Price Dispersion and Market Segmentation: Evidence from Bottled Water*

Joris Hoste†
University of Cambridge

February 11, 2025

Abstract

It is well known that substantial spatial price dispersion persists between countries in the European Single Market (ESM). However, whether this dispersion reflects destination-specific markups, differences in marginal costs, or trade barriers remains an open question. This paper leverages technological and institutional features of the European bottled water industry to estimate the magnitude of trade barriers between EU countries and their impact on price dispersion and consumer welfare. Using a partial equilibrium model of the bottled water industry, I estimate that barriers to trade between ESM countries are equivalent to a 20% import tariff. Relative to an integrated benchmark economy, these trade barriers increase equilibrium international spatial price dispersion by 5% and reduce consumer welfare by an amount 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`

1 Introduction

The same product often sells for different prices in different geographical markets (e.g. Handbury, 2021). Importantly, price differences that occur discontinuously at geographical boundaries are frequently attributed to the presence of barriers to trade, i.e. forces that segment markets by making it more costly to sell in a foreign market (e.g. Gopinath et al., 2011). In this light, large price differences between countries that do not share a unified market, such as the US and Canada, might be unsurprising (e.g. Engel & Rogers, 1996; Gorodnichenko & Tesar, 2009) but similar evidence among countries of the European Single Market (ESM) has cast doubt on the extent of integration in the ESM (e.g. Beck et al., 2020). As greater goods market integration has been cited to reduce spatial misallocation (e.g. Hornbeck & Rotemberg, 2024) and increase innovation (e.g. Andersson et al., 2023), estimating the size of trade barriers between EU countries is crucial to evaluate the potential gains for further market integration in the ESM.

However, using spatial price dispersion as a measure of trade barriers faces two major challenges. First, when markets are segmented, firms may exploit local monopoly power and price discriminate by charging destination-specific markups. This means that observed prices reflect not only destination-specific marginal costs, which are influenced by barriers to trade, but also differences in markups. To date, much of the literature has relied on restrictive assumptions about demand and market structure to link price variation 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, even if we could observe marginal costs, differences might arise from other factors unrelated to trade barriers. For instance, in a world dominated by multinationals, products sold in neighboring markets might traverse distinct value chains that are subject to different input prices and are operating at different scales of production. Without detailed information on the spatial structure of value chains, disentangling barriers to trade from other marginal cost differences typically requires making restrictive assumptions about the production technology firms use and the market structure in which they compete.

In this paper, I estimate the level of trade frictions between EU countries and quantify their effect on equilibrium spatial price dispersion and consumer welfare. To overcome the abovementioned identification challenges, I focus on the European bottled water industry for two reasons. First, bottled water is an ideal product for estimating destination-specific marginal costs with minimal assumptions on preferences, market structure, and technology. Since bottled water has a low value-to-weight ratio, cross-border arbitrage is unlikely, and consumers tend to buy it near their residence. Relying on household-level scanner data from multiple countries, I use this feature to construct a dataset on consumption and prices of bottled water at a spatially disaggregate level. Combined with a flexible model of demand and supply in the bottled water industry, these data enable me to estimate destination-specific marginal costs without relying on restrictive assumptions about preferences, market structure, or production technology.

Second, I exploit regulatory and technological features of the bottled water industry to isolate trade barriers from other sources of marginal cost differences. EU law mandates that bottled water producers specify the name of the water source on product labels, enabling me to hand-collect data on production locations of varieties sold across EU countries. The industry's relatively simple value chain – where water is bottled at the source, transported by truck, and distributed locally – further aids in disentangling trade barriers. After controlling for local distribution costs in retail and transport costs, variation in marginal costs across destinations for water from the same source provides a credible estimate of barriers to trade.

Three pieces of reduced-form evidence suggest that barriers to trade in the European bottled water industry could be substantial. First, spatial price dispersion is much larger between than within countries. Whereas absolute differences in after-tax consumer prices between regions of different countries are on average 33%, they average 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. Because systematically higher markups and distribution costs abroad seem unlikely, this is suggestive evidence that a non-trivial part of the cross-border price dispersion could be driven by barriers to trade. Finally, trade flows of bottled water have a gravity structure: regional trade flows fall with distance and drop discontinuously at country borders. Also, 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. While these findings provide suggestive evidence of substantial trade barriers in the European bottled water industry, they do not account for differences in preferences, market structure and other destination-specific costs.

To do so, 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. Destination-specific marginal costs are inferred by combining the data on destination-specific prices with first-order conditions for profit-maximizing prices. As manufacturers reach consumers via local distributors and there is uncertainty on how they precisely interact, I consider various models of vertical price-setting where either upstream manufacturers, downstream retailers or both set prices in a simultaneous move Bertrand game. Combined with information on production locations, data on local labor costs and transportation costs, the model enables me to disentangle barriers to trade from other sources of destination-specific marginal costs.

Estimating the model reveals two key insights. First, cross-border trade barriers within the European Single Market (ESM) are substantial. After controlling for local distribution and transport costs, the marginal cost of selling bottled water increases by an average of 9 cents per liter in foreign markets, making foreign prices 20% higher than domestic ones. Furthermore, trade barriers are lower between Eurozone countries (7 cents per liter) compared to countries where at least one does not use the Euro (23 cents per liter). Second, spatial price discrimination reduces price dispersion: manufacturing markups are 0.5 cents lower for bottled water varieties sold in foreign markets. This result aligns with the pricing-to-market literature (e.g., Atkeson & Burstein (2008); Fitzgerald &

Haller (2014); Corsetti et al. (2021)) while also showing that standard trade models, such as those assuming CES demand and monopolistic competition, would have underestimated this effect by only 0.5 cents.

To evaluate the effect of cross-border trade frictions on consumer welfare in the bottled water industry, I conduct a counterfactual exercise where trade frictions are removed while keeping the set of available products fixed. Comparing an integrated economy, where bottled water is supplied solely at destination-specific marginal costs, to one where exporting is subject to barriers to trade, I find that trade barriers increase cross-country price dispersion by 5%. Additionally, these frictions reduce consumer welfare by approximately 3.6 cents per liter, which is equivalent to a 10% consumption tax on bottled water. These findings highlight the economic cost of cross-border trade barriers within the ESM, even in an industry with a relatively simple value chain.

This paper contributes to four strands of literature, with the first focusing on deviations from the Law of One Price (LOP), beginning with Engel & Rogers (1996) and P. Goldberg & Knetter (1997). Recent studies, including 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 explanations such as spatial price discrimination (e.g., P. Goldberg & Verboven, 2001; Gopinath et al., 2011) and differences in non-traded input costs (e.g., Crucini et al., 2005; Parsley & Wei, 2007) have been proposed, separating trade frictions from these factors has remained challenging. This paper addresses this gap by leveraging the unique regulatory and technological features of the bottled water industry within a flexible structural model. Similar to Asplund & Friberg (2001), who study international price differences in duty-free shops, I use an industry with particular institutional features that help to isolate trade frictions from other sources of destination-specific marginal costs. Estimating the structural model separates destination-specific marginal costs from markups and provides novel evidence that barriers to trade are a significant contributor to LOP deviations between EU countries.

Second, this paper is related to an emerging literature that leverages within-country price dispersion to evaluate policies aimed at improving market integration and understanding the incidence of international shocks. Shiue & Keller (2007) and Donaldson (2018) focus on commodity markets, using within-country price dispersion to assess market integration in China and India, respectively. Atkin & Donaldson (2015) and Chatterjee (2023) examine how transport infrastructure and regulatory barriers, respectively, interact with intermediaries along the value chain to influence spatial price dispersion. Building on this literature, this paper uniquely combines within- and between-country price dispersion to estimate the magnitude of trade barriers between EU countries and quantifies the gap between current EU market integration and that of 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, Coşar et al. (2015), Head & Mayer (2021) and Santamaría et al. (2023) rely on the gravity framework to estimate the level of 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, the paper recovers 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, P. K. 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, P. 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 details the construction of the dataset. I begin with an overview of the cross-country scanner data on bottled water purchases. Next, I discuss key technological and institutional aspects of the industry, which inform the identification of production locations. Finally, I describe the data sources used to compile information on transportation costs, indirect taxes, and labor unit costs.

2.1 Consumption data

To construct the consumption dataset, I use household-level scanner data from eight countries. In each country, a market research firm equips households with scanning devices to record their purchases of grocery products. Between 2010 and 2019, the dataset captures 15 million bottled water transactions. Each transaction records the purchased barcode, the retail chain, a household identifier, the quantity and volume purchased, and the tax-inclusive value in local currency. By combining the reported data on units sold, volume sold, and expenditure with package information extracted from barcode descriptions, I standardize quantities as liters and consumer prices as price per liter in euros.[1]

[1] Daily FX rates from Eurostat are used to convert local currency units (LCUs) to euros.

Countries The sample includes Belgium, Germany, Denmark, France, The Netherlands, Poland, Sweden and the United Kingdom.[2] As founding members of the ESM and the Eurozone, Belgium, Germany, France, and The Netherlands arguably represent a set of relatively well-integrated countries. This suggests that the estimated barriers to trade in this dataset may be conservative relative to the broader ESM. Additionally, the sample includes four European countries that are not part of the Eurozone, enabling me to explore heterogeneity in the estimated barriers to trade.

Retail chains The dataset encompasses all types of retail chains, but for the main analysis, I restrict the store dimension in two ways. First, I focus exclusively on domestic purchases, excluding cross-border transactions. I designate a transaction as a cross-border transaction when it satisfies at least one of the following criteria. On the one hand, while the dataset does not specify the exact outlet of purchase, it indicates whether the outlet is domestic or foreign. If the outlet is foreign, the transaction is classified as cross-border. On the other hand, I also classify certain transactions with relatively low consumer prices near country borders as cross-border. Specifically, for contiguous country pairs, transactions are designated as cross-border if they involve products available in both countries, occur within 40 kilometers of the border, and have a consumer price below 75% of the country-specific median price for that product.[3] Table A.1 shows that cross-border shopping accounts for less than 0.6% of transactions, representing just over 1% of the purchased volume. However, it is more prevalent in Belgium and the Netherlands, at 4.6% and 2.8% of transactions, respectively.

Second, I exclude transactions from smaller retail chains, focusing on chains with a market share exceeding 1% in the bottled water industry. Purchases across different store formats within these chains are aggregated at the chain level. After excluding cross-border transactions and small retail chains, the sample retains around 92% of transactions and over 90% of the purchased volume, as documented in Table A.2.

Water varieties and firms Individual bottled water varieties are defined by a combination of brand, whether they are flavored or unflavored, and bottle size. First, each variety is linked to a brand and its corresponding producer. To establish this ownership matrix, I combine data from GS1 with hand-collected data. For A-label brands, the brand owner is also the producer. However, Table A.3 highlights the prevalence of private label brands—brands owned by retail chains—across countries. Consistent with Bonnet & Dubois (2010); Molina (2021), and the observation that private label water is generally cheaper (see Table A.5), I assume retail chains are vertically integrated with the producers of their private label varieties.

Second, while bottled water is predominantly unflavored in Belgium, Germany, and France, flavored varieties are much more common in Denmark, Sweden, and Great Britain (Table A.3).

[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 calculated based on transactions of the same product within the same retail chain and quarter.

Nevertheless, as most price variation across flavored and unflavored varieties is captured by a simple flavor indicator variable (Table A.4), I do not differentiate between specific flavors. Conditional on other product characteristics, flavored water is, on average, 34% more expensive than unflavored water (Table A.5).

Third, bottle size is a key dimension of horizontal product differentiation. The raw data includes over 20 distinct bottle sizes, and Table A.4 demonstrates that categorizing these as small ($\leq 750ml$), medium ($750ml{-}1500ml$), and large ($\geq 1500ml$) fails to adequately explain price variation. To avoid conflating price differences with unobserved product characteristics (D. Hummels & Skiba, 2004; Manova & Zhang, 2012; Bastos et al., 2018), I define varieties based on exact bottle sizes. Conditional on other characteristics, medium and large bottles are, respectively, 40% and 80% cheaper than small bottles (Table A.5).

Finally, I collapse two additional product characteristics —- still versus sparkling water and the material from which the bottles are made – since Tables A.3 and A.4 show that these factors do not significantly explain price variation when other characteristics are accounted for.

I impose two restrictions at the product level. First, I exclude observations with consumer prices that deviate by more than five times the median price for the variety in a given quarter.[4] Additionally, I winsorize prices that deviate by three to five times the median. As shown in Table A.7, these adjustments affect 1.3% of transactions and less than 1% of the volume sold.

Second, to ensure that the sample includes all the important local varieties and varieties that are traded internationally, I retain only varieties with a market share of at least 0.1% by volume or those sold in at least two countries and 40 markets. [5] Table A.6 indicates that these restrictions preserve over 80% of the volume sold in most countries and 69% of the aggregate volume. In Germany, however, the sample covers only 54% of the volume due to the fragmented market structure of the bottled water industry. Nevertheless, the reduced-form evidence in Section 3 is robust to loosening these sample restrictions.

Table 1 summarizes the final sample, which includes just under 10 million transactions covering 346 products owned by 70 firms and sold in 106 different chains. Crucially, excluding specific products or retail chains does not impact the ability to capture differences in the market structure across countries as I still use the excluded transactions when computing the outside option.

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 for capturing both within- and between-country variation in prices and consumption. As noted earlier, cross-border shopping is relatively unimportant in the bottled water industry, so I can reasonably assume that households purchase and consume bottled water near their residence. This assumption allows me to disaggregate consumer prices and purchased quantities at the regional level.

[4] The median price is computed across retail chains and regions.
[5] A market is defined as a region-quarter combination

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.

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 also categorize households into three age groups: $[\leq 34 \text{ years}, 35 - 64 \text{ years}, \geq 65 \text{ years}]$, and I categorize income into three categories: $[\leq 1,900 \text{ EUR}, 1,900 - 2,700 \text{ EUR}, \geq 2,700 \text{ EUR}]$. 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, making it essential to observe household characteristics to capture 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 these population weights to maximize the external validity of the results.

Regions To estimate barriers to trade, my strategy compares between- and within-country variation in the backed-out destination-specific marginal costs that need to incurred to supply different markets. In doing so, I balance the need for a fine spatial resolution with a desire for sufficiently high sample coverage in each regional market. For this reason, I define regional markets within each country using the NUTS 2 (rev. 2013) classification.[6] This approach results in 154 regional markets, with an average of around 600 sampled households per market per quarter (Table 1).

Potential market and outside option A key reason why prices and consumption of bottled water vary between countries is differences in the appeal of not purchasing bottled water. In regions where the (perceived) quality of 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-quarter combination, I calculate the sum of the population weights for consumers who record at least one purchase in the NARTD (Non-Alcoholic Ready-To-Drink) category, which includes bottled water, sodas, juices, and energy drinks. Second, I compute the sum of the population weights for consumers who also purchased water. The inside good share is then the ratio of the population weights of consumers who purchase water to the total population weights of consumers purchasing beverages. Finally, the potential market is determined by dividing total water consumption (in liters) by the inside good share. Table 1 shows substantial cross-country variation in the inside good share, which intuitively correlates with average consumer prices of bottled water and the number of local sources. For instance, the inside good share is around 80% in countries with many domestic water sources, such as France, Germany, and Poland, while it is below 40% in more expensive countries like Denmark and Sweden.

2.2 Production locations

Estimating trade frictions requires precise knowledge of the location of different steps in the value chain. This is particularly important in differentiated product markets, where multinational firms account for the bulk of trade flows (Antràs et al., 2024). These firms serve markets through direct exporting (e.g. Melitz, 2003), local affiliates (e.g. Helpman et al., 2004), or more sophisticated strategies involving export platforms (e.g. Tintelnot, 2016). The complexity of these organizational structures means that products sold in neighboring markets often traverse distinct value chains, which may operate at different production scales and utilize different input suppliers. Without detailed information on the spatial structure of these value chains, disentangling trade frictions from other sources of variation in destination-specific marginal costs would require making restrictive assumptions about production technologies and market structures.

To cut through this complexity, I leverage regulatory and technological features of the bottled

[6] The NUTS classification is a European standard for referencing administrative levels within countries. Following administrative reforms in 2016, France updated their NUTS classifications. Since the regional variable in the French dataset corresponds to the NUTS 2 level of the NUTS-2 (rev. 2013) version, I use the 2013 classification throughout the paper.

water industry. First, the European Commission mandates that manufacturers disclose the source of the water on the product's label. This ensures that the mineral content of the water matches what is advertised and reported to food safety authorities. Second, to comply with these regulations is a cost-efficient way, manufacturers typically bottle the water near its source and ship the sealed bottles to different markets. Using these two features, I take advantage of the registry under Directive 2009/54/EC, which records information on water brands, their sources, and the locations of those sources.[7] I augment this data with hand-collected information from company websites and product labels for varieties not included in the registry.[8]

For each country, Figure A.1 shows the distribution of total volume sold by sourcing mode -— domestic, foreign, or unknown. When the source is unknown, I further distinguish between A-level brands and private labels. In most countries, I classify about 80% of the total volume as either domestically sourced or sourced from abroad.[9] For a portion of the volume, I am unable to determine the production location. Figure A.1 indicates that the share of unknown sources is concentrated among private label varieties. Unlike A-level brands, which use their sources as a point of differentiation, private-label varieties tend to be more homogeneous and low-cost.[10] This gives retail chains more flexibility in choosing from where to source their water, making it sometimes impossible to pinpoint the production location.

Figure 1 provides an overview of water sources for the countries in the dataset. Figure 1a shows the water sources for bottled water consumed in Belgium (represented by dark dots), with lighter dots indicating other sources not used for Belgian consumption. Figure 1b displays the same information for German bottled water consumption. German consumers benefit from a larger number of local water sources compared to Belgian consumers, and they consume a higher proportion of domestically sourced bottled water. This is also consistent with lower average consumer prices in Germany (see Table 1).

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 vary significantly between European countries. Whereas Belgium and France impose a 6% VAT rate, Germany and Denmark charge 21% and 25%, respectively. In addition, excise duties often target beverages with added sugger or that are artificially

[7] Directive 2009/54/EC regulates the branding and information disclosure for bottled water sold in European countries.

[8] These are typically private label brands, for which I rely on the source disclosed on individual product labels to determine the source and municipality.

[9] The Netherlands is an outlier, with slightly less than 60% of the volume classified as either domestic or foreign sourcing.

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

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.

sweetened, though the level and timing of these taxes differ across countries. To account for these differences, I collect data on indirect taxes—including VAT, excise duties, and taxes on disposable packaging—from the European Commission's "Taxes in Europe Database" and from country-level regulatory authorities. Appendix A.5 provides an overview of the specific taxes applied in each country in the sample.

Labor unit costs A key source of differences in destination-specific marginal costs is differences in non-traded input prices (Crucini et al., 2005; Parsley & Wei, 2007; Burstein et al., 2005, e.g.). When manufacturers sell to final consumers through retailers, locally sourced labor is typically employed to handle inventories, restock shelves and provide customer service. Given that European labor markets tend to be segmented along country borders, differences in retail wages need to be accounted 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. This dataset is based 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, different components of their (non)-pecuniary income and the social security contributions paid by their employers. Appendix A.6 provides more detail on the constructions of these series.

Transportation costs Differences in transport costs are a second source of differences in destination-specific marginal 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.[11] 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 labor unit costs, and interact this with travel times to account for additional time-related variations in transport costs.

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 deviations from the Law of One Price (LOP) are significantly larger across countries than within them. Second, I use a border Regression Discontinuity Design (RDD) to demonstrate that a substantial portion of cross-country price dispersion arises discontinuously at national borders. Finally, I show that a structural gravity model explains the observed variation in prices and trade flows through the presence of considerable tariff-equivalent trade barriers ranging from 14.3% to 43%, depending on the assumed elasticity of substitution

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 shows that, after controlling for 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 aggregate variety-level prices within each quarter and region across the households and stores by weighting transactions using liters sold interacted with population weights. I then compute, for each region pair and quarter, the absolute price differences for all varieties sold 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. Although spatial price dispersion within countries averages 8.2%, absolute LOP deviations are close to zero in nearly 50% of cases. However, cross-country price dispersion is much larger, averaging 35.5%, or approximately 27 percentage points (ppt) higher than within-country dispersion.

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

One potential reason why spatial price dispersion is larger for international region pairs is that indirect taxes differ between countries. To assess this, figure 2b replicates figure 2a using after-tax consumer prices instead. While price differences between countries drop slightly, they remain about 25 percentage points higher than within-country dispersion, suggesting that differences in indirect taxes do not fully explain international price dispersion.

Figure 2: Spatial price dispersion

(a) Domestic vs. International - Tax-inclusive **(b)** Domestic vs. International - Tax-exclusive

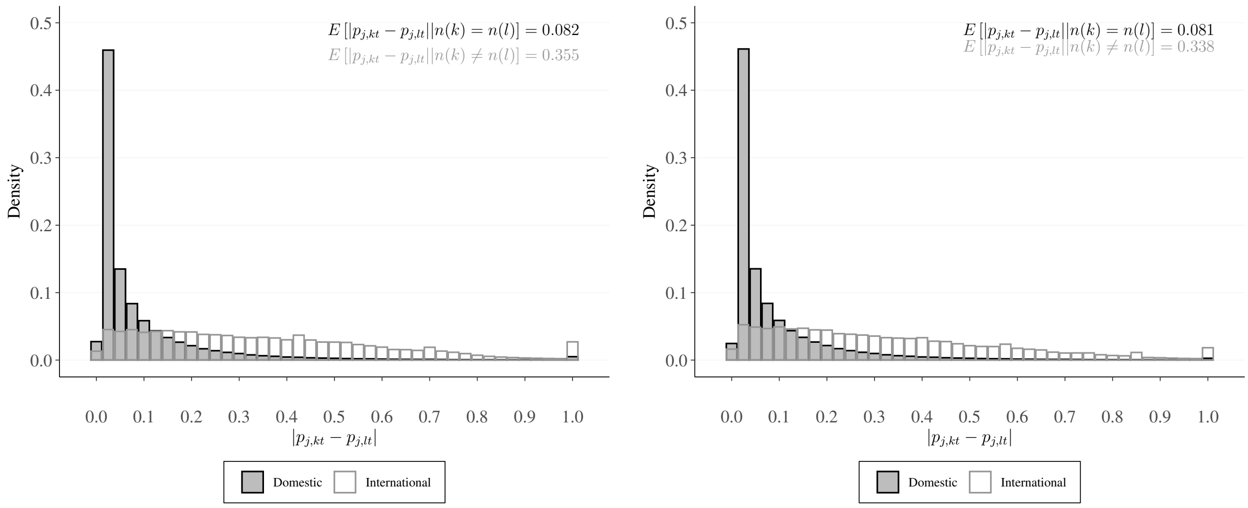

Notes: These figures plot conditional distributions of the absolute log LOP deviations between regional markets, where a regional market is defined at the NUTS2 level. Panel (a) compares the conditional distribution of absolute LOP deviations between domestic and international region pairs. Panel (b) 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.

Border regression discontinuity design Several factors could explain why absolute LOP deviations are larger for international region pairs than for domestic region pairs. A natural explanation is that domestic region pairs are geographically closer compared international region pairs. Other potential reasons are that spatial price discrimination, differences in destination-specific marginal costs or trade frictions are larger between countries, creating price differences that are discontinuous at country borders.

To disentangle these factors, I implement a border Regression Discontinuity Design (RDD). By comparing prices for the same product on either side of an international border, I can assess whether export prices, i.e. prices just across the border, are systematically higher than domestic prices in nearby regions of the same country. A significant price jump at the border would suggest the presence of trade frictions and rule out alternative explanations. First, focusing on price changes at country boundaries eliminates differences in average distance as a likely driver of international price dispersion. Second, standard spatial price discrimination models predict that markups should decrease when market access costs rise, making it unlikely that price jumps at the border are purely driven by pricing-to-market strategies.[12] Finally, while local distribution costs could vary, there is no clear reason why they should be systematically higher in foreign markets than in domestic ones.

I proceed in four steps. First, for each contiguous country pair, I identify products that are sold in both countries. Second, I compute ZIP-code-level weighted average consumer prices, using liters sold interacted with population weights as aggregation weights. Third, I rank observations based on their great-circle distance to the border, normalizing distances so that products sold domestically have negative values and those sold abroad have positive values. Finally, I estimate the following specification:

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

In this specification, $p_{j,t}^{s(j)l}$ is the consumer price of product j at time t when going from product j's source $s(j)$ to location l. Border$^{s(j)l}$ is an indicator for 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. Finally, I include $\lambda_{j,t}$, which are 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 overall RDD estimates, Figure 3 illustrates that spatial price dispersion occurs discontinuously at national borders. Figure 3a presents a binscatter with 30 bins and a first-order polynomial on either side of the French-German border. The figure indicates that while domestic prices gradually increase with distance from the source, prices jump discontinuously at the border, with foreign prices averaging 20% higher than domestic prices. Figure 3b replicates this analysis using after-tax consumer prices. The pattern persists: domestic prices continue to rise with distance, and foreign prices remain, on average, 12% higher than domestic prices.

Table 2 confirms that the insights of Figure 3 also hold in the main sample. The table reports the results from estimating equation (3.1) for after-tax consumer prices on the main sample. Standard errors are clustered at the product level, and their coverage accounts for potential misspecification in bandwidth selection, as argued in Calonico et al. (2014).

Using a first-order polynomial and including all observations within a 1,000 km range from the border, column (1) shows that export prices are, on average, 16.7% higher than domestic prices. This difference is statistically significant at conventional levels and rises slightly to 17.2% with a bandwidth of 500 KM and drops somewhat to 10.7% when I limit the range of observations to be within 100 KM of the border. When I apply the optimal bandwidth selection following Calonico et al. (2014), the estimated difference is 10.1%. Columns (5) to (8) show that these results remain largely unaltered when I use a second-order polynomial to control for distance to the border.

As mentioned in section 2, the main sample excludes observations classified as cross-border transactions. Still, when I also include cross-border transactions, Table B.2 shows that the results are

[12] A zero RDD estimate for price differences does not necessarily imply the absence of trade frictions, as firms could offset trade costs with lower markups. However, a positive RDD estimate provides sufficient evidence of trade frictions.

Figure 3: Border RDD: French-German border

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

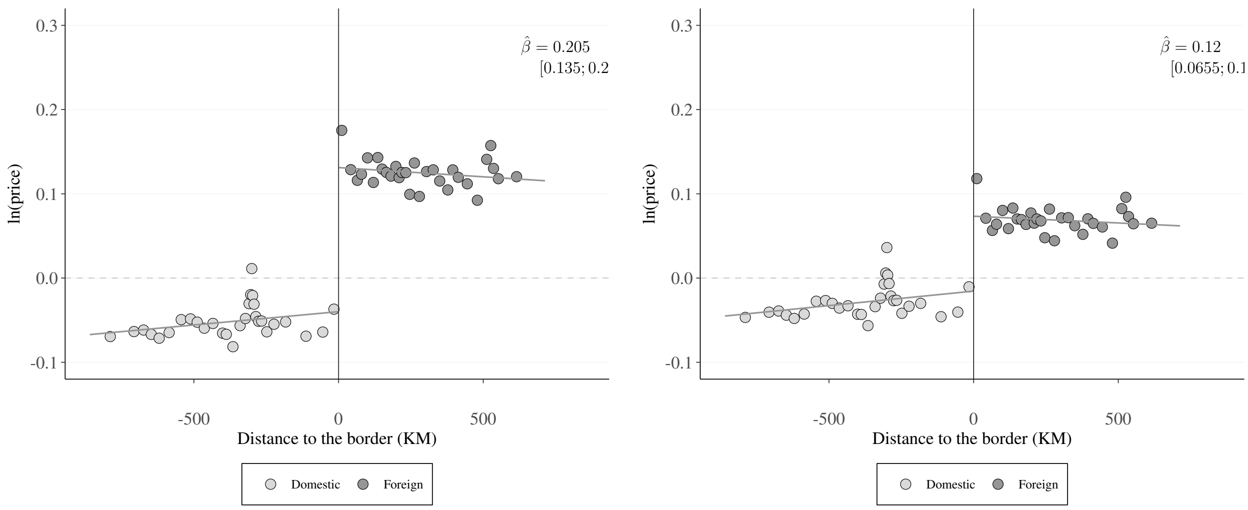

Notes: This figure shows the binscatter corresponding to estimating equation (3.1) with first-order polynomials where I have used 30 equally-spaced bins on either side of the border. I limit the sample to observations within a 1,000km range from the border and to the most important stores and products. 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: 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^{***}	0.172^{***}	0.107^{***}	0.101^{***}	0.179^{***}	0.155^{***}	0.0888^{***}	0.0927^{***}
	$[0.13; 0.203]$	$[0.134; 0.21]$	$[0.0803; 0.134]$	$[0.0755; 0.127]$	$[0.14; 0.218]$	$[0.119; 0.191]$	$[0.0657; 0.112]$	$[0.0683; 0.117]$
$\theta_{j,t}$	✓	✓	✓	✓	✓	✓	✓	✓
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.

very similar to the results presented in Table 2. Moreover, Tables B.1 and B.3 show that foreign prices are still significantly larger compared to domestic prices when I include all products in the restricted set of stores and all products in all stores (B.1). Taken together, these results suggest that an important determinant of between-country spatial price dispersion stems from a discontinuous jump in consumer prices at country borders, consistent with the presence of positive barriers to trade across EU borders.

3.2 Structural gravity

Because structural gravity equations arise 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)), they are often used to estimate the size of trade frictions (see Head & Mayer, 2014; Allen et al., 2020). In the gravity model, trade costs are estimated as the parameters needed to explain differences in trade flows that cannot be attributed to factors such as differences in country size and overall centrality in the global trading system as captured by the multilateral resistance terms.

To be able to benchmark the results of section 5 to this literature, I estimate the tariff equivalent trade frictions as implied by interpreting the variation in trade flows in the bottled water industry through the lens of a gravity model. I consider the following model:

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

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. 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 national 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 resistance terms. Their inclusion also controls for differences in 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, that account for destination-specific demand shocks, and a proxy for transport costs, the gravity equation attributes discontinuous differences in trade flows to barriers to trade across national borders.

Table 3 presents the results of estimating Equation (2) using the main sample and 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) shows that, on average, international trade flows, i.e. trade flows between international region pairs, are 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%.[13] 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; Santamaría et al., 2023). For instance, Santamaría et al. (2023) find that national borders in the EU reduce international trade flows to 9% of the size of domestic trade flows. 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%. Notably, this only holds in the sample with the main products as shown in 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 entirely attributed to trade frictions and not to differences in preferences, market structure or destination-specific marginal costs. Within the gravity framework, barriers to trade affect trade flows via their effect on bilateral prices. At the same time, this mechanism also suggests a test of the validity of the gravity structure. More specifically, after controlling for bilateral prices, the coefficients on Distance^{lk} and Border^{lk} should no longer be statistically different from zero. If they remain statistically significant predictors of trade flows, differences in the aforementioned factors must still be influencing trade flows. Typically, this test is infeasible as bilateral prices are unobserved. However, because I observe destination-specific prices, I can construct theory-consistent bilateral prices after assuming an elasticity of substitution.[14] In columns (4) and (5), I re-estimate Equation (2), now controlling for the vector of bilateral prices, using an elasticity of substitution of 8.28 (following Eaton & Kortum, 2002) and 2.12 (following Boehm et al., 2023), respectively.[15] As expected, higher prices reduce trade. More importantly, 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 these variables.

Two potential violations of the exclusion restriction come to mind. First, conditional on bilateral prices, spatial variation in consumer preferences for the same varieties could lead to differences in trade flows.[16] I address this concern below by estimating a flexible demand system for bottled water. Second, Helpman et al. (2008) shows that trade cost estimates are biased when firms endogenously choose which markets to serve. Column (6) provides suggestive evidence for this by estimating the

[13] 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).

[14] 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.

[15] 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.

[16] For instance, 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.

relationship between the similarity of the set of available varieties and the trade cost variables. I follow Hristakeva (2022) and measure the similarity as the dot product of two vectors normalized by the Euclidean 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. $(\boldsymbol{J}_l)_j = 1$. A similarity of 1 indicates that l and k share all varieties, and 0 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 and similarity drops by 74% when the two regions are separated by a country border within the Eurozone, and by an additional 50% when they also have different currencies.

Table 3: Gravity estimation

	X_t^{kl}					$\frac{\boldsymbol{J}_l \cdot \boldsymbol{J}_k}{\vert\vert\boldsymbol{J}_l\vert\vert \cdot \vert\vert\boldsymbol{J}_k\vert\vert}$
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(1+\text{Dis}^{lk})$	-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.

4 An empirical model of the bottled water industry

This section outlines the empirical model of demand and supply used to recover a model-implied estimate of destination-specific marginal costs, allowing for spatial differences in preferences, market structure and distribution. By combining these estimates with data on production locations, I estimate trade frictions as the residual variation in destination-specific marginal costs that cannot be explained by the other sources of destination-specific cost variation that I control for. In this way, this approach to recovering trade frictions closely follows the spirit of Hsieh & Klenow (2009) and Kalouptsidi (2018).

4.1 Preferences for bottled water

Instead of relying on demand systems traditionally used in international trade, I model consumer preferences for bottled water using a discrete choice framework.[17] An important advantage of 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).

More specifically, 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 from the set of bottled water varieties available in their regional market, which is denoted by $\mathcal{J}_{lt}$. One objection to 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).[18] However, aggregating the data to the quarterly level mitigates this concern, as Table 1 shows that in countries where the bulk of bottled water consumption occurs, on average, 90% of the households buy bottled water each quarter.

Also, defining the choice set as the set of varieties available in their regional market is potentially troublesome as well. This is because country borders within the EU Single Market do not require formal checks and consumers are in principle free to engage in cross-border shopping.[19] Yet, applying this assumption to the bottled water industry is reasonable 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 countries in which cross-border shopping occurs the most are Belgium and the Netherlands but the share of cross-border transactions is well below 5%.

[17] There are a few applications of discrete choice frameworks in international trade. For instance, Fajgelbaum et al. (2011) relies on a discrete choice framework to generate non-homothetic preferences for differentiated products and P. K. Goldberg (1995); P. Goldberg & Verboven (2001); Hellerstein (2008); Nakamura & Zerom (2010) use it to understand the transmission of exchange rate fluctuations into consumer prices.

[18] 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.

[19] 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.

More formally, consumers choose the variety $j \in \mathcal{J}_{lt}$ that maximizes their indirect utility $V_{ij,lt}$ which is given by:

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

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

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

4.2 Market structure and marginal costs

To estimate the distribution of destination-specific marginal costs from the distribution of destination-specific consumer prices, I specify a model of how consumer prices are determined in equilibrium. The implied first-order conditions will 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 I estimate the level of trade frictions. To fix ideas, I now discuss the vertical structure that nests all other variants considered below.

In each regional market and each quarter, consumer prices are determined by downstream retailers as the Nash equilibrium to a simultaneous move game. This implies the following first-order conditions for consumer prices:

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

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 $(\boldsymbol{\Delta}_{lt} \odot \boldsymbol{\Omega}_{lt}^{r})^{-1} \cdot \boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{m}; \boldsymbol{\Theta}^{d})$ is the vector of markups charged by retailers which depends on the matrix of first-order derivatives of demand with respect to consumer prices $\boldsymbol{\Delta}_{lt}$, the ownership matrix $\boldsymbol{\Omega}_{lt}^{r}$ that internalizes substitution towards other varieties sold by the same retailer and the vector of market shares, $\boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{r}; \boldsymbol{\Theta}^{d})$. Each of these objects depends on the demand parameters $\boldsymbol{\Theta}^{d}$, which are estimated below.

I assume that wholesale prices are set by upstream manufacturers as the Nash equilibrium to a simultaneous move game. Accordingly, the first-order conditions that determine wholesale prices are given by:

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

Here, $\boldsymbol{c}_{lt}^{w}$ 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 $(\mathbf{PT}_{lt} \cdot \boldsymbol{\Delta}_{lt} \odot \boldsymbol{\Omega}_{lt}^{w})^{-1}$ is the vector of markups charged by manufacturers. Like retail markups, this depends on the matrix of first-order derivatives of demand with respect to consumer prices $\boldsymbol{\Delta}_{lt}$, an ownership matrix $\boldsymbol{\Omega}_{lt}^{m}$ that internalizes substitution towards other varieties owned by the same manufacturing firm, the vector of market shares, $\boldsymbol{\sigma}_{lt}(\boldsymbol{p}^{r}; \boldsymbol{\Theta}^{d})$. In contrast to the determination of consumer prices, manufacturing markups also depend on the pass-through matrix, $\mathbf{PT}_{lt}$, which consists of the derivative of retail prices with respect to wholesale prices and which captures that manufacturers internalize that optimal retail prices change when wholesale prices change.

Alternative market structures The vertical structure laid out above corresponds to the case of double marginalization, which has been used, for instance, 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 in Atkeson & Burstein (2008); Edmond et al. (2015); Corsetti et al. (2021), but also accounts for differences in retail market structure across locations by allowing retail markups to vary as well. On the flip side, this setup also assumes that retailers take wholesale prices as given, hold no countervailing power relative to manufacturers, and do not internalize the fact that setting markups sequentially lowers overall profits compared to a case where they would maximize joint surplus. à 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 still compete à la Bertrand, 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. As the baseline case of double marginalization is isomorphic to the Nash bargaining in which manufacturers hold all the bargaining power, these 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 different distributions of bargaining power between retailers and manufacturers. Second, I consider the case in which upstream manufacturers still compete à la Bertrand, but retailers set prices equal to the sum of wholesale prices and local distribution costs. This variation will check whether the results are robust to assuming that retailers perfectly pass through costs, as in Corsetti & Dedola (2005) and Loecker & Scott (2022), and to certain forms of non-linear contracting and interlocking relationships that give rise to the same outcome (see Rey & Vergé (2010)).

Marginal costs Substituting equation 5 into 4 and straightforward rearranging allows me to back out destination-specific marginal costs as the sum of downstream distribution costs, production costs and the costs associated with delivering bottles to a 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, as the estimation of barriers to trade requires variation in destination-specific marginal costs, plant-level cost information is unlikely to yield the necessary destination-level variation in costs.

5 Structural estimation

In this section, I start by discussing the estimation and identification of the demand parameters. Armed with these estimated parameters, I turn to 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 to 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 implement this equation by equating $\sigma_{j,lt}$ to the observed probability of purchasing variety j, given by $S_{j,lt}$. I capture the other product characteristics by including different sets of fixed effects, such as 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 determined with knowledge of $\xi_{j,lt}$, creating a 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 in the origin location.[20] The suitability of transport costs as an instrument for consumer

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

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. This is especially true once I allow brand appeal to vary between countries.

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 households with population weights interacted with the number of purchased liters as weights. I cluster standard errors at the region level.

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 preferences for certain brands to vary between countries or when I include retail-chain-time fixed effects, the price coefficient remains virtually unchanged. While the OLS estimates yield downward-sloping demand curves, the implied average own-price elasticities are quite inelastic.

As mentioned before, the OLS results likely suffer from a simultaneity issue due to unobserved demand shocks. Therefore, I turn to the IV results in columns (4) to (6). Column (4) reports the results when I include brand and market fixed effects and instrument consumer prices with transport costs. The first stage F-statistics are large and confirm that, in line with Figure 3, transport costs are an important determinant of consumer prices. After instrumenting prices with transport costs, the estimated demand curves are also 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.

Implied markups To provide some validation for the estimated own- and cross-price elasticities, I benchmark the implied manufacturing markups to the literature. In the appendix, I plot these implied manufacturing markups, defined as the price-to-marginal cost ratio. Based on the estimate of column (6) of Table 4, I find 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 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, however, the appendix also 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.

Table 4: Estimation of preferences

	OLS			2SLS		
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(\text{P}_{j,lt})$	-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.

5.2 Estimation of barriers to trade

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 out in Equation (6). I use this distribution to estimate the magnitude of barriers to trade across national borders 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 highlights that both production costs, $\omega_{j,t}$, which are incurred at the water source and common to all destinations, as well as destination-specific marginal costs may lead to differences in marginal costs between varieties and markets. I consider three sources of destination-specific marginal costs. First, these costs may differ between destinations when local distribution costs differ Burstein et al. (e.g., 2005); Crucini et al. (e.g., 2005); Parsley & Wei (e.g., 2007). I capture such differences by including W_{lt}, which represents labor unit costs in the retail sector and is expressed in Euro per hour. 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. Second, transport costs vary across destinations because markets are located at different distances from the water source. I account for differences in transport costs by including $t_{s(j)lt}$, which are constructed as the interaction between the travel distance from the source to the destination market and the diesel price at the source. Finally, destination-specific marginal costs may differ due to barriers to trade across national borders. To this end, I incorporate $\tau_B \mathbb{1}\left(\text{Border}\right)_{s(j)l}$ and $\tau_C \mathbb{1}\left(\Delta\text{Currency}\right)_{s(j)l}$, which capture barriers to trade between Eurozone and non-Eurozone countries, respectively. As distribution costs, transport costs and, potential barriers to trade are essential for delivering bottled water to the final consumer, they enter the cost function in an additive or Leontief way.[21] Finally, $\eta_{j,lt}$ captures all variety- and market-specific costs that are not captured by the production and destination-specific costs picked up by the aforementiond variables. In addition to production costs and the three sources of destination-specific marginal costs, $\eta_{j,lt}$ captures the residual variation in variety- and market-level marginal costs.

Consistent estimation of the barriers to trade requires that the unobserved portion of the marginal costs, captured by $\eta_{j,lt}$, is not correlated with $\mathbb{1}\left(\text{Border}\right)_{s(j)l}$ and $\mathbb{1}\left(\Delta\text{Currency}\right)_{s(j)l}$. Three possible violations of this assumption come to mind. First, in many industries this restriction will be violated when firms fragment their value chain and opt for destination-specific production locations. However, this is not a concern in the bottled water industry where the production is tied to the source of the water. Moreover, by observing the same variety being consumed in multiple locations, I can flexibly account for the marginal costs of production through variety-time fixed effects and recover cross-border trade frictions under arbitrary returns to scale and scope. Second, there is a growing literature that documents that exporters export products of different quality to different countries by using an input mix that is contingent on the destination market (e.g., Manova & Zhang, 2012; Bastos et al., 2018). Differences in marginal costs across destinations due to differences in unobserved product quality are unlikely to affect my results as products are defined such that all product characteristics that affect prices and costs are accounted for (see Table A.5). Finally, exporters may use different modes of transportation for different destination markets (e.g., Harrigan, 2010; D. L. Hummels & Schaur, 2013). Given the proximity of the countries in my sample and the cost-effectiveness of truck transport at such distances, differences in the mode of transportation are highly unlikely.

Estimation results Table 5 presents the results from estimating Equation (5) under different assumptions about how consumer prices are determined in equilibrium. Each specification includes product-time fixed effects and 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 both the upstream and downstream markets were perfectly competitive such that retailers and manufacturers set prices equal to marginal costs. When I do not include any of the border dummies in column (1), consumer prices intuitively rise when transport costs and local distribution costs rise.[22] When I add a border dummy, which does not

[21] Assuming that distribution costs and transport costs enter the cost function additively is standard in the international trade literature (e.g., Corsetti & Dedola, 2005; Parsley & Wei, 2007; Irarrazabal et al., 2015).

[22] As column (1) effectively implements Equation (5) using final consumer prices, it can also be seen as the first-stage regression of the demand estimation.

discriminate between Eurozone and non-Eurozone countries, in column (2), I estimate that marginal costs rise by 8.6 eurocents when the product is exported versus when it is sold domestically.

Columns (3) to (6) present the results when I allow for deviations from perfect competition either in the upstream or downstream market. First, in column (3), I examine the case of Bertrand competition in both the upstream and downstream markets but without any border dummies yet. Column (3) confirms that marginal costs increase with transport costs and local distribution costs even when differences in retail and manufacturing markups are accounted for. Second, columns (4) to (6) highlight that accounting for spatial price discrimination increases the estimated tariff-equivalent barriers to trade.[23] In particular, I estimate that the marginal cost of supplying the same product to a foreign market is 9 eurocents higher compared to when that product would be supplied domestically. Importantly, this result holds across different vertical market structures, including upstream oligopoly (column (4)), downstream oligopoly (column (5)), and double marginalization (column (6)). Depending on the vertical market structure, the estimated domestic marginal cost of exported varieties averages between 33 and 38 eurocents, implying that the estimated barriers to trade are equivalent to a tariff of 23 and 28%. Notably, these estimates are in the range of the gravity estimates from Table 3.

Column (7) replicates the analysis from column (6) but also investigates heterogeneity in the estimated barriers to trade by distinguishing between trade barriers between Eurozone and non-Eurozone countries. Column (7) uncovers substantial heterogeneity in the estimated cross-border barriers to trade, showing that they amount to 23 eurocents between non-Eurozone countries. However, column (7) also confirms that estimated trade frictions are in large part also driven by barriers to trade between Eurozone countries as the estimated difference in the marginal cost between foreign and domestic markets remains high at around 7 eurocents.

Taking stock The results presented in Table 5 have two key implications. First, cross-border barriers to trade are lower between Eurozone countries than between non-Eurozone countries. This provides credence to models that consider a reduction in barriers to trade as a key gain of joining a currency union. Second, cross-border barriers to trade between Eurozone countries are also quite large. This suggests that further geographic market integration within the European Single Market could still generate significant welfare gains. I quantify the increase in consumer welfare that such an integration would bring in the next section.

6 Equilibrium effects of market segmentation

In this section, I explore the effect of barriers to trade on equilibrium price dispersion and consumer welfare through a counterfactual exercise. More specifically, I compare the current level of

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

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)
$t_t^{s(j),l}$	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)
$E\left[\cdot\|\text{Export}_j = 1\right]$	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.

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 one in which the current set of available varieties is supplied at destination-specific marginal costs. In this benchmark, marginal cost consists of two components.

First, there is the marginal cost of production. Relative to the estimation of the cost function, additional assumptions are necessary. So far, the assumptions on the marginal costs of production have been minimal as the product-time fixed effects flexibly accounted for production costs under arbitrary returns to scale (see Equation 5). However, in the absence of trade frictions in the efficient benchmark, the scale of production might differ from what is observed in the data. If so, the product-time fixed effects would no longer accurately capture production costs. For this reason, I now assume constant marginal costs, allowing me to reuse the product-time fixed effects $\lambda_{j,t}$ to construct marginal cost of production.

Second, while the marginal cost of production is identical across destination markets, there will

Table 6: Counterfactual excercises

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

be non-zero spatial price dispersion even in the integrated benchmark economy. This is because transport costs and local distribution costs in terms of labor significantly influence the destination-specific marginal costs (see Table 5) and differ between destinations.

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 international trade is subject to the estimated trade frictions. Two key findings emerge. First, cross-border barriers to trade increase between-country spatial price dispersion: 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 remains substantial, averaging 34% across countries. This has two implications. On the one hand, even if cross-border barriers to trade were eliminated, significant equilibrium price dispersion would persist purely due to differences in distribution and transport costs. On the other hand, the fact that substantial price dispersion exists even in the integrated benchmark suggests that the mere presence of spatial price differences at market boundaries is not necessarily evidence of trade barriers or geographic market segmentation along national borders.

I quantify the effect of cross-border trade barriers on consumer welfare by comparing consumer welfare in the integrated benchmark to that in the segmented economy, where international trade is subject to the estimated barriers to trade. Table 6 shows that expected consumer welfare in the segmented economy is 3.6 eurocents lower than in the integrated benchmark. While this difference may seem small in absolute terms, it is economically significant: given that the average consumer price is around 37 eurocents per liter, this loss is equivalent to a 10% tax on bottled water consumption.

7 Conclusion

Spatial price dispersion is common both within and between countries. Whether such dispersion, particularly at market boundaries, reflects geographic market segmentation and should be reduced depends on its underlying causes. If spatial price dispersion merely reflects differences in the marginal cost of supplying different markets, it will arise even in an integrated benchmark. Only when dispersion is driven by trade frictions does it indicate market segmentation, and only in that case would reducing it improve welfare.

This paper investigates whether trade within the European Single Market is still hindered by trade frictions and, if so, whether these frictions induce inefficient price dispersion and reduce consumer welfare. I address this question by studying the bottled water industry, leveraging institutional features to construct a unique dataset on prices, quantities, and production locations at the product level. I combine these data with an empirical model of demand and supply, identifying barriers to trade from model-implied marginal costs across domestic and foreign destinations while allowing for cross-country differences in preferences, market structure, and arbitrary returns to scale in production.

I estimate that cross-border trade barriers amount to 9 cents per liter between Eurozone countries and 32 cents per liter between non-Eurozone members. These barriers increase equilibrium price dispersion by 5% relative to an integrated benchmark in which bottled water is priced at its destination-specific marginal cost. The associated welfare loss is equivalent to a 10% tax on bottled water consumption.

While these results are derived from the bottled water industry, their implications extend beyond it. Cross-border trade barriers not only distort the marginal cost of exporting but also affect the geographic allocation of production. In industries with more flexibility in plant location, firms facing high export costs may opt for more dispersed—and potentially less efficient—production closer to consumers rather than fully exploiting economies of scale.

In this paper, I interpret residual differences in model-implied destination-specific marginal costs as cross-border trade frictions. Although my empirical model is more flexible than standard trade models —- allowing for differences in preferences, market structure, and production technology – these barriers remain a construct of the model used to rationalize the data. Crucially, such price differences persist because consumers and firms do not engage in profitable arbitrage. While arbitrage is unlikely in the bottled water industry due to its low value-to-weight ratio, an important open question is which economic or institutional features of the European Single Market limit arbitrage in industries with higher value-to-weight ratios that are subject to similarly large price differences between countries.

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. *The American Economic Review*, *98*, 1998-2031.

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

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

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

Beck, G., Kotz, 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.

Coşar, K., Grieco, P., 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., Telmer, C., & 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. (2015). Competition, markups, and the gains from international trade. *American Economic Review*, *105*, 3183-3221. doi: 10.1257/aer.20120549

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

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

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

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

Gorodnichenko, Y., & Tesar, L. (2009). Border effect or country effect? seattle may not be so far from vancouver after all. *American Economic Journal: Macroeconomics*, *1*, 219-241.

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

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., & Yeaple, S. (2004). Export versus fdi with heterogeneous firms. *American Economic Review*, *94*, 300-316. doi: 10.1257/000282804322970814

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. (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. doi: 10.1086/422562

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. doi: 10.1162/REST_a_00517

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., Khandelwal, A., & 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. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, *71*, 1695-1725. doi: 10.1111/1468-0262.00467

Miller, N., & Weinberg, M. (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., & 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., & 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., 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., & Keller, W. (2007). Markets in china and europe on the eve of the industrial revolution. *American Economic Review*, *97*, 1189-1216. doi: 10.1257/aer.97.4.1189

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. doi: 10.1093/qje/qjw037

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

$y_{ij,lt}$	$\mathbb{1}\left(\sum_{j\in\mathcal{J}} q_{ij,lt} > 0\right)$	$\ln p_{ij,lt}$	$\mathbb{1}\,(\text{Flavored})_j$	$\mathbb{1}\,((750,1500)\text{ml})_j$	$\mathbb{1}\,(\geq 1500\text{ml})_j$
	(1)	(2)	(3)	(4)	(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

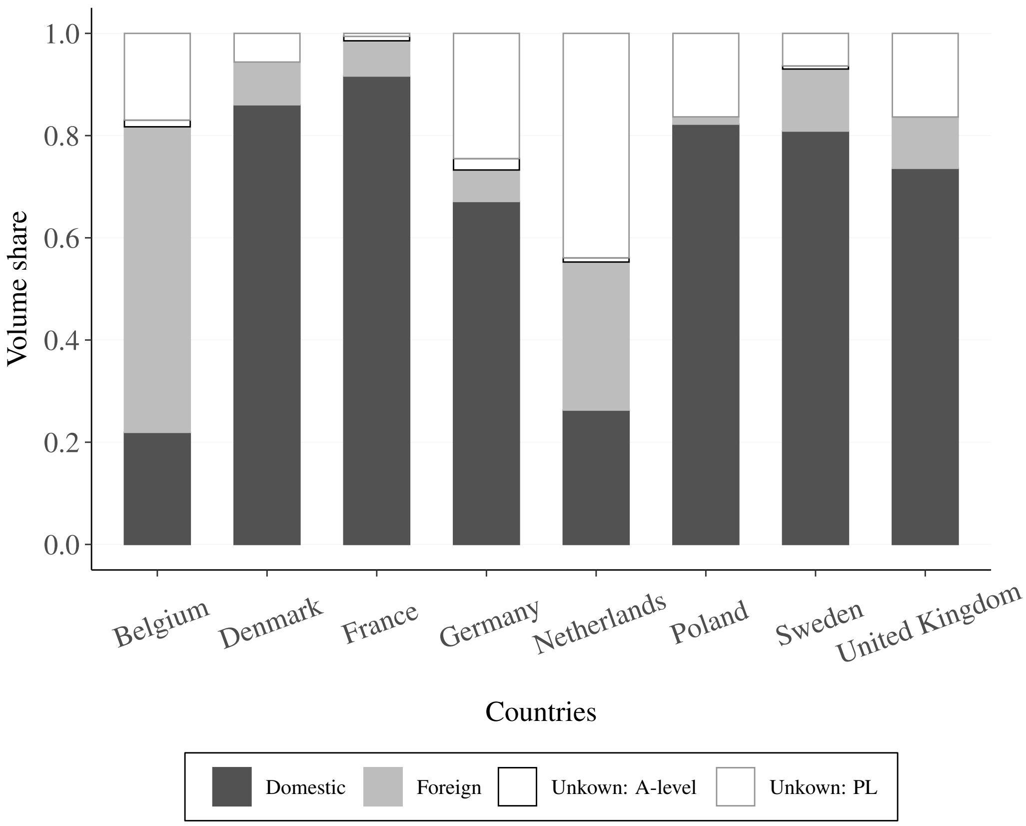

Notes: This figure plots for each country the distribution of volume sold across four groups: domestically sourced water in dark grey, foreign-sourced water in light grey, A-level brands with an unknown origin in white with a black contour and private label water with an unknown source in white with 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, $\mathbb{1}\left(i = \text{Male}\right)_{n(i)}$ 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

$p_{j,lt}$	1^{st}-order				2^{th} order			
	(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

$p_{j,lt}$	1^{st}-order				2^{th} order			
	(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}					$\frac{\boldsymbol{J}_l \cdot \boldsymbol{J}_k}{\lVert \boldsymbol{J}_l \rVert \cdot \lVert \boldsymbol{J}_k \rVert}$
	(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)
$\ln(\text{Price}_t^{lk})$				-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+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(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}					$\frac{\boldsymbol{J}_l \cdot \boldsymbol{J}_k}{\|\|\boldsymbol{J}_l\|\| \cdot \|\|\boldsymbol{J}_k\|\|}$
	(1)	(2)	(3)	(4)	(5)	(6)
$\ln(1+\text{Dis}^{lk})$	-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)
$\ln(\text{Price}_t^{lk})$				-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.