workshop for their helpful comments. I gratefully acknowledge financial support from The Research Foundation - Flanders (FWO) through fellowship 1169722N, the CTIP ESRC research grant, the Janeway Institute and the Keynes Fund.

[†] Electronic Adress: `jh2524@cam.ac.uk`

Contents

1 Introduction

Spatial price dispersion is commonplace both within and between countries (e.g. Handbury (2021); Fontaine et al. (2020)). Notably, price dispersion that occurs discontinuously at market boundaries is often attributed to positive trade frictions that segment markets geographically (e.g. Gopinath et al. (2011); Beck et al. (2020)). As greater goods market integration has been cited to reduce spatial misallocation (e.g. Hornbeck & Rotemberg (2024)) and increase innovation (Andersson et al. (2023)), understanding whether markets are segmented and the size of trade frictions is crucial to evaluate the gains of market integration.

Conceptually, however, whether spatial price dispersion reflects geographic market segmentation and should be reduced crucially depends on its sources. If price differences between markets only reflect differences in destination-specific marginal costs, spatial price dispersion may arise in equilibrium even when markets are geographically integrated and will be efficient (Goldberg & Knetter (1997)). When spatial price dispersion surpasses cost differences, e.g. due to trade frictions or geographic price discrimination, goods markets are geographically segmented and the gains to reducing price dispersion depend on the size of trade friction and the entry-enhancing effect of price discrimination.

Isolating trade frictions and geographic market segmentation from other sources of price dispersion is empirically challenging for two reasons. First, when firms price discriminate in space, only variation in destination-specific marginal costs can uncover trade frictions. Unfortunately, destination-specific marginal costs are almost always unobserved. This has forced researchers to rely on restrictive assumptions on final demand and market structure to map variation in prices or trade flows to differences in destination-specific marginal costs (e.g. McCallum (1995); Anderson & Wincoop (2003); Gopinath et al. (2011); Santamaria et al. (2020)). Second, in a world dominated by multinationals, knowing the spatial structure of supply chains is crucial. This is because products sold in adjacent markets potentially pass through different value chains which might operate at a different scale of production, use different local input prices and are subject to different levels of market power along the value chain. All these aspects could lead to differences in destination-specific marginal costs. Without information on the spatial structure of value chains, restrictions on destination-specific market structure and production technology are required to disentangle trade frictions from differences in destination-specific marginal costs.

In this paper, I estimate the level of trade frictions and quantify their effect on spatial price dispersion and consumer welfare in the European bottled water industry. The EU bottled water industry provides an intriguing setting for three reasons. First, I estimate destination-specific marginal costs by combining a spatially disaggregated dataset of consumer prices and consumption with a model of demand and supply of bottled water. I rely on cross-country household-level scanner data where the location of residence is used as a proxy for the location of purchase which is plausible for a low value-to-weight product like bottled water. I leverage this fine spatial distribution of prices and quantities within an empirical model of the industry to back out destination-specific marginal without stringent assumptions on preferences, market structure and technology.

Second, I make use of regulatory and technological features of the industry bottled water industry to account for spatial variation in supply chains without stringent assumptions on destination-specific market structure and technology. By EU law, bottled water producers must specify the name of the water source on the product label. I use this as a basis to hand-collect data on production locations of bottled water varieties sold across EU countries. From a technological point of view, the bottled water industry has a relatively simple value chain: water is bottled at the source, shipped to the destination market on trucks and distributed to consumers by combining it with local labor. After controlling for local distribution costs in retail and transport costs, I exploit spatial variation in model-implied marginal costs and estimate variable trade frictions under arbitrary returns to scale or scope.

Finally, focusing on European countries is interesting because a central goal of the European Single Market (ESM) is the integration of the goods markets of its individual member states. Whereas different types of formal trade barriers were gradually removed, research has repeatedly reaffirmed the presence of substantial cross-border price dispersion between EU countries (e.g. Fontaine et al. (2020) and Beck et al. (2020)). Hence, separating trade frictions from differences in destination-specific marginal costs is crucial to understanding whether trade frictions still segment the ESM and how close it is to the integrated frontier.

Three pieces of reduced-form evidence motivate my analysis. First, spatial price dispersion is much larger between than within countries. In particular, absolute after-tax price differences between regions of different countries are on average 33% compared to only 8% within countries. Second, using a border regression discontinuity design (RDD), I find that after-tax consumer prices are on average 10.7% larger abroad compared to the domestic market. As systematic increases in markups abroad and local distribution costs seem unlikely, this is suggestive evidence that a non-trivial part of the cross-border price dispersion could be driven by trade frictions. Finally, trade flows of bottled water have a gravity structure: regional trade flows fall with distance and drop discontinuously at country borders. The gravity model predicts that country borders have a trade-reducing tariff-equivalent effect between 14.3% and 43% depending on the assumed elasticity of substitution. At the same time, the gravity-based estimates should be interpreted with caution as preferences and the availability of different products might differ between countries.

To account for differences in preferences, market structure and other destination-specific costs, I specify a structural model of demand for and supply of bottled water. On the demand side, I consider a model of consumer preferences in which consumers make a discrete choice out of the set of water varieties that are available in the region in which they reside. The supply side accounts for the fact that manufacturers reach final consumers through retailers. To determine equilibrium consumer prices, I consider various vertical setups in which either upstream manufacturers, downstream retailers or both set prices in a simultaneous move Bertrand game. In each instance, I rely on the equilibrium conditions of the model to back out the distribution of destination-specific marginal costs. Combined with information on production locations, data on local labor costs and transportation time, cross-border variation in backed-out marginal costs identifies the cross-border trade frictions.

Estimating the model yields two insights. First, I recover substantial cross-border trade frictions between countries. After controlling local distribution costs and transport costs, the marginal cost of selling a bottle of water rises on average by 9 cents per liter. Put differently, when expressed relative to the marginal costs of selling the same bottle domestically, the foreign price is 20% higher. There is also substantial heterogeneity in trade frictions between Euro and non-Euro countries. Whereas they are a little over 7 cents per liter between countries of the Eurozone, they are 23 cents per liter when at least one country does not belong to the Eurozone. Second, I find that spatial price discrimination reduces price dispersion as manufacturing markups are 0.5 cents lower when the bottled water variety is sold in a foreign country. While this result aligns well with the pricing-to-market literature (e.g. Atkeson & Burstein (2008); Fitzgerald & Haller (2014); Corsetti et al. (2021)), it also shows that a standard trade model, i.e. CES-demand and monopolistic competition, would have only missed the mark by 0.5 cents.

To understand the effect of cross-border trade friction on spatial price dispersion and consumer welfare, I consider a counterfactual exercise in which I keep the set of available products fixed but remove the spatial trading frictions. In particular, I compare spatial price dispersion and consumer welfare in the integrated economy in which bottled water is supplied at the destination-specific marginal costs to an economy in which bottled water is also supplied at the destination-specific marginal costs but in which trade frictions also have to be incurred. Whereas trade frictions lead to an increase in cross-country price dispersion between countries by 5%, they lead to a reduction in consumer welfare of around 3.6 cents per liter which corresponds to a tax of roughly 10% on bottled water consumption.

This paper is related to four strands of literature. First, there is a vast literature on LOP deviations starting with Engel & Rogers (1996) and Goldberg & Knetter (1997). Recently, Cavallo et al. (2014), Dvir & Strasser (2018), Fontaine et al. (2020) and Beck et al. (2020) have reiterated the presence of large spatial price differences across countries. While, spatial price discrimination (e.g. Goldberg & Verboven (2001) and Gopinath et al. (2011)) and differences in local non-traded input prices (e.g. Crucini et al. (2005) and Parsley & Wei (2007)) have been suggested as possible explanations, separating trade frictions from both spatial price discrimination and other sources of destination-specific marginal costs has been elusive. One exception is Asplund & Friberg (2001) who focus on price setting by duty-free shops in which the aforementioned explanations can be ruled out. Whereas their focus is on price adjustment over time, this paper focuses on cross-sectional and spatial price dispersion.

Second, this paper is related to an emerging literature that leverages within-country price dispersion to evaluate policies geared towards improved market integration and the incidence of international shocks. Shiue & Keller (2007) and Donaldson (2018) focus on commodity markets and utilize within-country price dispersion to evaluate the level and changes in market integration in China and India. Atkin & Donaldson (2015) and Chatterjee (2023) study how transport infrastructure and regulatory barriers respectively interact with the competitiveness of intermediaries along the value chain in determining spatial price dispersion. Relative to this literature, this paper leverages within- and between-country price dispersion to estimate the level of trade frictions and evaluate the distance between the current level of EU market integration to a fully integrated economy.

Third, this paper connects to literature in international trade that leverages variation in trade flows and a gravity structure to estimate the level or changes in trade costs (e.g. McCallum (1995) and Anderson & Wincoop (2003)). For instance, K. Coşar et al. (2015), Head & Mayer (2021) and Santamaría et al. (2023) rely on the gravity framework to estimate the level and reduction in border frictions across European countries. While the gravity model is parsimonious and arises in the equilibrium of a broad class of models (see Allen et al. (2020)), symmetry assumptions on the economic environment are often necessary to map variation in trade flows into trade frictions. By combining spatial variation in both trade flows and consumer prices, I recover trade frictions while allowing for differences in consumer preferences, market structure and technology between markets.

Finally, methodologically, this paper connects with the literature on the intersection between international trade and industrial organization that focuses on particular industries to understand the effect of policies and distortions. For instance, Goldberg (1995), Verboven (1996), Berry et al. (1999) and Loecker (2011) study how changes in trade policy affected prices, productivity and consumer welfare in the US car and the Belgian textile industry. Relatedly, Goldberg & Verboven (2001), Hellerstein (2008) and Nakamura & Zerom (2010) rely on industry equilibrium models of the automobile, beer and coffee markets respectively to study exchange rate pass-through into final consumption prices. Closely related to this paper is Kalouptsidi (2018) which combines cross-country data on shipbuilding with an industry model to back out changes in the marginal cost of shipbuilding following Chinese government subsidies. This paper employs a similar empirical approach to estimate the level of trade frictions and understand how price dispersion and welfare would look in a fully integrated industry.

The rest of the paper is structured as follows. Section 2 provides an overview of the data sources we rely on and section 3 provides two pieces of motivational evidence about the European bottled water industry. Section 4 develops the structural model and section 5 discusses the identification and estimation of the key parameters of the structural model. Section 6 computes and discusses the two counterfactual exercises and section 7 concludes.

2 Data

This section elaborates on the construction of the dataset. I start by outlining the cross-country scanner data on purchases of bottled water. Hereafter, I provide more background on some of the technological and institutional aspects of the industry that I leverage to obtain production locations. Finally, I lay out the data sources used to construct the data on transportation costs, indirect taxes and labor unit costs.

2.1 Consumption data

To construct the consumption data, I rely on household-level scanner datasets from eight countries. In each country, a market research firm provides households with a scanning device to register all their purchases within a set of grocery products. From 2010 until 2019, the raw data registers a total of 15 million purchases of bottled water. Each transaction identifies the purchased barcode, the retail chain, the household, the number of units and volume purchased, and the tax-inclusive value in local currency. I express purchased quantities in terms of liters and consumer prices in terms of the price per liter in euros.[1] To this end, I combine package information scraped from barcode descriptions in the data with reported information on units sold, volume sold, and expenditure.

Countries The sample includes Belgium, Germany, Denmark, France, The Netherlands, Poland, Sweden and the United Kingdom.[2] As Belgium, Germany, France, and The Netherlands have been part of the formation of the ESM and the Eurozone from the outset, they arguably represent a set of relatively well-integrated countries. This makes it likely that the estimated cross-border frictions are conservative estimates relative to the broader ESM. The sample also includes four European countries that are not part of the Eurozone which I use to explore heterogeneity in the estimated trade frictions.

Retail chains The data cover all types of retail chains. In the main analysis, I restrict the store dimension in two ways. First, purchases can in principle be made both domestically and across the border but I only focus on domestic purchases. To identify cross-border transactions, I rely on outlet-level information and relative consumer prices. Even though the data does not identify the individual outlet in which the barcodes were purchased, the data indicates whether the outlet is domestic or foreign. I exclude transactions made in foreign stores. I also exclude a handful of transactions with relatively low consumer prices made close to country borders. In particular, for all country pairs with a contiguous border, I designate transactions on products available in both countries made within 40 kilometers from the country border with a consumer price that is less than 75% of the median consumer prices for that product as cross-border transactions.[3] Table A.1 provides an overview of the prevalence of cross-border shopping. Overall, the share of transactions that involve cross-border shopping is less than 0.6% representing a little over 1% of the purchased volume. Cross-border shopping is a little more pervasive in Belgium and the Netherlands at 4.6% and 2.8% of the total number of transactions respectively. Second, I do not include all retail chains in the final sample. I focus on purchases made in retail chains with a volume-based market share in the bottled water industry above 1% and I aggregate across different formats. Including all important retail chains allows me to capture differences in the local retail market structure. Moreover, Table A.2 documents that after excluding transactions involving cross-border transactions and small retail

---

[1] To convert between Euros and Local Currency Units (LCUs), I use daily FX rates obtained from Eurostat.

[2] The market research firm in Belgium, Denmark, Germany, the Netherlands, Poland and Sweden is GfK. In France and the UK, the data is gathered by Kantar.

[3] The median price is computed by considering all transactions that involve that particular product in that particular retail chain in the same quarter.

chains, the sample covers around 92% and more than 90% of the number of transactions and the volume sold respectively.

Water varieties and firms Individual bottled water varieties are defined as the combination of a brand, an indicator of whether the variety is flavored or not and a bottle size. First, each variety is associated with a brand and its respective source. To this end, I combine data from GS1 and hand-collected data to determine the ownership matrix of brands in each country. For A-label brands, the owner of the brand is also its producer. At the same time, Table A.3 shows that private label brands, brands owned by retail chains, are also pervasive across countries. Following Bonnet & Dubois (2010) and Molina (2021) and in line with the fact that private label water tends to be much cheaper all else equal (see Table A.5), I assume the retail chains are vertically integrated with the producer of that particular variety. Second, whereas water is mostly unflavored in Belgium, Germany and France, flavored bottled water is much more pervasive in Denmark, Sweden and Great Britain (see Table A.3). Nevertheless, Table A.4 shows that most price variation across varieties with different flavors is captured by an indicator variable capturing whether the water is flavored or not. Therefore, I do not distinguish between individual flavors. Conditional on all other product characteristics, flavored water is on average 34% more expensive than unflavored water (see Table A.5). Third, a key dimension of horizontal product differentiation is offering different bottle sizes. In the raw data, there are over 20 different bottle sizes and Table A.4 shows that simply dividing bottle sizes into small ($\leq 750ml,$), medium $(750ml, 1500ml)$ and large bottles ($\geq 1500ml$) does not capture the variation in prices across bottle sizes. To ensure that price differences do not stem from differences in observed product characteristics (e.g. D. Hummels & Skiba (2004), Manova & Zhang (2012) and Bastos et al. (2018)), I do not aggregate across different bottle sizes and define varieties at the exact bottle size. Conditional on other product characteristics, medium-sized bottles and large bottles are respectively 40% and 80% cheaper compared to small bottles (Table A.5). Finally, I collapse two other potential product characteristics: still versus sparkling water and the bottle material. Tables A.4 and A.3 show that these characteristics do not explain variation in consumer prices conditional on the other product characteristics.

I place two sample restrictions at the product level. On the one hand, I exclude observations with consumer prices that are more than five smaller or larger compared to the median price of that variety in that particular quarter.[4] In addition, I winsorize prices that are between three and five times smaller or larger compared to the median price. Table A.7 shows that 1.3% of the number of transactions and less 1 % of the volume sold is affected. On the other hand, I only include varieties that have a market share of at least 0.1% in terms of volume or that are sold in at least two countries and in at least 40 markets, where a market is defined as a quarter and region combination. In this way, the final sample includes both important local varieties and traded varieties with a minimum number of transactions. Table A.6 shows that this sample restriction preserves more than 80% of the volume sold in most countries and 69% of the volume sold on aggregate. Due to the fragmented market structure, the final

---

[4] In doing so, the median price is computed across retail chains and all regions.

sample only covers 54% of the volume sold in Germany. The reduced-form evidence presented in section 3 is nonetheless robust to the store- and product-level sample restrictions.

Table 1 provides an overview of the final sample. The final sample includes a little under 10 million transactions on 346 different products owned by 70 different firms bought in 106 different stores. Importantly, excluding certain products or retail chains from the final sample has no bearing on my ability to capture differences in market structure as all transactions not included in the final sample remain in the outside option.

Table 1: Sample overview

Variable	Overall	BEL	GER	DEN	FRA	NLD	PLN	SWE	UK
Regions	154	11	38	5	22	12	17	8	41
All sources	200	68	76	8	60	34	40	13	34
Local sources	-	8	60	5	50	3	33	4	23
Firms	127	32	41	8	22	23	42	12	20
Brands	267	59	94	23	53	40	55	20	49
Products	767	226	182	73	187	117	130	69	130
Stores	106	18	15	20	17	25	24	11	13
Households - All	704	352	730	361	561	754	404	299	984
Households - Water	436	286	539	125	470	399	349	93	445
Transactions ('1,000')	12,380	805	3,868	86	2,582	956	1,419	126	2,538
Share water - uncond.	0.23	0.25	0.29	0.11	0.32	0.11	0.36	0.09	0.12
Share water - cond.	0.34	0.32	0.40	0.30	0.39	0.21	0.42	0.30	0.27
Inside good share	0.64	0.79	0.72	0.34	0.83	0.51	0.86	0.31	0.44
Frequency of purchase	0.77	0.87	0.84	0.60	0.89	0.71	0.89	0.57	0.67
Unit price (incl.)	0.37	0.37	0.20	0.61	0.28	0.45	0.20	0.83	0.38
Unit price (excl.)	0.30	0.35	0.16	0.45	0.26	0.35	0.15	0.73	0.30

Notes: This table provides an overview of the dimensions of the dataset. For each of the eight countries and the overall dataset, I show the number of regions included, the number of bottled water sources and how many of those are local to the destination country. Also, I show the number of firms, brands and products, defined as the interaction between brands, whether it is sparkling, whether it is flavored, whether it is a plastic, glass or other package and whether it is a small, medium or large bottle. I show the number the total number of retail chains from which water is bought, the total number of households per NUTS2-region and quarter included in the sample and the total number of households per NUTS2-region and quarter that purchase water. I show that expenditure share of water purchases in total purchases of non-alcoholic beverages once averaged across all households ("uncond.") and once averaged across households that purchase water ("cond."). The inside good share is the ratio of the sum of the population weights of households that purchase and the sum of the population weights of households. Finally, I show the price per liter in Euros inclusive and exclusive of indirect taxes averaged across all transactions.

Household characteristics Observing the data at the household level offers several advantages. First, among other household characteristics, the data records the ZIP code where the household resides. This information is crucial to obtain within- and between-country variation in prices and consumption. Assuming that households purchase and consume bottled water near their residence, I disaggregate consumer prices and purchased quantities at the regional level. Second, in addition to ZIP codes, the dataset also records various other household characteristics, such as the size of the household, the age of the household head, the net monthly income of the household and the population weight. To ensure comparability across countries, I discretize the household characteristics and obtain 18 different consumer types. More specifically, I categorize households into two types: individuals or couples with one or two members, and families consisting of three or more members, I consider three age groups: $[\leq 34y, 35y - 64y, \geq 65y]$ and I distinguish between three income categories: $[\leq 1,900$ EUR$, 1,900 - 2,700$ EUR$, \geq 2,700$ EUR$]$. Table A.9 shows that the probability of buying bottled water is 30% larger for middle-aged families relative to young couples and that households with an income above $\geq 2,700$ EUR pay on average 4.1% more per liter for the same type of bottled water. Table A.8 documents considerable variation in household characteristics across countries, so observing household characteristics is essential to capture this cross-country consumer heterogeneity and discipline the model-implied substitution patterns. Finally, the dataset also records household-level population weights. When aggregating across transactions, I use population weights to maximize the external validity of the results.

Regions To estimate cross-border frictions, my strategy relies on comparing between- and within-country variation in backed-out marginal costs of supplying different regional markets. In doing so, I balance the desire for a fine spatial dimension with the need to have sufficient sample coverage within each regional market. For this reason, I define regional markets within each country at the NUTS 2 (rev. 2013) level.[5] Defining regional markets in this way yields 154 regional markets and around 600 sampled households per regional market and per quarter on average (Table 1).

Potential market and outside option A key reason why prices and consumption of bottled water differ between countries is differences in the appeal of not purchasing bottled water. In markets where the (perceived) quality of the available tap water is high, demand for bottled water may be lower. To account for regional differences in the importance of the outside option, I compute the outside option in three steps. First, for each region and quarter combination, I compute the sum of the population weights accounted for by consumers who record at least one purchase in the NARTD category.[6] Second, I compute the sum of the population weights accounted for by consumers who also purchased water. Finally, the inside good share is then simply the ratio of the sum of the population weights of consumers that purchase water over the sum of population weights of consumers that purchase beverages. Relatedly, the potential market is the ratio of the total water consumption (in liters) divided by the inside good share. Table 1 reports substantial cross-country variation in the inside good share that intuitively varies with average consumer prices of bottled water and the number of local sources. For instace, the inside good share is around 80% in countries with many domestic sources such as France, Germany and Poland and is below 40% in more expensive countries like Denmark and Sweden.

---

[5] The NUTS classification is a European standard used for referencing administrative levels within countries. After administrative reforms in 2016, France changed their NUTS classifications. Because the regional variable in the French dataset corresponds to the NUTS2-level of the NUTS-2 (rev. 2013) version, we use the 2013 version of the NUTS regions throughout the paper.

[6] The NARTD or Non-Alcoholic Ready-To-Drink category includes bottled water, sodas, juices and energy drinks.

2.2 Production locations

Estimating trade frictions requires knowledge about the location of different steps of the value chain. While data on trade flows typically contain production and consumption locations, scanner data do not.[7] This is particularly cumbersome in differentiated product markets where multinational firms account for the bulk of trade flows (Antràs et al. (2024)) and can serve markets through direct exporting (e.g. Melitz (2003)), local affiliates (e.g. Helpman et al. (2004)) or more sophisticated strategies based on export platforms (e.g. Tintelnot (2016)). The presence of these intricate firm-level organizational structures complicates isolating trade frictions from destination-specific marginal costs, such as differences in local input prices.

To cut through this complexity, I leverage regulatory and technological features of the bottled water industry. First, the European Commission requires manufacturers to disclose the source from which the water was extracted on the product's label. This requirement ensures that the mineral content of the water inside the bottle coincides with the one advertised on the label and reported to food safety authorities. Second, to abide by these rules in the most cost-efficient way, manufacturers bottle the water close to its source and ship the atmospherically sealed bottles to different destination markets. With these two features in hand, I take advantage of the registry of Directive 2009/54/EC that records information on water brands, their source and the location of the water source.[8] I augment these data with hand-collected information from company websites and product labels of varieties not included in the registry.[9]

For each country, Figure A.1 shows the distribution of total volume sold per sourcing mode, i.e. domestic, foreign or unknown. When the source is unknown I further distinguish between A-level brands and private labels. Across most countries, I classify around 80% of the total volume as either domestically sourced or as sourced from abroad.[10] For a portion of the volume, I am unable to determine the production location. Figure A.1 indicates the share of unknown sources is concentrated in the set of private label varieties. Unlike A-level brands, which use their particular source as a source of differentiation, private-label varieties tend to be more homogeneous and low-cost.[11] This gives retail chains more flexibility in terms of choosing from where to source their water, making it sometimes impossible to pinpoint the production location.

Figure 1 provides an overview of water sources from which bottled water is sourced for the countries in the dataset. Figure 1a indicates the water sources from which bottled water sold in

---

[7] Even in trade data the true country of origin may be incorrectly recorded as different datasets rely on different definitions, i.e. country of consignment or country with the last substantial transformation, to determine the country of origin. Cotterlaz & Vicard (2023) show this issue is especially prevalent for trade involving European countries that partake in the EU customs union and that it quantitatively affects the estimated trade enhancing effects of EU membership.

[8] Directive 2009/54/EC regulates the branding and information disclosure for bottled water to be sold to consumers across European countries.

[9] Most of the time these are private label brands and, in this case, I rely on the source disclosed on individual product labels to retrieve the source and the municipality.

[10] The Netherlands is an outlier as slightly less than 60% can be classified as either domestic or foreign sourcing.

[11] Table A.5 shows that conditional on product characteristics, private label varieties are on average 50% cheaper compared to varieties from A-level brands.

Belgium is sourced with a dark dot and other water sources not used for Belgian consumption with lighter dots. Figure 1b shows the same information for German bottled water consumption. German consumers benefit from a larger number of local water sources compared to Belgian consumers and consume relatively more (domestically sourced) bottled water which is consistent with lower average consumer prices in Germany (see Table 1).

Figure 1: Water sources

(a) Belgium **(b)** Germany

Notes: The figure plots the sources from which water is being sourced in Belgium and Germany respectively. In each country, we show active sourcing locations in darker colors and indicate the relative importance, based on volumes expressed in liters, for that destination country through the size of the dot. In addition to the active sourcing locations, we also show the other sourcing locations in the dataset which are used as a sourcing location by the other destination countries.

2.3 Other data sources

I complement the production and consumption data with several additional variables that need to control for differences in destination-specific marginal costs unrelated to cross-border frictions.

Indirect taxes It is well known that indirect taxes differ substantially between European countries. Whereas Belgium and France impose only a 6% VAT rate, Germany and Denmark charge a VAT rate of 21% and 25% respectively. In addition, excise duties often target beverages with added sugar or artificial sweeteners but their level and time of introduction differs from country to country. To account for such differences between countries, I collect data on indirect taxes, including Value Added Taxes (VAT), excise duties, and taxes on disposable packages, from the "Taxes in Europe Database" constructed by the European Commission and from country-level regulatory authorities. Appendix A.5 provides an overview of the different taxes applied by the different countries in the sample.

Labor unit costs A key source of differences in destination-specific marginal costs is differences in non-traded input prices (see Crucini et al. (2005), Parsley & Wei (2007) and Burstein et al. (2005)). When manufacturers sell to final consumers through retailers, locally sourced labor is typically employed to handle inventories, restock shelves and provide customer service. As European labor markets are also segmented along country borders, differences in retail wages need to be controlled for. To this end, I construct a panel of unit labor costs in the retail sector that varies at the NUTS-1 and year level. To this end, I rely on the microdata from the EU-SILC database in which households report, among other characteristics, their NACE 2-digit sector of occupation, their residence at the NUTS1 level and different components of their pecuniary income, non-pecuniary income and social security contributions by their employers. Appendix A.6 provides more detail on the constructions of these series.

Transportation costs Finally, a second source of differences in destination-specific marginal costs is differences in transport costs. I control for differences in transport costs by combining three datasets. The first dataset is a cross-sectional dataset that provides road travel distances and travel times for transport trucks between production and consumption locations at the ZIP code level.[12] Doing so, I abstract from the presence of distribution centers for different retailers. The second dataset is a panel of diesel prices at the country and the month level accessed through Eurostat. I interact travel distances with diesel prices to capture time variation in transport costs. Finally, I collect data on unit labor costs in the transportation sector from the same source as the labor unit costs. Interacted with the travel times, this is a second source of time variation in transport costs that is accounted for.

3 Motivating evidence

This section documents spatial price dispersion in the European bottled water industry using three pieces of reduced-form evidence. First, I show that absolute Law of One Price (LOP) deviations are much larger between countries compared to within countries. Second, I estimate a border Regression Discontinuity Design and illustrate that an important part of the price dispersion between countries occurs discontinuously at country borders. Finally, I establish a first estimate of the tariff-equivalent trade frictions necessary to rationalize the variation in trade flows and prices by estimating a structural gravity equation.

3.1 Spatial price dispersion

Table 1 documents substantial cross-country differences in the average price paid for bottled water, ranging from 29 cents per liter in Germany and Poland to 83 cents/L in Sweden. Also, Table A.5

---

[12] I purchased these from Localyze.eu which is a private company that provides logistics companies with data and insights.

shows that conditional on other product characteristics, water sourced from abroad sells at a 22% price premium. I now investigate spatial price dispersion and the potential contribution of cross-border frictions more thoroughly.

Absolute LOP deviations To compute LOP deviations, I first compute variety-level prices by aggregating transactions across the household and store dimensions using liters sold interacted with population weights within each quarter and region. Hereafter, I compute for each region pair and quarter the absolute price differences for all varieties that were sold at any point during that quarter in both regions.

Figure 2 compares the conditional distributions of absolute LOP deviations for domestic region pairs, pairs of regions belonging to the same country, and international region pairs, pairs of regions that do not belong to the same country. Figure 2a plots these distributions for final consumer prices, inclusive of taxes. Even though spatial price dispersion is on average 8.2% within countries, absolute LOP deviations are close to zero in around 50% of the cases. However, spatial price dispersion is much larger between European countries. On average, international spatial price dispersion is 35.5%, or roughly 27ppt higher.

One potential reason why spatial price dispersion is larger between regions belonging to different countries is that indirect taxes differ between countries. To underscore that differences in final consumer prices are not mainly driven by differences in indirect consumer taxes, figure 2b replicates figure 2a now relying on after-tax consumer prices. International spatial price dispersion drops slightly but international price dispersion is still roughly 25ppt higher compared to domestic price dispersion.

Border regression discontinuity design Several reasons could explain the shift in the distribution of absolute LOP deviations for international region pairs relative to one for domestic region pairs. A natural explanation is the fact the average pairwise distance between regions that belong to different countries relative to regions of the same country is larger. Other potential reasons are that spatial price discrimination, differences in destination-specific marginal costs or trade frictions are larger between countries.

To further discriminate between these explanations, I consider a border Regression Discontinuity Design. I rely on this specification since uncovering a positive and significant jump in prices right across the border relative to domestic prices provides suggestive evidence of positive trade frictions. This is because comparing prices country boundaries eliminates average differences in distance between international and domestic region pairs as a likely explanation. Also, standard spatial price discrimination arguments would predict markups to fall when supplying a certain market is more costly.[13] Furthermore, while certainly possible, it is unclear why local distribution costs would be systematically higher abroad compared to the domestic market.

---

[13] In this view, recovering a tightly estimated zero RDD estimate for differences in consumer prices does not provide evidence against the existence of positive trade frictions as they could be counteracted by lower markups. Therefore, this evidence is only sufficient and not necessary to test the existence of positive trade frictions.

Figure 2: Spatial price dispersion

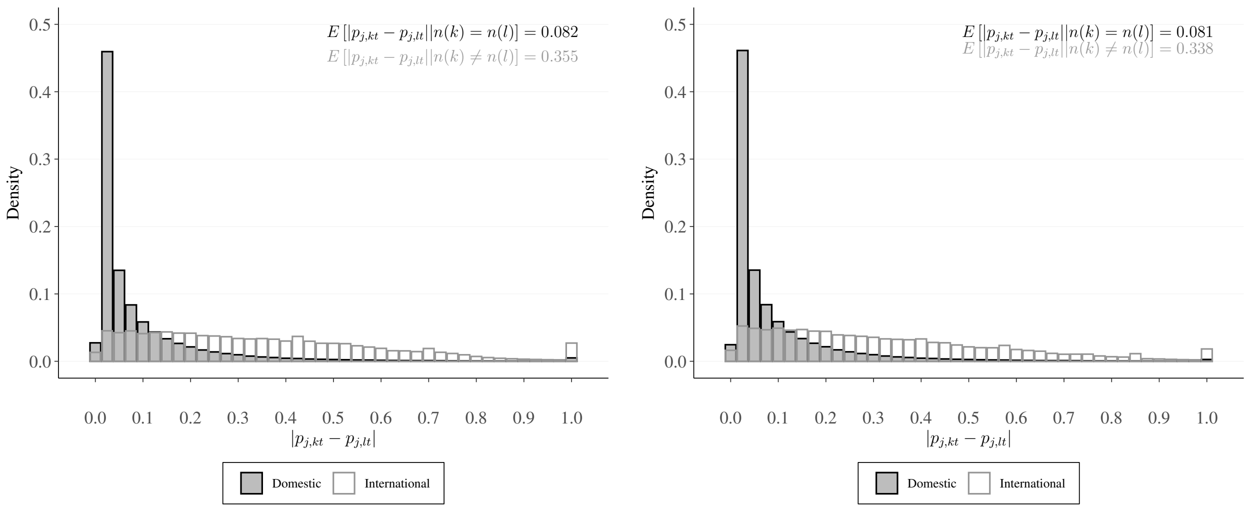

(a) Domestic vs. International - Tax-inclusive

(b) Domestic vs. International - Tax-exclusive

Notes: These figures plot conditional distributions of the absolute log LOP deviations between regional markets, where a regional market is defined at the NUTS2 level. Panel (a) compares the conditional distribution of absolute LOP deviations between domestic and international region pairs. Panel (b) compares the conditional distribution of absolute LOP deviations between international region pairs part of the Eurozone and international region pairs not part of the Eurozone. To compute absolute LOP deviations, I first obtain variety-level for every year and regional market by aggregating across households and stores using population weights. Hereafter, I compute for each region pair the absolute price differences for all varieties that were sold at any point during the year in both regional markets. Finally, I censor the absolute LOP differences at a standard deviation of 1.

To construct the RDD estimates, I proceed in four steps. First, for each contiguous country pair, I select the products that are sold in both countries. Second, I compute ZIPcode-level weighted average consumer prices where I use the interaction of liters sold and the population weight as the weights. Third, I rank observations in terms of their great circle distance to the border and normalize the distance to the border such that products that are sold domestically have a negative distance and goods sold abroad a positive distance. Finally, I consider the following RDD specification:

$$\ln\left(p_{j,t}^{s(j)l}\right)=\beta\text{Border}^{s(j)l}+f^{n}\left(\text{Dis}^{s(j)l};\boldsymbol{\gamma_0}\right)+f^{n}\left(\text{Dis}^{s(j)l};\boldsymbol{\gamma_1}\right)+\lambda_{j,t}+\varepsilon_{j,t}^{s(j)l}$$

where $p_{j,t}^{s(j)l}$ is the consumer price of product j at time t when going from the product j's source $s(j)$ to location l, Border$^{s(j)l}$ indicates whether $s(j)$ and l are in a different country, $f^{n}\left(\text{Dis}^{s(j)l};\boldsymbol{\gamma_0}\right)$ and $f^{n}\left(\text{Dis}^{s(j)l};\boldsymbol{\gamma_1}\right)$ are polynomials of order n in the distance to the border, which are allowed to differ on either side of the border. I include $\lambda_{j,t}$, which product-time fixed effects, to ensure that the polynomials and the border effect are identified only using spatial variation within product-quarter pairs.

Before turning to the formal RDD estimates, Figure 3 graphically illustrates that spatial price dispersion happens discontinuously at national market boundaries. In particular, Figure 3a shows a binscatter consisting of 30 bins and a first-order polynomial on either side of the French-German border. This figure shows that domestic prices tend to rise as the distance from the source rises. At the same, foreign prices discontinuously jump at the border and are on average 20% higher compared to domestic prices. Figure 3b replicates Figure 3a now using after-tax consumer prices to construct the binscatter. Domestic prices continue to rise as the distance from the source increases and foreign prices are now on average 12% higher compared to domestic prices.

Figure 3: Border RDD: French-German border

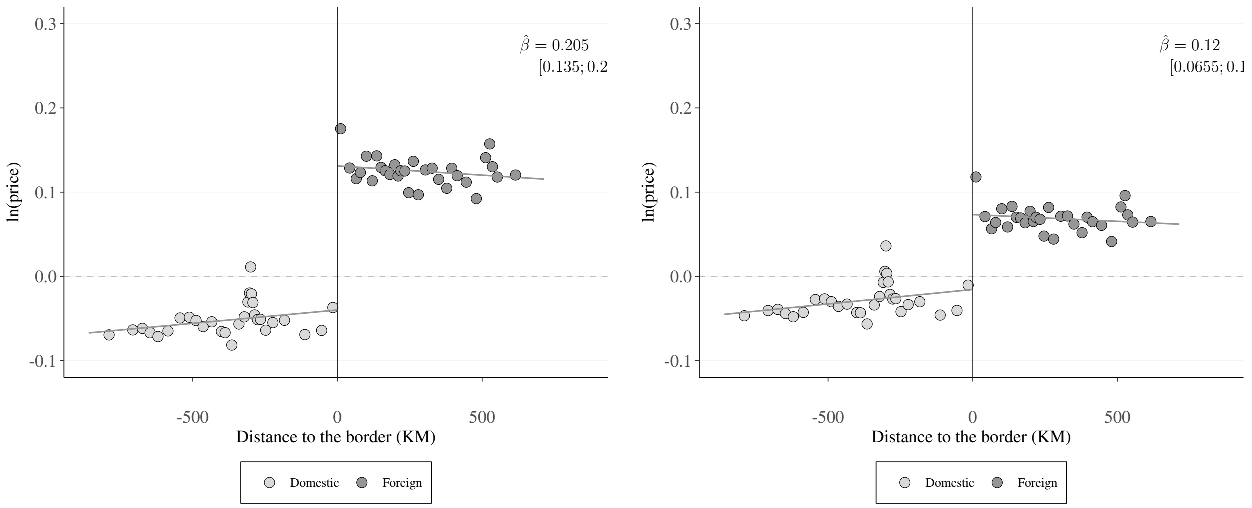
(a) Raw sample - Tax-inclusive consumer prices **(b)** Clean sample - Tax-exclusive consumer prices

Notes: This figure shows the binscatter corresponding to estimating equation 3.1 with first-order polynomials where I have used 30 equally-spaced bins on either side of the border. I limit the sample to observations within a 1,000km range from the border and to the most important stores and products. Panel (a) shows the binscatter results for tax-inclusive consumer prices and panel (b) for tax-exclusive consumer prices. In both cases, for each contiguous country pair, the set of products sold in both countries is selected and observations are ranked according to their great circle distance to the border. The resulting sample pools across these contiguous countries. ZIPcode-level prices are constructed as the weighted average across transactions where the weights are the population weights time the purchases liters associated with that transaction.

Table 2 shows that the insights of Figure 3 carry over to all observations in the main sample. The table reports the results from estimating equation 3.1 for after-tax consumer prices on the main sample. I cluster standard errors at the product level and construct confidence intervals that are robust to the fact that the chosen bandwidth may be too large for the traditional confidence intervals to have good coverage (see Calonico et al. (2014)). Considering first-order polynomials and including all observations within a 1,000 KM range from the border in column (1), I estimate that export prices are on average 16.7% higher compared to domestic prices. The difference slightly rises to 17.2% and drops somewhat to 10.7% when I limit the range of observations to be within a distance of 500 and 100 KM respectively. When I consider the optimal bandwidth selection following Calonico et al. (2014), the difference is estimated at 10.1%. In all cases, the difference is statistically significant at the conventional levels. Columns (5) to (8) show that these results are largely unaltered when I consider a second-order polynomial in distance to the border at either side of the border.

As mentioned in section 2, the main sample excludes observations that are designated as cross-border transactions. Although the structural model developed in section 4 will account for cross-border shopping through its effect on the outside option, Table B.2 shows that the results presented in Table 2 are robust to including cross-border transactions. When I include all transactions involving varieties in the main sample, the results are very similar to the results presented in Table 2. In addition, Tables B.1 and B.3 show that foreign prices are still significantly larger compared to domestic prices when I include all products (Tables B.1) and all products in all stores (B.1). Taken together, these

Table 2: Border Regression Discontinuity Design: Results

	1^{st}-order				2^{th} order			
$p_{j,lt}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\text{Border}^{s(j)l}$	0.167^{***} $[0.13; 0.203]$	0.172^{***} $[0.134; 0.21]$	0.107^{***} $[0.0803; 0.134]$	0.101^{***} $[0.0755; 0.127]$	0.179^{***} $[0.14; 0.218]$	0.155^{***} $[0.119; 0.191]$	0.0888^{***} $[0.0657; 0.112]$	0.0927^{***} $[0.0683; 0.117]$
$\theta_{j,t}$	✓	✓	✓	✓	✓	✓	✓	✓
Polynomial	1	1	1	1	2	2	2	2
Bandwidth	1,000	500	100	54.1	1,000	500	100	82.9
Optimal				✓				✓
No. obs	1,783,315	1,394,743	425,812	253,914	1,783,315	1,394,743	425,812	360,664

Notes: This table shows the results from estimating equation 3.1 using the robust Regression Discontinuity Design (RDD) estimator from Calonico et al. (2014) using the rdrobust package. Columns (1)-(4) show the results from estimating equation 3.1 using a first-order polynomial and Columns (5)-(8) for a second-order polynomial in the distance to the border. For each polynomial order, I show the results when I limit the sample in a range of 1,000, 500 and 100 kilometers to the border and when I consider the optimal bandwidth selection from Calonico et al. (2014). The sample underlying this table includes the most important stores and products and is constructed by ranking, for each contiguous country pair, prices according to their distance to the border. The resulting sample pools across these contiguous countries. Prices are constructed as the weighted average across transactions where the weights are the population weights time the purchases liters associated with that transaction. Standard errors are clustered at the product level reported and are robust to the fact that bandwidths that are far away from zero can lead to bad coverage of the confidence intervals (see Calonico et al. (2014)). I report the robust confidence intervals in square brackets and denote significance at the $p < 0.1^{*}$,$p < 0.05^{**}$ and $p < 0.01^{***}$ levels.

results suggest that an important determinant of between-country spatial price dispersion stems from a discontinuous jump in consumer prices at country borders that is consistent with the existence of positive trade frictions.

3.2 Structural gravity

A natural start to translate the between-country variation in prices and consumption into tariff-equivalent trade frictions is by estimating a gravity equation. In particular, I consider a gravity model of the following form:

$$X_t^{lk} = \left(\frac{\tau_t^{lk}}{\Pi_{lt} P_{kt}} \right)^{-\sigma} Q_{lt} Y_{kt} \tag{1}$$

where X_t^{lk} is the bilateral trade flow of bottled water from regional market l to regional market k, Π_{lt} is the outward multilateral resistance term capturing the capability of regional market l to supply bottled water to other regional markets, P_{kt} is the inward multilateral resistance summarizing the ease with which other regional markets reach market k and Q_{lt} and Y_{kt} are nominal output and demand in market l and k respectively. The bilateral trade cost vector is given by τ_t^{lk} which captures trade frictions that affect trade between the regional markets. Head & Mayer (2014) and Allen et al. (2020) show that this model arises in a wide array of international trade models, such as Armington (1969)-type (e.g. Anderson & Wincoop (2003)), Ricardian-type (e.g. Eaton & Kortum (2002)) and increasing returns to scale models (e.g. Krugman (1980)) models. To operationalize this model, I follow Silva & Tenreyro (2006) and consider the following specification:

$$X_t^{lk} = \exp\left(\beta \ln(1 + \text{Dis}^{lk}) + \gamma_B \text{Border}^{lk} + \gamma_C \text{Cur}^{lk} + \lambda_{lt} + \lambda_{kt}\right) + \varepsilon_t^{lk} \tag{2}$$

where Dis^{lk} is the population weighted distance between regional markets k and l, Border^{lk} is an indicator variable that equals one when k and l are separated by a country border. Cur^{lk} is an indicator variable that equals one when k and l do not use the same currency. I also include λ_{lt} and λ_{kt} which are origin-time and destination-time fixed effects that control for variation induced by the unobserved multilateral resistence terms and nominal output and demand in the origin and destination markets. Finally, ε_{lkt} is defined as the deviation of the observed trade flows X_t^{lk} from its model prediction $\mathbb{E}\left[X_t^{lk}|\text{Dis}^{lk}, \text{Border}^{lk}, \text{Cur}^{lk}, \lambda_{lt}, \lambda_{kt}\right]$. Hence, conditional on origin-time fixed effects, which control for origin-specific supply shocks affecting cost or product quality, destination-time fixed effects, accounting for destination-specific demand shocks, and a proxy for transport costs, the gravity equation attributes discontinuous differences in trade flows to cross-border trade frictions.

Table 3 presents the results of estimating Equation 2 on the main sample using the PPML estimator of Silva & Tenreyro (2006). I report two-way standard errors clustered at the region pair level. Column (1) shows that regional trade flows for bottled water fall with the distance between the origin and destination market. A 10% increase in distance leads to a drop in trade flows of bottled water by 8.15%. Column (2) also shows that international trade flows, trade flows between international region pairs are, on average, 69.7% lower compared to domestic trade flows. Depending on the assumed elasticity of substitution, this corresponds to a tariff-equivalent trade friction between 13.4% and 43.5%.[14] This sharp decrease in international trade flows relative to domestic flows is in line with the substantial literature on border effects in trade (e.g. McCallum (1995),Anderson & Wincoop (2003) and Santamaría et al. (2023)).[15] Column (3) explores heterogeneity in the border effect by distinguishing between country borders inside and outside of the Eurozone. When two markets do not use the same currency, trade flows fall by an additional 95%. However, this only holds in the sample with the main products (see Tables B.5 and B.6).

A shortcoming of the gravity framework is that, conditional on the aforementioned fixed effects, residual differences in trade flows are fully attributed to trade frictions which affect trade flows through their effect on bilateral prices. In other words, residual differences in preferences, market structure or destination-specific marginal costs cannot be accounted for. At the same time, this mechanism also suggests a test of the validity of the gravity structure. In particular, after controlling for bilateral prices, the coefficients on Distance^{lk} and Border^{lk} should no longer be statistically different from zero. If they still hold predictive power for trade flows, other primitives must also be changing in space. Typically, this test is infeasible as bilateral prices are unobserved but the detailed scanner data allow me to construct theory-consistent bilateral prices up to an elasticity of substitution.[16] In columns (4) and (5), I re-estimate equation 2 now controlling for the vector of bilateral prices where I use an elasticity of substitution of 8.28 (following Eaton & Kortum (2002))

---

[14] I consider elasticities of substitution at the upper and lower end of estimates in the literature. In particular, $\varepsilon_{EK} = 8.28$ is taken from Eaton & Kortum (2002) and $\varepsilon_{BLP} = 2.12$ is taken from Boehm et al. (2023).

[15] For instance, Santamaría et al. (2023) find that national borders reduce international trade flows to 9% of the size of domestic trade flows.

[16] More formally, bilateral prices P_t^{lk} are given by: $P_t^{lk} \equiv \left(\sum_{j \in \mathcal{J}_t^{lk}} \left(P_{j,t}^{kl}\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$ where $\mathcal{J}_t^{lk}$ is the set of varieties flowing from l to k.

and 2.12 (following Boehm et al. (2023)) respectively.[17] In line with expectations, higher prices have a trade-reducing effect. However, Distance^{lk} and Border^{lk} remain significant predictors of trade flows, suggesting that unobserved residual differences in preferences, market structure or destination-specific marginal costs are correlated with Distance^{lk} and Border^{lk}.

Two possible violations of the exclusion restriction come to mind. First, conditional on bilateral prices, spatial variation in consumer preferences for the same varieties will lead to differences in trade flows.[18] Below, I explicitly deal with this concern by estimating a flexible demand system for bottled water. Second, Helpman et al. (2008) shows that standard trade cost estimates are biased when firms endogenously choose which markets to serve. Column (6) provides suggestive evidence for selection across markets by estimating the relationship between the similarity of the set of available varieties and the trade cost variables. I measure the similarity in the set of available varieties as the dot product of two vectors normalized by the Euclidian norms of the respective vectors: $\frac{\boldsymbol{J}_l \cdot \boldsymbol{J}_k}{||\boldsymbol{J}_l|| \cdot ||\boldsymbol{J}_k||}$ where $\boldsymbol{J}_l$ is a vector constructed from variables indicating whether variety j was ever sold in regional market l, i.e. $(\boldsymbol{J}_l)_j = 1$.[19] When regional markets l and k only have varieties in common, this measure is one. It is zero otherwise. Column (6) shows that the similarity in the set of available varieties is strongly related to trade costs. A 10% increase in the distance leads to a reduction of 1.1% in the similarity of the available varieties. Also, the similarity in the set of available varieties drops by 74% when the two regions are separated by a country border within the Eurozone and by an additional 50% when the two regions are separated by a country border and the two regions have a different currency.

4 An empirical model of the bottled water industry

This section discusses the empirical model of demand and supply I use to recover a model-implied estimate of destination-specific marginal costs that allows for spatial differences in preferences, market structure and distribution. I combine this distribution with information on production locations to estimate trade frictions as the residual variation in destination-specific marginal costs that are not explained by other sources of destination-specific marginal costs. In this way, my approach to recovering trade frictions is similar to Hsieh & Klenow (2009) and Kalouptsidi (2018).

4.1 Preferences for bottled water

Instead of working with demand systems conventionally used in international trade, I model consumer preferences for bottled water using a discrete choice framework.[20] An important advantage of

---

[17] The drop in the number of observations in columns (4) and (5) relative to the first three columns is due to the fact that prices cannot be computed for observation with zero trade.

[18] For instance, A. K. Coşar et al. (2018) shows that in the automobile industry, 50% of the differences in market shares of domestic and foreign brands can be explained by home-biased preferences.

[19] This measure is also used in Hristakeva (2022).

[20] That being said, discrete choice frameworks have been used in international trade as well. For instance, Fajgelbaum et al. (2011) relies on a discrete choice framework to generate non-homothetic preferences for differentiated products and Goldberg (1995), Goldberg & Verboven (2001), Hellerstein (2008) and Nakamura & Zerom (2010) use it to understand

Table 3: Gravity estimation

	X_t^{kl}					$\frac{\boldsymbol{J}_l \cdot \boldsymbol{J}_k}{\lvert\lvert \boldsymbol{J}_l \rvert\rvert \cdot \lvert\lvert \boldsymbol{J}_k \rvert\rvert}$
	(1)	(2)	(3)	(4)	(5)	(6)
ln(1+Dis^{lk})	-0.605^{***}	-0.519^{***}	-0.518^{***}	-0.458^{***}	-0.383^{***}	-0.108^{***}
	(0.078)	(0.077)	(0.077)	(0.073)	(0.069)	(0.015)
Border^{lk}		-1.2^{***}	-1.19^{***}	-1.01^{***}	-0.811^{***}	-1.35^{***}
		(0.236)	(0.237)	(0.239)	(0.268)	(0.057)
Cur^{lk}			-3.09^{***}	-2.53^{**}	-2.15^{**}	-0.694^{***}
			(1.178)	(1.117)	(1.043)	(0.071)
ln(P_t^{lk})				-0.537^{**}	-0.776^{***}	
				(0.247)	(0.208)	
$e^{\hat{\beta}} - 1$	-	-69.8%	-69.6%	-63.7%	-55.6%	-74.2%
$e^{\frac{\hat{\beta}}{\varepsilon_{EK}}} - 1$	-	-13.5%	-13.4%	-11.5%	-	-
$e^{\frac{\hat{\beta}}{\varepsilon_{BLP}}} - 1$	-	-43.2%	-43.0%	-	-31.8%	-
$e^{\hat{\gamma}} - 1$	-	-	-95.4%	-92.0%	-88.4%	-50.0%
$e^{\frac{\hat{\gamma}}{\varepsilon_{EK}}} - 1$	-	-	-31.1%	-26.3	-	-
$e^{\frac{\hat{\gamma}}{\varepsilon_{BLP}}} - 1$	-	-	-76.7%	-	-63.8%	-
λ_l						✓
λ_k						✓
$\lambda_{l,t}$	✓	✓	✓	✓	✓	
$\lambda_{k,t}$	✓	✓	✓	✓	✓	
No. obs	73,323	73,323	73,323	73,323	73,323	23,409

Notes: This table presents the results from estimating equation 2 using Pseudo Poisson Maximum Likelihood (PPML) to account for zeros in the dependent variable. Columns (1) to (5) present various specifications with bilateral trade flows, X_t^{kl}, as the dependent variable. I compute bilateral trade flows by computing for each destination market the tax-inclusive expenditures denominated in EUR flowing in from all origin markets included in the sample via aggregation across transactions using population weights. Bilateral prices P_t^{lk} are computed by aggregating variety-level prices using a CES-aggregator: $P_t^{lk} \equiv \left(\sum_{j \in \mathcal{J}_t^{lk}} \left(P_{j,t}^{kl} \right)^{1-\sigma} \right)^{\frac{1}{1-\sigma}}$ across all varieties that flow from l to k indicated by the set $\mathcal{J}_t^{lk}$. In doing so, I include observations from the main stores and products and consider the following elasticities$\varepsilon_{EK} = 8.28$ taken from Eaton & Kortum (2002) and $\varepsilon_{BLP} = 2.12$ taken from Boehm et al. (2023). Column (6) considers the similarity in the set of available varieties between any two given markets as the dependent variable. The similarity is computed as the similarity between two vectors that indicate whether a given variety is present in a given market, i.e. $(\mathcal{J}_l)_j = 1$ if and only if j is sold in regional market l at time t. Columns (6) further collapses the time dimension by computing similarity measures at the region pair level. I present two-way clustered standard errors at the origin and destination market level and denote the significance levels at the 0.1^{*}, 0.05^{**}, 0.01^{***} levels.

working with a discrete choice framework is that it provides a natural way to account for the option of not purchasing bottled water, which varies substantially across markets (see Table 1).

Formally, I assume that consumer $i \in 1, \ldots, N_{lt}$ in market lt, where the market is defined as a regional market l in a quarter t, makes a static discrete choice over the set of bottled water varieties available in their regional market, which is denoted by $\mathcal{J}_{lt}$. One issue with modeling demand for bottled water as a static choice is that bottled water is a storable good, which could lead to biased substitution patterns Hendel & Nevo (2006).[21] However, aggregating the data to the quarterly level eases this concern. This is because Table 1 shows that in countries in which bottled water is appealing relative to the outside option on average 90% of the households buy bottled water. With this in mind,

---

the transmission of exchange rate fluctuations into consumer prices.

[21] Ignoring such dynamic concerns typically leads to an overestimating of demand elasticities and an underestimation of the degree to which manufacturers will want to raise prices above marginal costs.

consumers choose the variety $j \in \mathcal{J}_{lt}$ that maximizes gives them the highest indirect utility $V_{ij,lt}$ given by:

$$V_{ij,lt} = \alpha P_{j,lt} + \boldsymbol{\beta}' \boldsymbol{X}_j + \lambda_{lt} + \xi_{j,lt} + \varepsilon_{ij,lt}$$

where $P_{j,lt}$ is the price of a variety j in market lt, $\boldsymbol{X}_{j,lt}$ is the vector variety-specific product characteristics, λ_{lt} are market-specific effects that capture variation in the importance of the outside good, $\xi_{j,lt}$ is an unobserved demand shifter and $\varepsilon_{ij,lt}$ captures idiosyncratic tastes that are individually and independently distributed across consumers and products following an EV(1) distribution. I normalize the valuation of the outside good to zero in all markets. The assumption that idiosyncratic tastes follow an EV(1) distribution implies that market shares, $\sigma_{j,lt}$, take the following form:

$$\sigma_{j,lt} = \frac{\alpha P_{j,lt} + \boldsymbol{\beta}' \boldsymbol{X}_j + \lambda_{lt} + \xi_{j,lt}}{1 + \sum_{j \in \mathcal{J}_{lt}} \alpha P_{j,lt} + \boldsymbol{\beta}' \boldsymbol{X}_j + \lambda_{lt} + \xi_{j,lt}} \tag{3}$$

The definition of the choice set implies that consumers only choose among varieties available in their regional market. However, as country borders within the EU Single Market do not require formal checks, consumers are in principle free to engage in cross-border shopping.[22] I decide to model preferences in this way for two reasons. First, bottled water is an inexpensive, voluminous and heavy product which makes engaging in arbitrage for bottled water unlikely to be very profitable. Hence, in case water was purchased abroad, it is likely that it would be bought alongside other products. Second, cross-border shopping is marginal in the data. In particular, Table A.1 shows that less than 1% of all transactions involve cross-border shopping. The only countries in which cross-border shopping are Belgium and the Netherlands but the share of cross-border transactions is well below 5%.

4.2 Market structure and marginal costs

To transition from the distribution of destination-specific consumer prices to a distribution of destination-specific marginal costs, I specify a model of how consumer prices are determined in equilibrium. The implied first-order conditions allow me to map the distribution of consumer prices to the distribution of marginal costs.

Market structure Manufacturers of bottled water reach final consumers through retailers. Because I do not observe how wholesale prices are determined I consider several possible ways in which manufacturers and retailers vertically interact when recovering trade frictions. To fix ideas, I now discuss the vertical structure that nests all I consider below.

In each regional market and each quarter, downstream retailers engage in price or Bertrand competition such that equilibrium consumer prices are determined as the outcome of a simultaneous move Nash equilibrium. Doing so, consumer prices are determined as the solution to the following

---

[22] Recently, Auer et al. (2023) investigates how the incidence of the appreciation of Swiss Franc differs between consumers that live close to and far away from the country's borders.

system of first-order conditions:

$$\boldsymbol{p}_{lt}^{r} = \boldsymbol{c}_{lt}^{r} + \boldsymbol{p}_{lt}^{w} - \left(\boldsymbol{\Delta}_{lt} \odot \boldsymbol{\Omega}_{lt}^{r}\right)^{-1} \cdot \boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{r}; \boldsymbol{\Theta}^{d}). \tag{4}$$

In this expression, $\boldsymbol{c}_{lt}^{r}$ and $\boldsymbol{p}_{lt}^{w}$ are the vectors of distribution costs and wholesale prices incurred by retailers. The term $\left(\boldsymbol{\Delta}_{lt} \odot \boldsymbol{\Omega}_{lt}^{r}\right)^{-1} \cdot \boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{m}; \boldsymbol{\Theta}^{d})$ is the vector of markups charged by retailers which depends on the matrix of first-order derivatives of demand relative to consumer prices $\boldsymbol{\Delta}_{lt}$, the ownership matrix $\boldsymbol{\Omega}_{lt}^{r}$ that internalizes substitution towards other varieties sold by the retailers and by the vector of market shares, $\boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{r}; \boldsymbol{\Theta}^{d})$, which depends on the demand parameters $\boldsymbol{\Theta}^{d}$.

Given this determination of consumer prices, I assume that upstream manufacturers compete in Bertrand competition such that wholesale prices are determined as the results of a simultaneous move Nash equilibrium as well. In particular, the first-order conditions that determine wholesale prices are given by:

$$\boldsymbol{p}_{lt}^{w} = \boldsymbol{c}_{lt}^{w} - \left(\mathbf{PT}_{lt} \cdot \boldsymbol{\Delta}_{lt} \odot \boldsymbol{\Omega}_{lt}^{w}\right)^{-1} \cdot \boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{r}; \boldsymbol{\Theta}^{d}) \tag{5}$$

where $\boldsymbol{c}_{lt}^{w}$ is the vector of destination-specific marginal costs, which include production costs and the costs associated with delivering varieties to the particular market. The term $\left(\mathbf{PT}_{lt} \cdot \boldsymbol{\Delta}_{lt} \odot \boldsymbol{\Omega}_{lt}^{w}\right)^{-1}$ is the vector of markups charged by manufacturers. Like with retail markups, this depends on the matrix of first-order derivatives of demand relative to consumer prices $\boldsymbol{\Delta}_{lt}$, an ownership matrix $\boldsymbol{\Omega}_{lt}^{m}$ that internalizes substitution towards other varieties owned by the manufacturing firm, the vector of market shares, $\boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{r}; \boldsymbol{\Theta}^{d})$. Different before markups, manufacturing markups also depend on the $\mathbf{PT}_{lt}$ pass-through matrix which consists of the derivative of retail prices with respect to wholesale prices and which captures the idea that manufacturers internalize that retail prices can change when wholesale prices change.

Marginal costs Substituting equation 5 into 4 and straightforward re-arranging yields allows me to back out destination-specific marginal costs as the sum of downstream distribution costs, production costs and the costs associated with delivering varieties to the particular market:

$$\boldsymbol{c}_{lt}^{m} + \boldsymbol{c}_{lt}^{r} = \boldsymbol{p}_{j,lt}^{r} + \left(\boldsymbol{\Delta}_{lt} \odot \boldsymbol{\Omega}_{lt}^{r}\right)^{-1} \cdot \boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{r}; \boldsymbol{\Theta}^{d}) + \left(\mathbf{PT}_{lt} \cdot \boldsymbol{\Delta}_{lt} \odot \boldsymbol{\Omega}_{lt}^{w}\right)^{-1} \cdot \boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{r}; \boldsymbol{\Theta}^{d}) \tag{6}$$

This approach to recovering destination-specific marginal costs heavily relies on a correct specification of preferences for bottled water and the market interactions through which consumer prices are determined. An alternative approach to obtaining marginal cost estimates is using data on the production costs of bottled water and exploiting assumptions about the technology that is used to produce bottled water. However, bottled water sold in different markets mostly originates from the same. For this reason, plant-level cost information is unlikely to yield the necessary destination-level variation in costs.

Alternative market structures To ease some of the concerns surrounding the assumptions on the market interactions that shape consumer prices, I consider a variety of possible setups. The vertical structure laid out above corresponds to the case of double marginalization and has been deployed to study mergers in the US beer market Miller & Weinberg (2017) and entry into the French telecommunications industry Bourreau et al. (2021). An important advantage of this vertical structure is that it not only accounts for spatial price discrimination on the part of manufacturers as Atkeson & Burstein (2008), Edmond et al. (2015), Corsetti et al. (2021) but also accounts for differences in the destination-specific retail market structure by allowing retail markups to vary as well. At the same time, this setup assumes that retailers take wholesale prices as given and hold no countervailing power relative to manufacturers and that retailers and manufacturers do not internalize the fact that setting markups sequentially lowers overall profits relative to a case where they would maximize joint surplus. Moreover, it restricts the type of contracting between retailers and manufacturers to linear contracting only.

To accommodate these concerns, I consider two other vertical market structures that are nested within the double marginalization case. First, I consider a case in which retailers compete à la Bertrand and set prices accordingly but manufacturers set prices equal to the destination-specific marginal costs. Comparing this setup to the baseline model is useful as this simplified version of the model is isomorphic to the case in which retailers and manufacturers jointly maximize surplus through Nash bargaining but retailers hold all the bargaining power. At the same time, the baseline case of double marginalization is isomorphic to the Nash bargaining in which manufacturers hold all the bargaining power. Considering the two polar outcomes of the situation in which retailers and manufacturers jointly maximize surplus should provide reassurance that the recovered trade frictions are robust to the assumption of separate profit maximization. Second, I consider the case in which upstream manufacturers compete à la Bertrand but retailers set prices equal to the sum of wholesale prices and local distribution costs. Considering whether the results of robust to this variation is useful as it would show that the results are robust to assuming that retailers perfectly pass-through costs as in Corsetti & Dedola (2005) and Loecker & Scott (2022) and that the results are robust to certain forms of non-linear contracting and interlocking relationships that give rise to the same outcome (see Rey & Vergé (2010))

5 Structural estimation

In this section, I start by discussing the estimation and identification of the demand parameters. Conditional on these parameters, I consider the estimation and identification of the cost parameters and the cross-border trade frictions in particular.

5.1 Estimation of preferences

Identification To estimate consumer preferences for bottled water, I rely on equation 3 and apply the inversion from Berry (1994) to obtain the following estimating equation:

$$\ln\left(\frac{\sigma_{j,lt}}{\sigma_{0,lt}}\right) = \alpha P_{j,lt} + \theta_{b(j),n(l)} + \theta_{c(j),t} + \lambda_{lt} + \xi_{j,lt}. \tag{7}$$

where $\sigma_{0,lt}$ is the market share of the outside good, which corresponds with the probability of buying other non-alcoholic beverages or tap water. The inclusion of the market fixed effects, λ_{lt}, however, subsumes the variation in relative market shares that stems from variation in the outside option. I operationalize this equation by equating $\sigma_{j,lt}$ to the observed probability of purchasing variety j, given by $S_{j,lt}$. I capture non-price product characteristics by including different sets of fixed effects at the brand $\theta_{b(j),n(l)}$ and retail chain level $\theta_{b(j),n(l)}$.

An important empirical challenge lies in consistently estimating α which governs the sensitivity of demand to changes in prices. This challenge arises because $\xi_{j,lt}$ is unobserved and equilibrium prices are likely determined with knowledge of $\xi_{j,lt}$, creating a typical simultaneity issue. To address this challenge, I leverage the information about production locations and instrument consumer prices with the interaction of the distance between the production location and the destination market and diesel prices from the origin location.[23] The suitability of transport costs as an instrument for consumer prices is that it affects consumer prices by affecting the marginal cost of delivering the variety to the destination. At the same time, conditional on brand fixed effects it is unlikely that variation in transport costs is correlated with the demand shifters, especially once I allow brand appeal to vary between countries. In this way, the demand-side parameters are identified separately from the model of market structure and marginal costs.

Estimation results Table 4 reports the results from estimating Equation with OLS and with 2SLS using transport costs as the instrument for consumer prices. The estimations are based on the main sample where a product is the interaction of a store and a bottled water variety. I obtain the prices and probabilities of purchasing by aggregating transactions across the households by using population weights interacted with the number of purchased liters as weights. Alongside the coefficient estimated, I report clustered standard errors at the level of the regional destination markets.

First, Column (1) presents the OLS results with brand and market fixed effects. In this case, I recover a precisely estimated negative price coefficient of -1.76. This estimate implies a rather inelastic quantity-weighted own-price elasticity of -0.35. When I allow for the preference for certain brands to vary between countries or when I include retail chain time fixed effects, the price coefficient is virtually unchanged. So, whereas the OLS estimates yield downward-sloping demand curves, the implied average own-price elasticities are inelastic.

Second, as mentioned before the OLS results likely suffer from a simultaneity issue due to

---

[23] The instrument is equally powerful when I use diesel prices in the destination market and the estimated coefficients are also quantitatively very similar.

unobserved demand shocks. Therefore, I turn to the IV results in columns (4) to (6). As in column (1), column (4) reports the results with brand and market fixed effects but now instruments consumer prices with transport costs. The first stage F-statistics are large and confirm that, in line with Figure 3, transport costs co-determine consumer prices significantly. After instrumenting prices with transport costs, the estimated demand curves are much more elastic. Column (4) shows that the estimated price coefficient falls to -22.4, remains precisely estimated and implies a quantity-weighted own-price elasticity of demand of -4.47. The estimated coefficient is almost unchanged when I allow brand appeal to differ by the destination country in column (5) and slightly increases to -20.2 when I also include retail chain-time fixed effects. In this case, the implied quantity-weighted own-price elasticity of demand equals -4.03.

Table 4: Estimation of preferences

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(P_{j,lt})$	-1.76^{***}	-1.75^{***}	-1.76^{***}	-22.4^{***}	-22.3^{***}	-20.2^{***}
	(0.184)	(0.185)	(0.184)	(1.640)	(1.190)	(1.310)
1^{st} Stage F-stat	-	-	-	4,287.6	5,633.6	6,040.6
$\mathbb{E}\left[\varepsilon_{j,lt}\right]$	-0.35	-0.35	-0.35	-4.47	-4.46	-4.03
$\theta_{b(j)}$	✓			✓		
$\theta_{b(j),n(l)}$		✓	✓		✓	✓
$\lambda_{c,t}$			✓			✓
λ_{lt}	✓	✓	✓	✓	✓	✓
N	952,970	952,970	952,970	786,735	786,735	786,735

Notes: This table presents the results from estimating on variety-level data. I obtain the variety-level prices and probabilities of purchasing by aggregating transactions across the household and store level using population weights. Columns (1) and (2) present estimates using the OLS estimator. Columns (3) and (4) present the results using the 2SLS estimator where I instrument variety-level prices with the interaction of diesel prices and the distance between the destination location and the production location. The reported F-statistic is the cluster F-statistic from Montiel-Olea & Pflueger (2013). $\mathbb{E}\left(\varepsilon_{j,lt}\right)$ is the unweighted average over the product- and market-specific own-price elasticities, which are given by $\varepsilon_{j,lt} = \alpha P_{j,lt}(1 - S_{j,lt})$. Alongside the coefficient estimates, I report cluster-robust standard errors computed at the regional market level and significance levels at the 0.1^{*}, 0.05^{**}, 0.01^{***} levels.

Implied markups To provide some validation for the estimated own- and cross-price elasticities of demand, I benchmark the implied manufacturing markups to the literature. Figure **??** plots these implied manufacturing markups, defined as the price-to-marginal cost ratio. Based on the estimate of column (6) of Table 4, Figure **??** shows that the estimated own- and cross-price elasticities yield a quantity-weighted markup of around 1.68. To put this number in perspective, Loecker et al. (2016) recover a median markup of 1.34 from production data across manufacturing firms in India. Relative to the extensive literature that relies framework that imposes CES-demand and monopolistic competition, the same price-cost ratio would arise with an elasticity of substitution that is a little lower than 5. In contrast to this literature, Figure **??** shows that there is substantial heterogeneity around this mean of 1.68 because markups are allowed to co-vary with the environment in each of the destination markets.

5.2 Estimation of cross-border frictions

Armed with the estimated demand parameters, I am in a position to back out the distribution of destination-specific marginal costs from first-order conditions of retailers and manufacturers as spelled in Equation 6. I use this distribution to estimate the level of cross-border trade frictions in the EU bottled water industry.

Identification Estimating the level of trade frictions requires parameterizing the marginal cost function. In particular, I assume that the destination-specific marginal cost function takes the following form:

$$\boldsymbol{c}_{lt}^{r} + \boldsymbol{c}_{lt}^{m} = \beta W_{lt} + \lambda_{c(j),t} + \gamma_t t_{s(j)lt} + \tau_B \mathbb{1}\left(\text{Border}\right)_{s(j)l} + \tau_C \mathbb{1}\left(\Delta\text{Currency}\right)_{s(j)l} + \omega_{j,t} + \eta_{j,lt} \quad (8)$$

This cost function accounts for three sources of destination-specific marginal costs. First, destination-specific marginal costs may differ between destination markets when local distribution costs differ (e.g. Burstein et al. (2005); Crucini et al. (2005); Parsley & Wei (2007)). I capture differences in distribution costs by including W_{lt} which are labor unit costs in the retail sector, expressed in Euro per hour, and which are allowed to vary over time and by destination market. In addition, I include $\lambda_{c(j),t}$ which are chain-time fixed effects and which account for differences in retail costs between retailers over time. The second source of differences in destination-specific marginal costs is the fact that transport costs, given by $t_{s(j)l,t}$, differ for markets at different distances from the water source. As mentioned before, I account for differences in transport costs by including the travel distance between the origin and source ZIPcodes by truck interacted with diesel prices and the travel time between the origin and source ZIPcodes interacted with trucker wages. Finally, destination-specific marginal costs differ when there are positive cross-border trade frictions $\tau_B \mathbb{1}\left(\text{Border}\right)_{s(j)l}$ and $\tau_C \mathbb{1}\left(\Delta\text{Currency}\right)_{s(j)l}$ for Eurozone and outside-of-Eurozone borders respectively. Besides these three sources of destination-specific marginal costs, the cost function captures two additional sources of variation: production costs are captured through $\omega_{j,t}$ and finally $\eta_{j,lt}$ which is an unobserved variety- and market-specific cost shifter that varies over time.

Two features of this parameterization stand out. First, having data on variety-level production locations and observing the same variety being consumed in multiple locations allows me to fully and flexibly account for the marginal costs of production through variety-time fixed effects. For this reason, this specification allows for very flexible substitution patterns on the production side and allows me to recover cross-border trade frictions under arbitrary returns to scale and scope. Second, to arrive at this specification, I have implicitly that the production function of bottled water takes the Leontief form in which water bottles, local labor and transportation costs all enter as perfect complements. In my setting, these assumptions feel quite natural as both transport and distribution are indispensable to be able to deliver bottled water to the final consumer. Moreover, assuming that distribution costs and transport costs enter the cost function additively are standard in the international trade literature (see Corsetti & Dedola (2005); Parsley & Wei (2007); Irarrazabal et al. (2015)).

Consistent estimation of the cross-border frictions requires that the time and location-specific cost shifter, $\eta_{j,lt}$ is uncorrelated with $\mathbb{1}\left(\text{Border}\right)_{s(j)l}$ and $\mathbb{1}\left(\Delta\text{Currency}\right)_{s(j)l}$. Two potential violations come to mind. First, there is a growing literature in international trade that documents variation in export prices across different destination markets (see e.g. Manova & Zhang (2012)). One explanation is that the same exporters export products of different quality to different countries by using an input mix that is contingent on the destination market (see Bastos et al. (2018)). Price variation due to differences in unobserved product characteristics is unlikely to drive my results as bottled water has a limited number of dimensions of product differentiation and because I define products in such a way that all dimensions of product differentiation that co-determine prices are accounted for (see Table A.5). A second source of destination-specific marginal costs is when exporters use different modes of transportation for different destination markets (e.g. Harrigan (2010) and D. L. Hummels & Schaur (2013)). Because foodstuffs, including bottled water, are shipped on trucks and because I consider a set of geographically close countries, differences in costs due to differences in the mode of transportation are inconceivable.

Estimation results Table 5 presents the results from estimation equation 5 under different assumptions about how consumer prices are determined in equilibrium. Each specification includes product-time fixed effects chain time fixed effects. Below the coefficient estimates, I report standard errors which are clustered at the level of the regional destination market.

Columns (1) and (2) show the results when I implement Equation 5 using final consumer prices. This coincides with the case in which both retailers and manufacturers would set prices equal to marginal costs, for instance when both the upstream and downstream markets would be perfectly competitive. Column (1) considers a specification without including any of the border dummies and therefore can be interpreted as the first-stage regression of the demand estimation. Final consumer prices intuitively rise when transport costs and local distribution costs rise. Adding a border dummy in column (2) indicates that consumer prices are roughly 8.6 eurocents higher abroad compared to the domestic price.

Columns (3) to (6) contain the results when I allow for deviations from perfect competition either in the upstream or downstream market. First, Column (3) considers the same specification as column (1) assuming Bertrand competition both in the upstream and downstream market. Column (3) confirms that consumer prices increase with transport costs and local distribution costs because marginal costs are in those and not because markups are higher. This result supports the identification strategy of the demand parameters. Second, columns (4) to (6) highlight that accounting for spatial price discrimination increases the estimated tariff-equivalent trade friction.[24] In particular, I

---

[24] This result aligns well with the pricing-to-market literature (e.g. Atkeson & Burstein (2008); Fitzgerald & Haller (2014); Corsetti et al. (2021)). As foreign firms are at a relative cost disadvantage due to higher transport costs and/or cross-border trading frictions, their optimal markup is lower compared to the one at home.

estimate that the marginal cost of supplying the same product to a foreign market increases by a little over 9 eurocents. This result is consistent across the different vertical market structures: upstream oligopoly (column (4)), downstream oligopoly (column (5)) and double marginalization (column (6)). Given that the domestic marginal cost of exported varieties is on average between 33 and 38 eurocents, depending on the vertical market structure, this implies tariff-equivalent trade frictions between 23 and 28%. These estimates are somewhat larger compared to the RDD estimates from Table 2 but on the lower end of the gravity estimates from Table 3.

Finally, column (7) replicates the analysis from column (6) but also investigates heterogeneity in the trade frictions by distinguishing between trade barriers between country pairs that are both part of the Eurozone and country pairs where at least one country is not part of the Eurozone. On the one hand, column (7) confirms that substantial trade frictions exist between countries that are both part of the Eurozone as the estimated difference in the marginal cost between foreign and domestic markets remains high at a little over 7 eurocents. On the other hand, column (7) also uncovers substantial heterogeneity in the estimate of cross-border trade frictions. When at least one country is not part of the Eurozone, the trade friction rises by 23 eurocents.

Table 5: Estimation of cost function

	$p^r_{j,lt}$		$mc^w_{j,lt} + mc^r_{j,lt}$				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Border$^{s(j),l}$	-	0.0864***	-	0.0901***	0.088***	0.0918***	0.0712***
	-	(0.007)	-	(0.007)	(0.007)	(0.007)	(0.007)
Cur$^{s(j),l}$	-	-	-	-	-	-	0.233***
	-	-	-	-	-	-	(0.038)
$t^{s(j),l}_t$	0.00849**	0.00247	0.00901**	0.00307	0.002	0.00261	0.000486
	(0.003)	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	(0.001)
w_{lt}	0.0151***	0.00818***	0.0156***	0.00818***	0.00828***	0.00828***	0.0138***
	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)	(0.001)	(0.001)
$E\left[\cdot\middle\vert\text{Export}_j = 1\right]$	0.44	0.44	0.33	0.38	0.38	0.33	0.33
τ_B	-	0.20	-	0.24	0.23	0.28	0.22
τ_C	-	-	-	-	-	-	0.58
$\mu^r_{j,lt}$			✓		✓	✓	✓
$\mu^w_{j,lt}$			✓	✓		✓	✓
$\omega_{j,t}$	✓	✓	✓	✓	✓	✓	✓
$\lambda_{c(j),t}$	✓	✓	✓	✓	✓	✓	✓
R^2	0.92	0.93	0.92	0.93	0.93	0.93	0.93
N	645,227	645,227	645,227	645,227	645,227	645,227	645,227

Notes: This table presents a decomposition of the border effects for consumer prices into a border effect for marginal costs and manufacturing markups. Columns (1) and (2) present the estimates for consumer prices. Columns (3) and (4) present the results for marginal costs where marginal costs are obtained from backing them out using Equation **??** using the estimates from column (4) of Table 4. These columns thus represent the results from estimating . Columns (5)-(6) present the results for manufacturing markups. All results are computed based on variety-level prices which are obtained by aggregating transactions across the household and store level using population weights. In addition, I report the conditional unweighted average of the dependent variable where the conditioning pertains to only including varieties that are exported and only including observations in the home country of that respective variety. Alongside the coefficient estimates, I report cluster-robust standard errors computed at the regional destination market level and significance levels at the 0.1^{*}, 0.05^{**}, 0.01^{***} levels.

Taking stock The results presented in Table 5 have two main implications. On the one hand, the fact that cross-border frictions are lower between country pairs part of the Eurozone compared to pairs of countries where at least one of them is outside of the Eurozone provides credence to models that consider the reduction in trading frictions between countries with a different currency as a key source of welfare gains. On the other hand, the fact that cross-border frictions between countries of the European Single Market and the Eurozone are also quite large suggests that the potential welfare gains of further geographic market integration within the Eurozone are potentially large as well. This is the subject of section 6.

6 Equilibrium effects of market segmentation

[VERY MUCH UNDER CONSTRUCTION, TAKE WITH A GRAIN OF SALT]

In this section, I explore the effect of trade frictions on equilibrium price dispersion and consumer welfare through a counterfactual exercise. Doing so, I compare the current level of equilibrium price dispersion to the level of equilibrium price dispersion that would arise in an integrated benchmark economy. In addition, I compute how much lower consumer welfare is relative to this integrated benchmark economy.

Integrated benchmark I define the integrated benchmark economy as an economy in which the current set of available varieties is supplied at the marginal cost of supplying those particular varieties to that particular market. This marginal cost comprises two separate components. First, there is the marginal costs of production. Relative to the estimation of the cost function, I require additional assumptions on the marginal cost of production. So far, the assumptions on the marginal costs of production were minimal as arbitrary returns to scale were flexibly accounted for through product-time fixed effects (see Equation 5). In the absence of trade frictions in the efficient benchmark, the counterfactual scale of production might differ from the one observed in the data. However, under the assumption of constant returns to scale, the scale of production does not affect the marginal cost of production and the current estimates, i.e. the product-time fixed effects $\lambda_{j,t}$ from Equation 5, can be used to construct marginal cost of supplying different markets.

Second, whereas the marginal cost of production is the same across destination markets, the costs of delivering bottled water to different destination markets are not the same. As shown in Table 5, transport costs to the destination market and local distribution costs in terms of labor will lead to non-zero spatial price dispersion even in the integrated benchmark economy.

Price dispersion and welfare Table 6 compares the counterfactual levels of equilibrium cross-country price dispersion between the integrated economy and the segmented economy in which varieties are supplied at their marginal costs but supplying them is also subject to the estimated trade frictions. Two things stand out. First, cross-border trade frictions increase

Table 6: Counterfactual excercises

			$\mathbb{E}\left[\lvert p_{j,lt} - p_{j,kt}\rvert n(k) \neq n(j)\right]$		
Counterfactual	$\boldsymbol{\tau}$	$\boldsymbol{\mu}$	Level	Change	ΔCS
Integrated economy	0	0	34%	-	-
Segmented - No Market power	$\hat{\boldsymbol{\tau}}$	0	39%	+ 5%	- 0.036 EUR/L
Segmented - No price discrimination	$\hat{\boldsymbol{\tau}}$	$\boldsymbol{\mu}_{j,t}(\hat{\boldsymbol{\tau}})$		Loading ...	
Segmented - Price discrimination	$\hat{\boldsymbol{\tau}}$	$\boldsymbol{\mu}_{j,lt}(\hat{\boldsymbol{\tau}})$		Loading ...	

between-country spatial price dispersion. This is because equilibrium between-country spatial price dispersion is 39% in the segmented economy and only 34% in the integrated economy. Second, even in the integrated benchmark economy, price dispersion is considerable as it is on average 34% among international region pairs. This has two implications. On the one hand, this implies that even though cross-border trade frictions would be removed in the integrated benchmark, one can expect significant equilibrium spatial price dispersion to remain purely due to differences in destination-specific marginal costs. On the other hand, the substantial equilibrium between-country spatial price dispersion in the integrated benchmark economy highlights that the existence of spatial price dispersion at market boundaries per se is not a reason for the existence of trade frictions and geographic market segmentation.

To quantify the effect of cross-border frictions on consumer welfare, I compute the difference in consumer welfare between the integrated benchmark and the segmented economy in which goods are delivered at the destination-specific marginal costs, inclusive of the estimated cross-border trade frictions. Doing so, I compute the percentage change in consumer surplus in the segmented economy relative to the integrated benchmark. Table 6 shows that the expected consumer welfare from purchasing water in the segmented economy is 3.6 eurocents lower compared to the segmented economy. Whereas this number is small in absolute terms, given that the average consumer price is around 37 eurocents per liter, this constitutes roughly a 10% tax on bottled water consumption.

7 Conclusion

Spatial price dispersion is commonplace both within and between countries. Whether such spatial price dispersion, especially at market boundaries, reflects geographic market segmentation and should be reduced, crucially depends on the origins of such price dispersion. When spatial price dispersion only reflects differences in the marginal cost of supplying different markets, equilibrium spatial price dispersion will arise even in the integrated benchmark. Only when spatial price dispersion is driven by trade frictions will it be reflective of market segmentation and will reducing it be welfare improving.

This paper examines whether trade between members of the European Single Market is still hampered by trade frictions and if so whether these frictions induce inefficient price dispersion and subdued consumer welfare. I make progress on this question by focusing on the bottled water industry where I leverage institutional features to construct a unique dataset on prices, quantities and production locations at the variety level. I combine these data with an empirical model of demand and supply in the bottled water industry and identify trade frictions from cross-border variation in model-implied marginal costs while allowing for between-country differences in preferences, market structure and arbitrary returns to scale in production.

I estimate that within the Eurozone cross-border trading frictions are around 9 cents per liter and that they are 32 cents per liter when trade happens outside of the Eurozone. Equilibrium price dispersion is 5% higher compared to an integrated benchmark in which bottled water is delivered at the marginal cost of supplying different markets and its welfare cost is equivalent to a 10% tax on bottled water consumption.

My results have two implications. First, the presence of trade trading frictions likely has implications beyond the bottled water industry. While production locations are fixed in the bottled water industry, the presence of cross-border frictions likely distorts the allocation of production units across European countries in industries with more flexibility in terms of plant location. Instead of fully exploiting returns to scale, such industries might opt for more dispersed (and more costly) production closer to the final market.

Second, I interpret the residual differences in model-implied destination-specific marginal costs as cross-border trading frictions. While the empirical model is more flexible in terms of allowing differences in preferences, market structure and production technology relative to standard trade models, the cross-border frictions are a construct of the empirical model to rationalize the data within the model. Crucially such price differences can only be sustained because consumers do not engage in profitable arbitrage. Whereas arbitrage through cross-border shopping is not very profitable in the bottled water industry due to its low value-to-weight ratio, an important outstanding question is which economic primitives generate such limited arbitrage and large equilibrium price dispersion for more valuable products as well.

References

Allen, T., Arkolakis, C., & Takahashi, Y. (2020). Universal gravity. *Journal of Political Economy*, *128*, 393-433.

Anderson, J. E., & Wincoop, E. V. (2003). Gravity with gravitas: A solution to the border puzzle. *American Economic Review*, *93*, 170-192.

Andersson, D., Berger, T., & Prawitz, E. (2023, 3). Making a market: Infrastructure, integration, and the rise of innovation. *The Review of Economics and Statistics*, *105*, 258-274. doi: 10.1162/rest_a_01067

Antràs, P., Fadeev, E., Fort, T. C., & Tintelnot, F. (2024). Exporting, global sourcing and multinational activity: Theory and evidence from the united states. *Review of Economics and Statistics*, *Forthcoming*, 1-45.

Armington, P. S. (1969). A theory of demand for products distinguished by place of production. *IMF Staff Papers*, *16*, 159-178.

Asplund, M., & Friberg, R. (2001). The law of one price in scandinavian duty-free stores. *American Economic Review*, *91*, 1072-1083.

Atkeson, A., & Burstein, A. (2008). Pricing-to-market, trade costs, and international relative prices. *American Economic Review*, *98*, 1998-2031.

Atkin, D., & Donaldson, D. (2015). *Who's getting globalized? the size and implications of intranational trade costs.*

Auer, R., Burstein, A., Lein, S., & Vogel, J. (2023). Unequal expenditure switching: Evidence from switzerland. *The Review of Economic Studies*, *Forthcoming*, 1-35.

Bastos, P., Silva, J., & Verhoogen, E. (2018, 2). Export destinations and input prices. *American Economic Review*, *108*, 353-392.

Beck, G. W., Kotz, H. H., & Zabelina, N. (2020, 11). Price gaps at the border: Evidence from multi-country household scanner data. *Journal of International Economics*, *127*, 1033-1068.

Berry, S. (1994). Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, *25*, 242-262.

Berry, S., Levinsohn, J., & Pakes, A. (1999). Voluntary export restraints on automobiles: Evaluating a trade policy. *The American Economic Review*, *89*, 400-431.

Boehm, C. E., Levchenko, A. A. A., & Pandalai-Nayar, N. N. (2023). The long and short (run) of trade elasticities. *American Economic Review*, *Forthcoming*, 1-46.

Bonnet, C., & Dubois, P. (2010). Inference on vertical contracts between manufacturers and retailers allowing for nonlinear pricing and resale price maintenance. *RAND Journal of Economics*, *41*, 139-164.

Bourreau, M., Sun, Y., & Verboven, F. (2021, 11). Market entry, fighting brands, and tacit collusion: Evidence from the french mobile telecommunications market. *American Economic Review*, *111*, 3459-3499. doi: 10.1257/aer.20190540

Burstein, A., Eichenbaum, M., & Rebelo, S. (2005). Large devaluations and the real exchange rate. *Journal of Political Economy*, *113*, 742-784.

Calonico, S., Cattaneo, M. D., & Titiunik, R. (2014, 11). Robust nonparametric confidence intervals for regression-discontinuity designs. *Econometrica*, *82*, 2295-2326. doi: 10.3982/ecta11757

Cavallo, A., Neiman, B., & Rigobon, R. (2014). Currency unions, product introductions, and the real exchange rate. *The Quarterly Journal of Economics*, *129*, 529-595.

Chatterjee, S. (2023, 8). Market power and spatial competition in rural india. *The Quarterly Journal of Economics*, *138*, 1649-1711. doi: 10.1093/qje/qjad004

Corsetti, G., Crowley, M., Han, L., & Song, H. (2021). Markets and markups: A new empirical framework and evidence on exporters from china. *mimeo*.

Corsetti, G., & Dedola, L. (2005, 9). A macroeconomic model of international price discrimination. *Journal of International Economics*, *67*, 129-155.

Cotterlaz, P., & Vicard, V. (2023). *Why origin matters in trade data* working paper cepii working paper why origin matters in trade data.*

Coşar, A. K., Grieco, P. L., Li, S., & Tintelnot, F. (2018). What drives home market advantage? *Journal of International Economics*, *110*, 135-150.

Coşar, K., Grieco, P., & Tintelnot, F. (2015, 7). Borders, geography, and oligopoly: Evidence from the wind turbine industry. *The Review of Economics and Statistics*, *97*, 623-637.

Crucini, M. J., Telmer, C. I., & Zachariadis, M. (2005). Understanding european real exchange rates. *American Economic Review*, *95*, 724-738.

Donaldson, D. (2018, 4). Railroads of the raj: Estimating the impact of transportation infrastructure. *American Economic Review*, *108*, 899-934.

Dvir, E., & Strasser, G. (2018, 5). Does marketing widen borders? cross-country price dispersion in the european car market. *Journal of International Economics*, *112*, 134-149.

Eaton, J., & Kortum, S. (2002). Technology, geography and trade. *Econometrica*, *70*, 1741-1779.

Edmond, C., Midrigan, V., & Xu, D. Y. (2015). Competition, markups, and the gains from international trade. *American Economic Review*, *105*, 3183-3221.

Engel, C., & Rogers, J. H. (1996). How wide is the border? *American Economic Review*, *86*, 1112-1125.

Fajgelbaum, P., Grossman, G., & Helpman, E. (2011, 8). Income distribution, product quality, and international trade. *Journal of Political Economy*, *119*, 721-765.

Fitzgerald, D., & Haller, S. (2014). Pricing-to-market : Evidence from plant-level prices. *The Review of Economic Studies*, *81*, 761-786.

Fontaine, F., Martin, J., & Mejean, I. (2020, 5). Price discrimination within and across emu markets: Evidence from french exporters. *Journal of International Economics*, *124*, 1-19.

Goldberg, P. K. (1995). Product differentiation and oligopoly in international markets : The case of the us automobile industry. *Econometrica*, *63*, 891-951.

Goldberg, P. K., & Knetter, M. M. (1997). Goods prices and exchange rates: What have we learned? *Journal of Economic Literature*, *35*, 1243-1272.

Goldberg, P. K., & Verboven, F. (2001). The evolution of price dispersion in the european car market. *The Review of Economic Studies*, *68*, 811-848.

Gopinath, G., Gourinchas, P. O., Hsieh, C. T., & Li, N. (2011). International prices, costs, and markup differences. *American Economic Review*, *101*, 2450-2486.

Handbury, J. (2021). Are poor cities cheap for everyone? non-homotheticity and the cost of living across u.s. cities. *Econometrica*, *89*, 2679-2715. doi: 10.3982/ecta11738

Harrigan, J. (2010, 11). Airplanes and comparative advantage. *Journal of International Economics*, *82*, 181-194. doi: 10.1016/j.jinteco.2010.07.002

Head, K., & Mayer, T. (2014). Gravity equations: Workhorse,toolkit, and cookbook. In (Vol. 4, p. 131-195). Elsevier B.V.

Head, K., & Mayer, T. (2021). The united states of europe : A gravity model evaluation of the four freedoms. *Journal of Economic Perspectives*, *35*, 23-46.

Hellerstein, R. (2008, 9). Who bears the cost of a change in the exchange rate? pass-through accounting for the case of beer. *Journal of International Economics*, *76*, 14-32.

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

Helpman, E., Melitz, M. J., & Yeaple, S. R. (2004). Export versus fdi with heterogeneous firms. *American Economic Review*, *94*, 300-316.

Hendel, I., & Nevo, A. (2006). Measuring the implications of sales and consumer inventory behavior. *Econometrica*, *74*, 1637-1673.

Hornbeck, R., & Rotemberg, M. (2024). Growth off the rails: Aggregate productivity growth in distorted economies. *Journal of Political Economy*, 1-34.

Hristakeva, S. (2022, 12). Vertical contracts with endogenous product selection: An empirical analysis of vendor allowance contracts. *Journal of Political Economy*, *130*, 3202-3252.

Hsieh, C.-T., & Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly Journal of Economics*, *124*, 1403-1448.

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

Hummels, D. L., & Schaur, G. (2013). Time as a trade barrier. *American Economic Review*, *103*, 2935-2959.

Irarrazabal, A., Moxnes, A., & Opromolla, L. D. (2015). The tip of the iceberg: A quantitative framework for estimating trade costs. *Review of Economics and Statistics*, *97*, 777-792.

Kalouptsidi, M. (2018, 4). Detection and impact of industrial subsidies: The case of chinese shipbuilding. *The Review of Economic Studies*, *85*, 1111-1158.

Krugman, P. (1980). Scale economies, product differentiation, and the pattern of trade. *American Economic Review*, *70*, 950-959.

Loecker, J. D. (2011). Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica*, *79*, 1407-1451.

Loecker, J. D., Goldberg, P. K., Khandelwal, A. K., & Pavcnik, N. (2016). Prices, markups, and trade reform. *Econometrica*, *84*, 445-510.

Loecker, J. D., & Scott, P. T. (2022). *Markup estimation using production and demand data. an application to the us brewing industry *.*

Manova, K., & Zhang, Z. (2012, 2). Export prices across firms and destinations. *The Quarterly Journal of Economics*, *127*, 379-436.

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

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, *71*, 1695-1725.

Miller, N. H., & Weinberg, M. C. (2017). Understanding the price effects of the millercoors joint venture. *Econometrica*, *85*, 1763-1791.

Molina, H. (2021). *Buyer alliances in vertically related markets.*

Montiel-Olea, J. L., & Pflueger, C. (2013, 7). A robust test for weak instruments. *Journal of Business and Economic Statistics*, *31*, 358-369.

Nakamura, E., & Zerom, D. (2010). Accounting for incomplete pass-through. *The Review of Economic Studies*, *77*, 1192-1230.

Parsley, D. C., & Wei, S.-J. (2007). A prism into the ppp puzzles: The micro-foundations of big mac real exchange rates. *The Economic Journal*, *117*, 1336-1356.

Rey, P., & Vergé, T. (2010, 12). Resale price maintenance and interlocking relationships. *Journal of Industrial Economics*, *58*, 928-961. doi: 10.1111/j.1467-6451.2010.00439.x

Santamaria, M. A., Ventura, J., & Yesilbayraktar, U. (2020). Borders within europe. *NBER Working Paper Series*, *1-87*, 1689-1699.

Santamaría, M., Ventura, J., & Yeşilbayraktar, U. (2023). Exploring european regional trade. *Journal of International Economics*.

Shiue, C. H., & Keller, W. (2007). Markets in china and europe on the eve of the industrial revolution. *The American Economic Review*, *97*, 1189-1216.

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

Tintelnot, F. (2016). Global production with export platforms. *The Quarterly Journal of Economics*, *132*, 157-209.

Verboven, F. (1996). International price discrimination in the european car market. *The RAND Journal of Economics*, *27*, 240-268.

A Data

A.1 Retail stores

Table A.1: Cross-border shopping

Country	All	BEL	GER	DEN	FRA	NLD	PLN	SWE	UK
Transactions (Count)									
·Domestic	15,027,795	859,605	5,120,630	90,889	3,012,993	1,005,842	1,965,480	134,482	2,837,874
·Cross-border	96,685	41,629	10,390	502	15,241	28,743	98	1	81
·Undisclosed	3,838	150	0	2,541	830	7	0	281	29
Transactions									
·Domestic	0.993	0.954	0.998	0.968	0.995	0.972	1.000	0.998	1.000
·Cross-border	0.006	0.046	0.002	0.005	0.005	0.028	0.000	0.000	0.000
·Undisclosed	0.000	0.000	0.000	0.027	0.000	0.000	0.000	0.002	0.000
Liters (liters)									
·Domestic	109,359,137	7,616,084	51,753,208	291,570	27,055,500	3,529,122	9,845,240	312,938	8,955,475
·Cross-border	1,246,252	587,835	187,959	5,249	299,837	164,323	609	0	440
·Undisclosed	18,344	721	0	9,893	7,253	4	0	419	54
Liters									
·Domestic	0.989	0.928	0.996	0.951	0.989	0.956	1.000	0.999	1.000
·Cross-border	0.011	0.072	0.004	0.017	0.011	0.044	0.000	0.000	0.000
·Undisclosed	0.000	0.000	0.000	0.032	0.000	0.000	0.000	0.001	0.000
Price (EUR/L)									
·Domestic	0.431	0.380	0.251	0.618	0.285	0.456	0.210	0.849	0.396
·Cross-border	0.298	0.306	0.293	0.665	0.191	0.301	0.158	0.000	0.470
·Undisclosed	0.683	0.636	0.000	0.646	0.333	2.000	0.000	1.182	0.664
Exp. share - NARTD									
·Unconditional	0.098	0.253	0.281	0.085	0.312	0.110	0.363	0.097	0.098
·Conditional	0.252	0.318	0.397	0.243	0.377	0.217	0.424	0.318	0.252
Exp. share - All									
·Unconditional	0.003	0.015	0.017	0.003	0.019	0.005	0.012	0.005	0.003
·Conditional	0.008	0.019	0.024	0.008	0.022	0.010	0.014	0.018	0.008

Notes: This table provides an overview of the extent of cross-border shopping in the data. For each country and overall, I show the total number of transactions and the share in the number of transactions that involve consumption of bottled water either in a domestic store, in a store in an adjecent country or in an undisclosed store. I also show the total number of liters and the share of in the number of liters and the average price liters weigthed by the number of liters times the population weights across each of these three groups of stores. I also show the average expenditure share on water in terms of the non-alcoholic ready-to-drink (NARTD) category and in total grocery expenditure, once for all consumers ("unconditional") and once for the group of household that purchase bottled water ("conditional"). Beyond bottled water, the NARTD category includes sodas, energy drink, juices and ice teas. To compute the shares and the average price per liter results in the column "All" we compute simple average across countries.

Table A.2: Sample selection: Stores

Country	All	BEL	GER	DEN	FRA	NLD	PLN	SWE	UK
Transactions (count)									
·Included	15,699,510	840,591	4,522,146	89,901	2,906,213	991,998	1,650,407	130,217	2,704,314
·Excluded	1,365,408	60,793	608,874	4,031	122,851	42,594	315,171	4,547	133,670
Transactions (share)									
·Included	0.920	0.933	0.881	0.957	0.959	0.959	0.840	0.966	0.953
·Excluded	0.080	0.067	0.119	0.043	0.041	0.041	0.160	0.034	0.047
Liters (liters)									
·Included	100,852,157	7,438,186	44,241,841	289,649	26,376,079	3,490,325	8,832,911	306,676	8,657,338
·Excluded	10,522,905	766,454	7,699,327	17,062	986,510	203,123	1,012,939	6,682	298,632
Liters (share)									
·Included	0.906	0.907	0.852	0.944	0.964	0.945	0.897	0.979	0.967
·Excluded	0.094	0.093	0.148	0.056	0.036	0.055	0.103	0.021	0.033

Notes: This table provides an overview of the extent of cross-border shopping in the data. For each country and overall, I show the total number of transactions and the share in the number of transactions that involve the consumption of bottled water either in a domestic store, in a store in an adjacent country or in an undisclosed store. I also show the total number of liters and the share of the number of liters and the average price liters weighted by the number of liters times the population weights across each of these three groups of stores. I also show the average expenditure share on water in terms of the non-alcoholic ready-to-drink (NARTD) category and in total grocery expenditure, once for all consumers ("unconditional") and once for the group of households that purchase bottled water ("conditional"). Beyond bottled water, the NARTD category includes sodas, energy drinks, juices and ice teas. To compute the shares and the average price per liter results in the column "All" we compute simple averages across countries.

A.2 Product-level consumption data

Table A.3: Product characteristics

Variable	BEL	GER	DEN	FRA	NLD	PLN	SWE	UK
Brand type								
·Branded	0.626	0.637	0.740	0.802	0.563	0.672	0.694	0.449
·White label	0.374	0.342	0.249	0.197	0.437	0.328	0.294	0.546
·Undisclosed	0.000	0.021	0.010	0.001	0.000	0.000	0.011	0.005
Water type								
·Still	0.654	0.250	0.278	0.649	0.447	0.600	0.052	0.690
·Sparkling	0.346	0.710	0.722	0.349	0.553	0.400	0.948	0.309
·Undisclosed	0.000	0.041	0.000	0.002	0.000	0.000	0.000	0.001
Flavor								
·Unflavored	0.940	0.945	0.484	0.945	0.509	0.810	0.303	0.598
·Flavored	0.060	0.054	0.516	0.053	0.491	0.190	0.697	0.402
·Undisclosed	0.000	0.000	0.000	0.002	0.000	0.000	0.000	0.001
Bottle size								
· $\leq$ 750ml	0.276	0.311	0.361	0.116	0.317	0.139	0.431	0.500
·(750ml-1500ml)	0.230	0.303	0.225	0.383	0.094	0.010	0.009	0.214
· $\geq$ 1500ml	0.494	0.386	0.406	0.478	0.589	0.851	0.543	0.286
·Undisclosed	0.000	0.000	0.008	0.024	0.000	0.000	0.017	0.000
Bottle package								
·Plastic	0.940	0.780	0.000	0.976	0.839	0.992	0.000	0.990
·Glass	0.056	0.196	0.000	0.008	0.002	0.005	0.000	0.008
·Other	0.004	0.001	0.000	0.011	0.159	0.003	0.000	0.002
·Undisclosed	0.000	0.024	1.000	0.005	0.000	0.000	1.000	0.000

Notes: This table provides an overview of the distribution of product characteristics across countries. The distributions are based on the count of each of the bottled water varieties that fall within a particular level of the categories. The category "Undisclosed" captures the share of varieties for which the product category is missing.

Table A.4: Product characteristics: R^2

$p_{j,lt}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)
PANEL (A): RAW DATA							
R^2	0.538	0.539	0.591	0.572	0.744	0.759	0.763
N	2,522,717	2,522,717	2,519,952	2,522,717	2,504,743	2,504,743	2,479,554
No. FEs	124,344	124,347	124,636	124,350	124,121	124,140	123,881
PANEL (B): CLEANED SAMPLE							
R^2	0.552	0.552	0.616	0.592	0.786	0.802	0.805
N	1,657,717	1,657,717	1,657,717	1,657,717	1,657,605	1,657,605	1,654,563
No. FEs	79,316	79,319	79,491	79,321	79,341	79,357	79,349
·Region × Time × Chain	✓	✓	✓	✓	✓	✓	✓
·Carbonated		✓	✓	✓	✓	✓	✓
·Flavor			✓				
·$\mathbb{1}$ (Flavored)				✓	✓	✓	✓
·Bottle size					✓	✓	✓
·Bottles per pack						✓	✓
·Package type							✓

Notes: This table provides the results from regressing tax-inclusive consumer prices on fixed effects capturing various levels of variation in the data. To conduct these estimations, I exclude data from Denmark and Sweden as the package type is not reported in these countries. If not, I would not be able to reliably compare the adjusted R^2 across specifications due to the sharp reduction in observations and, more importantly, the number of fixed effects for which the adjusted R^2 corrects. For each of the eight regressions, I show the resulting adjusted R^2, the number of observations included in the regression, which of the fixed effects are included and the total number of included fixed effects.

Table A.5: Hedonic price regression

	Raw sample				Cleaned sample			
$p_{j,lt}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}\,(\text{Sparkling})_j$	-0.0429	-0.0182	-0.0216	-0.00811	0.00629	0.00714	0.0269	0.0125
	(0.027)	(0.015)	(0.027)	(0.015)	(0.033)	(0.019)	(0.035)	(0.020)
$\mathbb{1}\,(\text{Flavored})_j$	0.528^{***}	0.484^{***}	0.465^{***}	0.39^{***}	0.541^{***}	0.508^{***}	0.493^{***}	0.401^{***}
	(0.033)	(0.024)	(0.036)	(0.025)	(0.048)	(0.033)	(0.049)	(0.033)
$\mathbb{1}\,(\text{Glass bottle})_j$	-0.103^{*}	-0.125^{***}	-0.0945	-0.148^{***}	-0.273^{*}	-0.138	-0.229	-0.158
	(0.059)	(0.037)	(0.058)	(0.037)	(0.149)	(0.114)	(0.152)	(0.110)
$\mathbb{1}\,(\text{Other package})_j$	0.411^{***}	0.138^{***}	0.386^{***}	0.071	0.323^{***}	0.0994	0.342^{***}	0.0265
	(0.053)	(0.049)	(0.066)	(0.053)	(0.105)	(0.071)	(0.119)	(0.091)
$\mathbb{1}\,((750\text{ml}, 1500\text{ml}))_j$	-0.249^{***}	-0.369^{***}	-0.239^{***}	-0.356^{***}	-0.231^{***}	-0.334^{***}	-0.243^{***}	-0.348^{***}
	(0.037)	(0.020)	(0.037)	(0.021)	(0.049)	(0.031)	(0.049)	(0.032)
$\mathbb{1}\,(\geq 1500\text{ml})_j$	-0.832^{***}	-0.778^{***}	-0.847^{***}	-0.794^{***}	-0.735^{***}	-0.735^{***}	-0.767^{***}	-0.769^{***}
	(0.033)	(0.018)	(0.035)	(0.018)	(0.043)	(0.025)	(0.045)	(0.027)
$\mathbb{1}\,(\text{Private label})_j$	-0.53^{***}		-0.507^{***}		-0.595^{***}		-0.559^{***}	
	(0.028)		(0.029)		(0.037)		(0.037)	
$\mathbb{1}\,(\text{Foreign})_j$			0.287^{***}	0.178^{***}			0.261^{***}	0.218^{***}
			(0.043)	(0.052)			(0.051)	(0.062)
Region-Time FEs	✓	✓	✓	✓	✓	✓	✓	✓
Brand FEs		✓		✓		✓		✓
$Adj.R^2$	0.62	0.84	0.62	0.85	0.69	0.87	0.70	0.88
No. obs	917,894	920,722	742,693	742,693	535,497	536,247	439,772	439,772

Notes: This table provides the results from regressing tax-inclusive consumer prices at the variety- and market-level on product characteristics. Because the package type is unreported in Denmark and Sweden, these countries are excluded from the estimation sample. To obtain variety-level prices, we aggregate transaction-level prices from the transaction level to the variety level using population weights to ensure that prices are representative. I cluster standard errors at the variety level. Reported significance levels are at the $p < 0.1^{*}$,$p < 0.05^{**}$ and $p < 0.01^{***}$ levels.

Table A.6: Sample selection: Products

Country	All	BEL	GER	DEN	FRA	NLD	PLN	SWE	UK
Transactions (count)									
·Included	12,394,789	806,698	3,875,517	86,186	2,583,913	958,184	1,419,445	126,282	2,538,564
·Excluded	2,733,529	94,686	1,255,503	7,746	445,151	76,408	546,133	8,482	299,420
Transactions (share)									
·Included	0.819	0.895	0.755	0.918	0.853	0.926	0.722	0.937	0.894
·Excluded	0.181	0.105	0.245	0.082	0.147	0.074	0.278	0.063	0.106
Liters (liters)									
·Included	91,244,705	7,259,419	38,305,894	283,658	25,273,701	3,420,255	8,108,307	301,970	8,291,501
·Excluded	19,379,029	945,221	13,635,274	23,053	2,088,889	273,193	1,737,543	11,388	664,468
Liters (share)									
·Included	0.825	0.885	0.737	0.925	0.924	0.926	0.824	0.964	0.926
·Excluded	0.175	0.115	0.263	0.075	0.076	0.074	0.176	0.036	0.074

Notes: This table provides an overview of the extent of the product-level sample restrictions. For each country and overall, I show the total number of transactions and the share in the number of transactions that are included in the final sample. I also show the total number of liters and the share of the number of liters accounted for products included in the final sample.

Table A.7: Sample selection: Transactions

Country	All	BEL	GER	DEN	FRA	NLD	PLN	SWE	UK
Transactions (count)									
·Included	15,001,811	894,044	5,126,332	92,868	2,930,570	1,030,844	1,964,823	130,323	2,832,007
·Winsorized	19,351	2,195	3,892	467	3,764	1,801	505	2,280	4,447
·Excluded	107,156	5,145	796	597	94,730	1,947	250	2,161	1,530
Transactions (share)									
·Included	0.992	0.992	0.999	0.989	0.967	0.996	1.000	0.967	0.998
·Winsorized	0.001	0.002	0.001	0.005	0.001	0.002	0.000	0.017	0.002
·Excluded	0.007	0.006	0.000	0.006	0.031	0.002	0.000	0.016	0.001
Liters (liters)									
·Included	110,451,086	8,132,361	51,924,396	296,661	27,323,238	3,687,670	9,839,625	311,073	8,936,063
·Winsorized	116,533	42,739	13,651	1,640	29,054	4,722	5,849	1,595	17,284
·Excluded	56,115	29,541	3,121	8,411	10,298	1,056	376	690	2,622
Liters (share)									
·Included	0.998	0.991	1.000	0.967	0.999	0.998	0.999	0.993	0.998
·Winsorized	0.001	0.005	0.000	0.005	0.001	0.001	0.001	0.005	0.002
·Excluded	0.001	0.004	0.000	0.027	0.000	0.000	0.000	0.002	0.000

Notes: This table provides an overview of the extent of the transaction-level sample restrictions. For each country and overall, I show the total number of transactions and the share in the number of transactions that are subject to winsorizing and are excluded because they seem implausibly small or large. I also show the total number of liters and the share of the number of liters accounted for by transactions that are subject to winsorizing and are excluded because they seem implausibly small or large.

A.3 Household-level consumption data

Table A.8: Household characteristics

Countries	BEL	GER	DEN	FRA	NLD	PLN	SWE	UK
Household income								
$\cdot \leq 1,900$EUR	0.157	0.332	0.284	0.304	0.159	0.546	0.329	0.346
$\cdot 1,900-2,700$EUR	0.548	0.331	0.216	0.257	0.489	0.306	0.322	0.359
$\cdot \geq 2,700$EUR	0.296	0.337	0.500	0.439	0.352	0.149	0.350	0.295
Household age								
$\cdot \leq 34y$	0.196	0.203	0.189	0.180	0.213	0.240	0.181	0.191
$\cdot 35-64y$	0.573	0.548	0.601	0.547	0.596	0.593	0.573	0.591
$\cdot \geq 65y$	0.232	0.250	0.210	0.273	0.191	0.167	0.246	0.219
Household size								
$\cdot 1-2$	0.614	0.704	0.668	0.651	0.630	0.455	0.703	0.570
$\cdot \geq 3$	0.386	0.296	0.332	0.349	0.370	0.545	0.297	0.430

Notes: This table shows the distribution of household characteristics for households that purchase water in a given quarter across countries. To compute the distribution, I compute for each of the household characteristics and for each quarter-regional market the share of total population weight accounted for by each level of the characteristics. To arrive at the final number, I take the average over regions and quarters by weighting each region-quarter pair equally.

Table A.9: Household heterogeneity

$y_{ij,lt}$	$\mathbb{1}\left(\sum_{j\in\mathcal{J}} q_{ij,lt} > 0\right)$ (1)	$\ln p_{ij,lt}$ (2)	$\mathbb{1}$ (Flavored)$_j$ (3)	$\mathbb{1}$ ((750, 1500)ml)$_j$ (4)	$\mathbb{1}$ ($\geq$ 1500ml)$_j$ (5)
$\mathbb{1}(1,900-2,700\text{EUR})_i$	0.029^{**}	-0.0015	-0.087^{***}	-0.0129	-0.0139
	(0.012)	(0.010)	(0.017)	(0.015)	(0.026)
$\mathbb{1}(\geq 2,700\text{EUR})_i$	0.0126	0.041^{***}	-0.138^{***}	-0.0205	-0.0271
	(0.012)	(0.011)	(0.021)	(0.016)	(0.030)
$\mathbb{1}(35-64y)_i$	0.124^{***}	0.00504	0.0202	0.00296	0.0561^{**}
	(0.011)	(0.009)	(0.017)	(0.014)	(0.027)
$\mathbb{1}(\geq 64y)_i$	-0.0324^{*}	0.0376^{**}	-0.0221	-0.00668	0.0923^{**}
	(0.017)	(0.015)	(0.026)	(0.019)	(0.039)
$\mathbb{1}(\geq 3\text{members})_i$	0.219^{***}	-0.0324^{***}	0.0924^{***}	0.0372^{***}	0.0587^{**}
	(0.009)	(0.010)	(0.016)	(0.011)	(0.023)
Region-Time FEs	✓	✓	✓	✓	✓
Flavored FE		✓		✓	✓
Bottle size FEs		✓	✓		
Estimator	Poisson	OLS	Poisson	Poisson	Poisson
$Adj.R^2$	-	0.84	-	-	-
No. obs	4,137,347	3,498,132	3,380,741	3,498,129	3,497,464

Notes: This table shows household heterogeneity in the consumption of bottled water. The first column estimates the probability of purchasing bottled water as a function of household characteristics using a Poisson MLE estimator. The second column estimates the conditional expectation of the log of the consumer price paid as a function of consumer characteristics using an OLS estimator. Columns (3) to (6) estimate the probability of purchasing sparkling water, flavored water, water in mid-to-large sized bottles and water packaged in glass, metal or cardboard packages respectively as a function of consumer characteristics using an OLS estimator. In doing so, I aggregate the transaction-level data to the household and variety level by aggregating over transactions using population weights. In the estimation, I weight observations by the corresponding population weights. I cluster standard errors at the household level and report significance levels are at the $p < 0.1^{*}$,$p < 0.05^{**}$ and $p < 0.01^{***}$ levels.

A.4 Production Locations

Figure A.1: Water sources - Accuracy

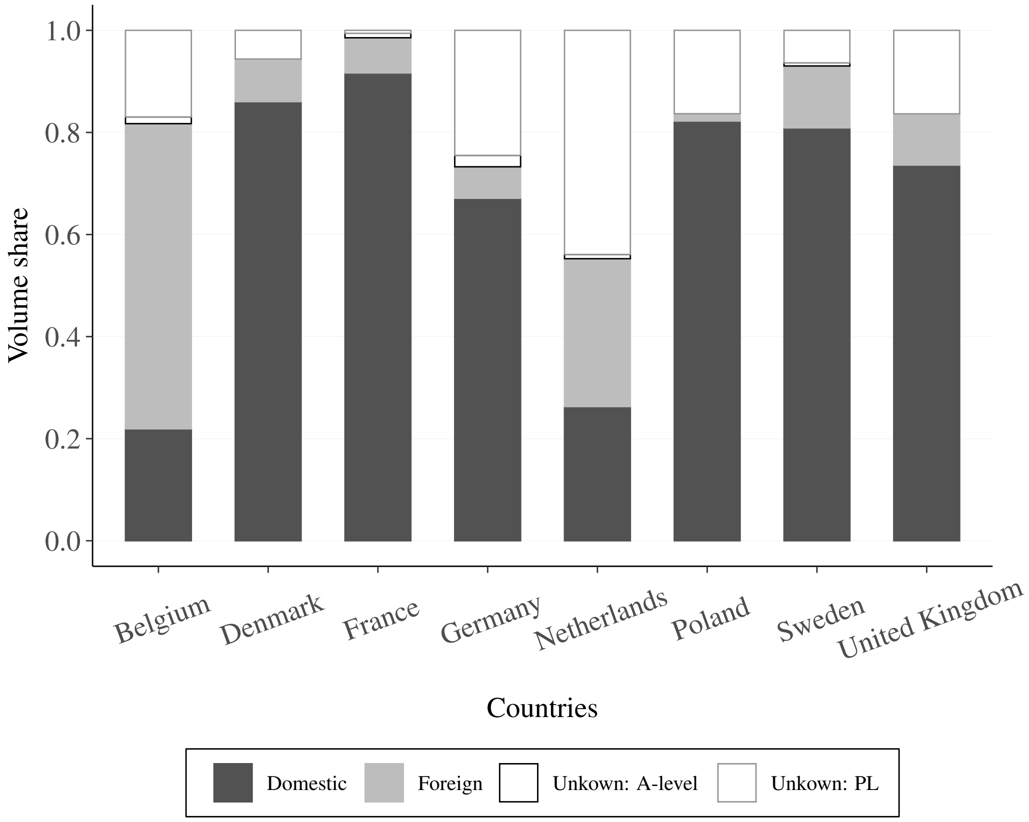

Notes: This figure plots for each country the distribution of volume sold across four groups: domestically sourced water in dark grey, foreign-sourced water in light grey, A-level brands with an unknown origin in white with a black contour and private label water with an unknown source in white with a grey contour. The weights are computed as the total volume in liters that fall in each category divided by the total volume sold. By doing so, we pool across all periods and all regions within a country.

A.5 Consumption taxes

Consumption taxes differ widely across European countries. First, across all countries in the sample, consumers are required to pay Value-Added Taxes (VAT) even though it is typically collected and transferred by the seller of the product. VAT uses the value of the purchase as the tax base. Even though VAT is somewhat homogenized across European countries in terms of the standard, reduced and super-reduced rates, countries only need to comply with certain boundaries of the tax within each rate. Moreover, depending on whether products are considered essential products, countries have discretionary power in terms of which of three VAT rates they apply to products. Second, many countries also collect excise duties on bottled water which is determined by the volume of the product. Often, the excise duties have a "health" aspect as they will differ between products with and without added sugars or artificial sweeteners. Finally, since 2008 Belgium and the Netherlands have had a disposable package tax that that depends on the weight and type of the package. Belgium, Denmark, Germany, the Netherlands and Sweden also have a deposit-return system in which consumers are reimbursed for each package returned to the retail store. France, Poland the UK will introduce similar systems in 2024, 2025 and 2025 respectively. We do not take this system into account as the reimbursement included when purchasing the product is almost quoted separately on the receipt and therefore not included in the unit price of the product. We provide a detailed description for each country below.

Belgium Belgium applies all three types of taxes. Throughout the sample period, Belgium has taxed bottled water under the super-reduced rate VAT which has been at 6% since 2000. Since 2011, Belgian consumers are also liable to pay an excise duty amounting to 0.0373 EUR/liter if the product is sweetened or flavored. The duty was raised to 0.0681 EUR/liter in 2016 and in 2018 it was further raised for 0.11923 EUR/liter for bottled water with added sugar. Finally, Belgium also introduced a so-called package tax in 2011. This tax depends on the volume of the product and on the product's package type. The rates are 0.0986 EUR/liter for water bottled in plastic bottles or metal cans and 0.0141 EUR/liter for glass bottles.

Denmark Denmark levies VAT and temporarily levied health-based excise duties. Denmark is one of the few countries that has a standard VAT rate and therefore bottled water is taxed at the standard rate which has been 25% since 1992. In 2008, Denmark introduced their version of a so-called sugar tax and started taxing artificially flavored bottled water and bottled water with added sugar at 0.57 DKK/liter. This excise tax was reduced to 0.30 DKK/liter in 2013 and finally abolished in 2014.

France Like Denmark, France levies both VAT and health-inspired excise taxes. France has, quite uniquely, four different VAT rates and taxes bottled water at the super-reduced rate. This rate has been 5.5% since 1982. Since 1942, France has had a excise dutie of 0.0058 EUR/liter on mineral waters. France also introduced their health-based excise duty scheme in 2012 in which artificially flavored bottled water and bottled water with added sugar were taxed at 0.076 EUR/liter. However, in 2018 this tax was reduced to 0.031 EUR/liter.

Germany Germany only taxes bottled water consumption through VAT taxes at the standard rate. Before July, 2020 the standard rate was 19%. This was temporarily reduced to 16% between July, 2020 and January, 2021 as part of one of the fiscal stimulus packages introduced by the German government to support aggregate demand in the wake of the COVID-19 pandemic. In January of 2021, the standard VAT rate returned to 19%.

The Netherlands The Netherlands also applies all three types of taxes to the consumption of bottled water. The structure of VAT in The Netherlands is also made up from three VAT rates. Like Belgium and France, bottled water is taxed at the super-reduced rate. Until 2019 this rate was 6%, currently it stands at 9%. In addition to VAT, The Netherlands also levies excise duties but does not differentiate those with respect to added sugar or artificial flavors. In 2008, this excise was 0.0413 EUR/liter. It

Table A.10: Package tax in The Netherlands

Date	Plastic	Glass	Metal	Cardboard
01/01/2008	0.4813	0.0734	0.9726	0.014
01/01/2013	0.3876	0.0595	0.0212	0.0233
01/01/2016	0.38	0.056	0.02	0.022
01/01/2020	0.34	0.056	0.05	0.022
01/01/2021	0.41	0.056	0.11	0.022
01/01/2022	0.44	0.048	0.16	0.022
01/01/2023	0.79	0.06	0.16	0.012

Notes: This table lays out the rate structure of the package tax in the Netherlands per package type over time. The rates are expressed in EUR/KG. To compute the tax liability for a given package in the data, we use the following conversion rates between package size and weight of the package. For plastic bottles, we consider 20gr, 24gr, 33gr, 35gr and 43gr for bottles sizes of 330ml, 500ml, 750ml, 1000ml and 1500ml respectively. We consider 10.6gr and 13.4gr for metal 330ml and 500ml metal cans. For glass bottles, we consider 200gr, 280gr, 400gr and 450gr for 330ml, 500ml, 750ml, 1000ml bottles respectively.

was subsequently raised to 0.057 EUR/liter and to 0.0883 EUR/liter in 2014 and 2016 respectively. Finally, The Netherlands levies a tax on disposable packages that is differentiated by package type.

Poland Poland taxes bottled water consumption using VAT and excise taxes. The VAT rate is the standard rate which was 22% until 2011. From 2011 onwards, the VAT rate has been 23%. In 2021, Poland also introduced health-based excise taxes. From that moment, bottled water with added sugar or artificial flavors is taxed at a rate of 0.5 PLN/liter.

Sweden Sweden only levies VAT on the consumption of bottled water. Sweden also has a three-part VAT structure and taxes bottled water consumption at the reduced rate of 12%, which has been in place since 1996.

United Kingdom Like Sweden the United Kingdom only levies VAT on the consumption of bottled water. The United Kingdom has a dual VAT structure with a standard rate and a reduced rate. Bottled water is taxed at the standard rate which has fluctuated over time. Until December of 2008, the standard rate was 17.5%. From that moment until January of 2010, the standard rate was lowered to 15% after which it was raised again to 17.5%. In April of 2011, the standard rate was increased to 20% at which it has remained ever since.

A.6 Labor unit costs

We construct a cross-country panel of labor unit costs in the beverage production, transportation and retail sectors by relying on data from the Survey of Income and Living Conditions compiled by Eurostat (EU-SILC). The EU-SILC database is constructed from surveys of individual households in all countries that are part of the European Union in which households provide information about the

Table A.11: EU-SILC variables

Variable	2010	2011-2019	BE	DK	FR	DE	NL	PL	SE	UK
Year	DB010	DB010	✓	✓	✓	✓	✓	✓	✓	✓
Country	DB020	DB020	✓	✓	✓	✓	✓	✓	✓	✓
NUTS1	DB040	DB040	✓	✓	✓			✓	✓	✓
Population weight	RB030	RB030	✓	✓	✓	✓	✓	✓	✓	✓
Birth year	RB080	RB080	✓	✓	✓	✓	✓	✓	✓	✓
Sex	RB090	RB090	✓	✓	✓	✓	✓	✓	✓	✓
Employment status	RB210	RB210	✓	✓	✓	✓	✓	✓	✓	✓
Education	PE040	PE040	✓	✓	✓	✓	✓	✓	✓	✓
Occupation	PL050	PL051	✓	✓	✓	✓	✓	✓	✓	✓
Hours worked	PL060	PL060	✓	✓	✓	✓	✓	✓	✓	✓
Full-time months	PL073	PL073	✓	✓	✓	✓	✓	✓	✓	✓
Part-time months	PL074	PL074	✓	✓	✓	✓	✓	✓	✓	✓
Unemployed months	PL080	PL080	✓	✓	✓	✓	✓	✓	✓	✓
Retired months	PL085	PL085	✓	✓	✓	✓	✓	✓	✓	✓
Out-for-health months	PL086	PL086	✓	✓	✓	✓	✓	✓	✓	✓
Study months	PL087	PL087	✓	✓	✓	✓	✓	✓	✓	✓
Public service months	PL088	PL088	✓	✓	✓	✓	✓	✓	✓	✓
Aid-at-home months	PL089	PL089	✓	✓	✓	✓	✓	✓	✓	✓
Other months	PL090	PL090	✓	✓	✓	✓	✓	✓	✓	✓
Sector	PL111	PL111	✓	✓	✓	✓	✓	✓	✓	✓
Pecuniary gross income	PY010G	PY010G	✓	✓	✓	✓	✓	✓	✓	✓
Non-pecuniary gross income (car)	PY020G	PY020G	✓	✓	✓	✓	✓	✓	✓	✓
Non-pecuniary gross income (other)	PY021G	PY021G	✓	✓	✓	✓	✓	✓	✓	✓
Employer Social sec. contribution	PY030G	PY030G	✓	✓	✓	✓	✓	✓	✓	✓

Notes: This table shows which variables in the EU-SILC surveys we use to construct the labor unit cost measure. Employment status refers to whether the person is employed, unemployed or not in the labor force. We convert the education levels into four levels: “Pre-primary”, “Primary”, “Secondary”, “Higher”. We consider the following occupations: “Military”, “Managers”, “Middle management”, “Professionals”, “Clerks”, “Service workers”, “High-skilled Laborers”, “Low-skilled Laborers” and “Other”. Based on these denominations, we group “Managers”, “Middle management”, “Professionals” and “Clerks” into white-collar workers and “Service workers”, “High-skilled Laborers” and “Low-skilled Laborers” into blue-collar workers. The sector of occupation is defined at the NACE 1 level. Hours worked refers to the average hours worked in a typical week. Full-time months, Part-time months, Unemployed months, Retired months, Out-of-health months, Study months, Public service months and Aid-at-home months refer to the number of months in that year that the person in each situation. The table also indicates which variables are available for which country.

personal characteristics and employment outcomes of all members of the household that are eligible to work.

Variables Besides information on personal characteristics, such as age, sex, educational attainment, etc., the dataset also registers the region of residence (at the NUTS1-level), the sector of occupation (at the NACE 1 level), the occupation and information about hours worked during the year. Finally, for each eligible worker, the dataset records information about pecuniary gross cash income, non-pecuniary gross non-cash income (such as health insurance or a company car) and social security contributions by employers. We proxy labor costs by the sum of these components. Table A.11 provides an overview of the variables that we use and their availability across countries.

Labor unit costs To compute labor unit costs, we take three steps. First, we estimate a hedonic wage model in which we predict hourly based on the following personal characteristics: sex, educational attainment, occupation, tenure, sector of employment, year and region of residence. We compute hourly labor costs by computing for each person who worked full-time during the past 12 months the effective hours worked in a given year by multiplying the average hours worked per week times 52. In turn, we estimate the following regression:

$$\ln W_{i,t} = \mathbb{1}\left(i = \text{Male}\right)_{n(i)} + \theta_{e(i)n(i)} + \theta_{o(i)n(i)} + \theta_{s(i)n(i)} \cdot \text{Tenure}_{i,t} + \lambda_{s(i),t} + \lambda_{r(i),t} + \varepsilon_{i,t}$$

where $W_{i,t}$ is the hourly wage, $\nVdash \left(i = \text{Male}\right)_{n(i)}$ is fixed effect for being male, $\theta_{e(i)n(i)}$ are educational attainment fixed effects, $\theta_{o(i)n(i)}$ are occupation fixed effects, $\theta_{s(i)n(i)}$ are sector fixed effects, $\lambda_{s(i),t}$ are sector-year fixed effects and $\lambda_{r(i),t}$ are region-time fixed effects. We interact all effects with country-fixed effects which allows for each of these personal characteristics to have heterogeneous effects across countries. Second, we use the estimated fixed effects to predict unit labor costs for each person in the data. Finally, we compute population-weighted median unit labor costs within country-year-industry-occupation cells.

B Reduced-form Evidence

B.1 Spatial price dispersion

Table B.1: Regression Discontinuity - Main stores

	1^{st}-order				2^{th} order			
$p_{j,lt}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Border$^{s(j)l}$	0.167^{***}	0.172^{***}	0.107^{***}	0.101^{***}	0.179^{***}	0.155^{***}	0.0885^{***}	0.0923^{***}
	$[0.13; 0.203]$	$[0.133; 0.21]$	$[0.08; 0.134]$	$[0.0751; 0.127]$	$[0.14; 0.218]$	$[0.119; 0.191]$	$[0.0655; 0.112]$	$[0.068; 0.117]$
$\theta_{j,t}$	✓	✓	✓	✓	✓	✓	✓	✓
Polynomial	1	1	1	1	2	2	2	2
Bandwidth	1,000	500	100	53.9	1,000	500	100	82.8
Optimal				✓				✓
No. obs	1,789,002	1,398,951	427,061	252,424	1,789,002	1,398,951	427,061	361,652

Notes: This table shows the results from estimating equation 3.1 using the robust Regression Discontinuity Design (RDD) estimator from Calonico et al. (2014) using the rdrobust package. Columns (1)-(4) show the results from estimating equation 3.1 using a first-order polynomial and Columns (5)-(8) for a second-order polynomial in the distance to the border. For each polynomial order, I show the results when I limit the sample in a range of 1,000, 500 and 100 kilometers to the border and when I consider the optimal bandwidth selection from Calonico et al. (2014). The sample underlying this table includes the most important stores, all products and is constructed by ranking, for each contiguous country pair, prices according to their distance to the border. The resulting sample pools across these contiguous countries. Prices are constructed as the weighted average across transactions where the weights are the population weights time the purchases liters associated with that transaction. Standard errors are clustered at the product level reported and are robust to the fact that bandwidths that are far away from zero can lead to bad coverage of the confidence intervals (see Calonico et al. (2014)). I report the robust confidence intervals in square brackets and denote significance at the $p < 0.1^{*}$,$p < 0.05^{**}$ and $p < 0.01^{***}$ levels.

Table B.2: Regression Discontinuity - Main products

	1^{st}-order				2^{th} order			
$p_{j,lt}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\mathbb{1}(\text{Border}^{s(j)l} = 1)$	0.157***	0.154***	0.074***	0.0581***	0.158***	0.133***	0.0426**	0.0482**
	$[0.124; 0.19]$	$[0.12; 0.189]$	$[0.0406; 0.107]$	$[0.0228; 0.0934]$	$[0.123; 0.193]$	$[0.0996; 0.166]$	$[0.00427; 0.0809]$	$[0.0111; 0.0852]$
$\theta_{j,t}$	✓	✓	✓	✓	✓	✓	✓	✓
Polynomial	1	1	1	1	2	2	2	2
Bandwidth	1,000	500	100	45.2	1,000	500	100	76.9
Optimal				✓				✓
No. obs	2,302,791	1,834,173	536,661	272,617	2,302,791	1,834,173	536,661	430,165

Notes: This table shows the results from estimating equation 3.1 using the robust Regression Discontinuity Design (RDD) estimator from Calonico et al. (2014) using the rdrobust package. Columns (1)-(4) show the results from estimating equation 3.1 using a first-order polynomial and Columns (5)-(8) for a second-order polynomial in the distance to the border. For each polynomial order, I show the results when I limit the sample in a range of 1,000, 500 and 100 kilometers to the border and when I consider the optimal bandwidth selection from Calonico et al. (2014). The sample underlying this table includes the most important products, all stores and is constructed by ranking, for each contiguous country pair, prices according to their distance to the border. The resulting sample pools across these contiguous countries. Prices are constructed as the weighted average across transactions where the weights are the population weights time the purchases liters associated with that transaction. Standard errors are clustered at the product level reported and are robust to the fact that bandwidths that are far away from zero can lead to bad coverage of the confidence intervals (see Calonico et al. (2014)). I report the robust confidence intervals in square brackets and denote significance at the $p < 0.1^*$,$p < 0.05^{**}$ and $p < 0.01^{***}$ levels.

Table B.3: Regression Discontinuity - Raw sample

	1^{st}-order				2^{th} order			
$p_{j,lt}$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$\text{Border}^{s(j)l}$	0.156***	0.153***	0.0727***	0.0568***	0.157***	0.132***	0.0416**	0.0472**
	$[0.123; 0.189]$	$[0.119; 0.188]$	$[0.0395; 0.106]$	$[0.0215; 0.0921]$	$[0.122; 0.192]$	$[0.0987; 0.165]$	$[0.00326; 0.0799]$	$[0.0101; 0.0842]$
$\theta_{j,t}$	✓	✓	✓	✓	✓	✓	✓	✓
Polynomial	1	1	1	1	2	2	2	2
Bandwidth	1,000	500	100	45.0	1,000	500	100	76.9
Optimal				✓				✓
No. obs	2,328,140	1,854,847	542,517	274,635	2,328,140	1,854,847	542,517	435,208

Notes: This table shows the results from estimating equation 3.1 using the robust Regression Discontinuity Design (RDD) estimator from Calonico et al. (2014) using the rdrobust package. Columns (1)-(4) show the results from estimating equation 3.1 using a first-order polynomial and Columns (5)-(8) for a second-order polynomial in the distance to the border. For each polynomial order, I show the results when I limit the sample in a range of 1,000, 500 and 100 kilometers to the border and when I consider the optimal bandwidth selection from Calonico et al. (2014). The sample underlying this table includes the all products and stores and is constructed by ranking, for each contiguous country pair, prices according to their distance to the border. The resulting sample pools across these contiguous countries. Prices are constructed as the weighted average across transactions where the weights are the population weights time the purchases liters associated with that transaction. Standard errors are clustered at the product level reported and are robust to the fact that bandwidths that are far away from zero can lead to bad coverage of the confidence intervals (see Calonico et al. (2014)). I report the robust confidence intervals in square brackets and denote significance at the $p < 0.1^*$,$p < 0.05^{**}$ and $p < 0.01^{***}$ levels.

Table B.4: Gravity estimation - Main products

	X_t^{kl}					$\frac{\boldsymbol{J}_l \cdot \boldsymbol{J}_k}{\lvert\lvert\boldsymbol{J}_l\rvert\rvert \cdot \lvert\lvert\boldsymbol{J}_k\rvert\rvert}$
	(1)	(2)	(3)	(4)	(5)	(6)
ln$(1+\text{Dis}^{lk})$	-0.724^{***}	-0.613^{***}	-0.611^{***}	-0.554^{***}	-0.472^{***}	-0.108^{***}
	(0.098)	(0.098)	(0.098)	(0.094)	(0.089)	(0.015)
Border^{lk}		-1.36^{***}	-1.35^{***}	-1.19^{***}	-0.976^{***}	-1.35^{***}
		(0.267)	(0.267)	(0.251)	(0.260)	(0.057)
Cur^{lk}			-2.83^{**}	-2.21^{*}	-1.72^{*}	-0.694^{***}
			(1.171)	(1.133)	(1.037)	(0.071)
ln(Price_t^{lk})				-0.507^{**}	-0.797^{***}	
				(0.233)	(0.209)	
$e^{\hat{\beta}}-1$	-	-74.4%	-74.1%	-69.7%	-62.3%	-74.2%
$e^{\frac{\hat{\beta}}{\varepsilon_{EK}}}-1$	-	-15.2%	-15.0%	-13.4%	-	-
$e^{\frac{\hat{\beta}}{\varepsilon_{BLP}}}-1$	-	-47.4%	-47.1%	-	-36.9%	-
$e^{\hat{\gamma}}-1$	-	-	-94.1%	-89.1%	-82.2%	-50.0%
$e^{\frac{\hat{\gamma}}{\varepsilon_{EK}}}-1$	-	-	-29.0%	-23.4	-	-
$e^{\frac{\hat{\gamma}}{\varepsilon_{BLP}}}-1$	-	-	-73.7%	-	-55.6%	-
λ_l						✓
λ_k						✓
$\lambda_{l,t}$	✓	✓	✓	✓	✓	
$\lambda_{k,t}$	✓	✓	✓	✓	✓	
No. obs	76,660	76,660	76,660	76,660	76,660	23,409

Notes: This table presents the results from estimating equation 2 using Pseudo Poisson Maximum Likelihood (PPML) to account for zeros in the dependent variable. Columns (1) to (5) present various specifications with bilateral trade flows, X_t^{kl}, as the dependent variable. I compute bilateral trade flows by computing for each destination market the tax-inclusive expenditures denominated in EUR flowing in from all origin markets included in the sample via aggregation across transactions using population weights. Bilateral prices P_t^{lk} are computed by aggregating variety-level prices using a CES-aggregator: $P_t^{lk} \equiv \left(\sum_{j \in \mathcal{J}_t^{lk}} \left(P_{j,t}^{kl}\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$ across all varieties that flow from l to k indicated by the set $\mathcal{J}_t^{lk}$. In doing so, I include observations from all stores and the main products and consider the following elasticities$\varepsilon_{EK} = 8.28$ taken from Eaton & Kortum (2002) and $\varepsilon_{BLP} = 2.12$ taken from Boehm et al. (2023). Column (6) considers the similarity in the set of available varieties between any two given markets as the dependent variable. The similarity is computed as the similarity between two vectors that indicate whether a given variety is present in a given market, i.e. $(\mathcal{J}_l)_j = 1$ if and only if j is sold in regional market l at time t. Columns (6) further collapses the time dimension by computing similarity measures at the region pair level. I present two-way clustered standard errors at the origin and destination market level and denote the significance levels at the 0.1^{*}, 0.05^{**}, 0.01^{***} levels.

Table B.5: Gravity estimation - Main stores

	X_t^{kl}					$\frac{\boldsymbol{J}_l\cdot\boldsymbol{J}_k}{\lVert\boldsymbol{J}_l\rVert\cdot\lVert\boldsymbol{J}_k\rVert}$
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(1+\text{Dis}^{lk})$	-0.815^{***}	-0.743^{***}	-0.742^{***}	-0.643^{***}	-0.464^{***}	-0.108^{***}
	(0.106)	(0.110)	(0.111)	(0.103)	(0.080)	(0.015)
Border^{lk}		-1.19^{***}	-1.19^{***}	-0.952^{***}	-0.678^{**}	-1.35^{***}
		(0.251)	(0.251)	(0.243)	(0.292)	(0.057)
Cur^{lk}			-0.507	-0.202	-0.123	-0.694^{***}
			(0.797)	(0.694)	(0.624)	(0.071)
$\ln(\text{Price}_t^{lk})$				-0.684^{***}	-1.05^{***}	
				(0.218)	(0.184)	
$e^{\hat{\beta}}-1$	-	-69.7%	-69.6%	-61.4%	-49.2%	-74.2%
$e^{\frac{\hat{\beta}}{\varepsilon_{EK}}}-1$	-	-13.4%	-13.4%	-10.9%	-	-
$e^{\frac{\hat{\beta}}{\varepsilon_{BLP}}}-1$	-	-43.1%	-42.9%	-	-27.4%	-
$e^{\hat{\gamma}}-1$	-	-	-39.7%	-18.3%	-11.6%	-50.0%
$e^{\frac{\hat{\gamma}}{\varepsilon_{EK}}}-1$	-	-	-5.9%	-2.4	-	-
$e^{\frac{\hat{\gamma}}{\varepsilon_{BLP}}}-1$	-	-	-21.3%	-	-5.6%	-
λ_l						✓
λ_k						✓
$\lambda_{l,t}$	✓	✓	✓	✓	✓	
$\lambda_{k,t}$	✓	✓	✓	✓	✓	
No. obs	80,102	80,102	80,102	80,102	80,102	23,409

Notes: This table presents the results from estimating equation 2 using Pseudo Poisson Maximum Likelihood (PPML) to account for zeros in the dependent variable. Columns (1) to (5) present various specifications with bilateral trade flows, X_t^{kl}, as the dependent variable. I compute bilateral trade flows by computing for each destination market the tax-inclusive expenditures denominated in EUR flowing in from all origin markets included in the sample via aggregation across transactions using population weights. Bilateral prices P_t^{lk} are computed by aggregating variety-level prices using a CES-aggregator: $P_t^{lk} \equiv \left(\sum_{j\in\mathcal{J}_t^{lk}}\left(P_{j,t}^{kl}\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$ across all varieties that flow from l to k indicated by the set $\mathcal{J}_t^{lk}$. In doing so, I include observations from the main stores and all products and consider the following elasticities$\varepsilon_{EK} = 8.28$ taken from Eaton & Kortum (2002) and $\varepsilon_{BLP} = 2.12$ taken from Boehm et al. (2023). Column (6) considers the similarity in the set of available varieties between any two given markets as the dependent variable. The similarity is computed as the similarity between two vectors that indicate whether a given variety is present in a given market, i.e. $(\mathcal{J}_l)_j = 1$ if and only if j is sold in regional market l at time t. Columns (6) further collapses the time dimension by computing similarity measures at the region pair level. I present two-way clustered standard errors at the origin and destination market level and denote the significance levels at the 0.1^{*}, 0.05^{**}, 0.01^{***} levels.

Table B.6: Gravity estimation - Raw sample

	X_t^{kl}					$\frac{\boldsymbol{J}_l \cdot \boldsymbol{J}_k}{\lVert\boldsymbol{J}_l\rVert \cdot \lVert\boldsymbol{J}_k\rVert}$
	(1)	(2)	(3)	(4)	(5)	(6)
ln(1+Dis^{lk})	-1^{***}	-0.913^{***}	-0.912^{***}	-0.804^{***}	-0.567^{***}	-0.108^{***}
	(0.123)	(0.132)	(0.132)	(0.124)	(0.096)	(0.015)
Border^{lk}		-1.41^{***}	-1.4^{***}	-1.16^{***}	-0.848^{***}	-1.35^{***}
		(0.311)	(0.311)	(0.266)	(0.280)	(0.057)
Cur^{lk}			-0.803	-0.386	-0.305	-0.694^{***}
			(1.043)	(0.818)	(0.631)	(0.071)
ln(Price$_t^{lk}$)				-0.684^{***}	-1.16^{***}	
				(0.210)	(0.195)	
$e^{\hat{\beta}} - 1$	-	-75.6%	-75.3%	-68.7%	-57.2%	-74.2%
$e^{\frac{\hat{\beta}}{\varepsilon_{EK}}} - 1$	-	-15.7%	-15.5%	-13.1%	-	-
$e^{\frac{\hat{\beta}}{\varepsilon_{BLP}}} - 1$	-	-48.6%	-48.3%	-	-33.0%	-
$e^{\hat{\gamma}} - 1$	-	-	-55.2%	-32.0%	-26.3%	-50.0%
$e^{\frac{\hat{\gamma}}{\varepsilon_{EK}}} - 1$	-	-	-9.2%	-4.6	-	-
$e^{\frac{\hat{\gamma}}{\varepsilon_{BLP}}} - 1$	-	-	-31.5%	-	-13.4%	-
λ_l						✓
λ_k						✓
$\lambda_{l,t}$	✓	✓	✓	✓	✓	
$\lambda_{k,t}$	✓	✓	✓	✓	✓	
No. obs	85,774	85,774	85,774	85,774	85,774	23,409

Notes: This table presents the results from estimating equation 2 using Pseudo Poisson Maximum Likelihood (PPML) to account for zeros in the dependent variable. Columns (1) to (5) present various specifications with bilateral trade flows, X_t^{kl}, as the dependent variable. I compute bilateral trade flows by computing for each destination market the tax-inclusive expenditures denominated in EUR flowing in from all origin markets included in the sample via aggregation across transactions using population weights. Bilateral prices P_t^{lk} are computed by aggregating variety-level prices using a CES-aggregator: $P_t^{lk} \equiv \left(\sum_{j \in \mathcal{J}_t^{lk}} \left(P_{j,t}^{kl}\right)^{1-\sigma}\right)^{\frac{1}{1-\sigma}}$ across all varieties that flow from l to k indicated by the set $\mathcal{J}_t^{lk}$. In doing so, I include observations from all stores and products and consider the following elasticities$\varepsilon_{EK} = 8.28$ taken from Eaton & Kortum (2002) and $\varepsilon_{BLP} = 2.12$ taken from Boehm et al. (2023). Column (6) considers the similarity in the set of available varieties between any two given markets as the dependent variable. The similarity is computed as the similarity between two vectors that indicate whether a given variety is present in a given market, i.e. $(\mathcal{J}_l)_j = 1$ if and only if j is sold in regional market l at time t. Columns (6) further collapses the time dimension by computing similarity measures at the region pair level. I present two-way clustered standard errors at the origin and destination market level and denote the significance levels at the 0.1^{*}, 0.05^{**}, 0.01^{***} levels.