KU LEUVEN

DEPARTMENT OF ECONOMICS

Evaluating the impact of market integration - Banning online trade restrictions in the EU portable PC market

Néstor DUCH-BROWN, Lukasz GRZYBOWSKI, André ROMAHN and Frank VERBOVEN

FACULTY OF ECONOMICS AND BUSINESS

DISCUSSION PAPER SERIES DPS20.10

JUNE 2020

Evaluating the Impact of Market Integration – Banning Online Trade Restrictions in the EU Portable PC Market*

Néstor Duch-Brown† Lukasz Grzybowski‡ André Romahn§ Frank Verboven¶

June 9, 2020

Abstract

We develop a framework to evaluate the impact of market integration, accounting for spillovers between multiple distribution channels. Our adaptation of the standard random coefficients logit demand model allows for substitution between distribution channels and incorporates consumer arbitrage across countries. We apply our framework to the European portable PC market, where geo-blocking practices that restrict online trade have recently been banned. The total consumer and welfare gains from reducing cross-border arbitrage costs are relatively modest, and entirely due to increased product choice rather than reduced price discrimination. At the same time, the distributional effects from the cross-country price convergence are substantial. Consumers in high income countries gain most, while consumers in medium and low income countries are only marginally better or even worse off.

*We would like to thank Joel Stiebale, Jo Van Biesebroeck and conference participants at EEA/ESEM (Zurich), CRESSE (Greece), Workshop in IO and Economic Theory (Chile), 10th Digital Economics Conference (Télécom ParisTech) and seminar participants at MINES ParisTech for helpful comments and suggestions. All remaining errors are our own.

†European Commission, Joint Research Center Seville; nestor.duch-brown@ec.europa.eu

‡Télécom ParisTech; lukasz@mushroomski.com

§University of Düsseldorf, DICE; romahn@hhu.de

¶KU Leuven, Department of Economics; frank.verboven@kuleuven.be; corresponding author

1 Introduction

Many consumer goods markets may remain nationally segmented, even without import tariffs because of important non-tariff barriers to trade. Such barriers are not only due to protective national government regulations. They can also be the result of deliberate strategies by firms to raise consumer cross-border trade costs, for example through distribution agreements that restrict retailers to sell in other countries. Free trade areas have taken various actions against firms that engage in such restrictive practices.[1] Policy makers that actively promote cross-border trade by consumers make the implicit presumption that full market integration would make markets more competitive to the benefit of all consumers. Economic theory, however, suggests that this is not so obvious. First, removing opportunities to engage in price discrimination may benefit consumers in some countries at the expense of consumers in other countries. Second, the impact on overall welfare is ambiguous, especially in the presence of oligopolistic behavior.

In this paper, we develop an empirical framework to assess the impact of reducing cross-border trade costs in nationally segmented markets. We are particularly interested in the situation where only one distribution channel (the online channel) becomes more integrated, while other channels (traditional "brick-and-mortar" channels) remain segmented.

We are inspired by a recent policy in 2018 of the European Commission, which put a ban on widespread geo-blocking practices. Such practices restrict consumers from purchasing products online in other countries. They were held responsible for the limited cross-border trade in online markets and for preventing the rise of a single European digital market. A ban on geo-blocking can thus make online markets more integrated, without directly affecting segmentation in the traditional distribution channel. In earlier investigations, the Commission indeed found that online cross-border shopping was very limited despite large cross-country price differences, notably in the markets for consumer electronics. For example, according to Eurostat in 2015 only 1.6 percent of consumers had ordered computer hardware from a different EU country. A mystery shopping survey carried out in 2015 on behalf of the European Commission found that 79 percent of cross-border shopping attempts for consumer electronics products were geo-blocked.

Our framework to assess the impact of reducing cross-border trade costs in online markets starts from a differentiated products demand model. We explicitly model the fact that consumers can purchase their products at two distribution channels: the traditional and the online channel. Furthermore, after integration consumers can purchase their products online in the other EU countries. We could in principle make use of the standard random coefficients logit demand model of Berry (1994) and Berry, Levinsohn and Pakes (1995; henceforth BLP). However, such an approach would not be warranted in our setting. It would involve a very high dimensional (type 1 extreme

[1] A well-known example is the automobile industry, with cases in both Europe and North-America on vertical restraints that were imposed by manufacturers on their dealers to restrict selling to customers in other countries.

value) individual taste term, that is not only specific to each product but also to the distribution channel and country of purchase where each product is available. This would lead to undesirable substitution patterns between the same products sold in different distribution channels and (after integration) in different countries. Even more importantly, it would lead to misleading welfare conclusions when enlarging the choice set to the same products sold online in other countries. To tackle these issues we propose an adapted BLP model, which reduces the dimensionality of the individual taste term to only the product, rather than the product, channel and country. We show how to accurately approximate this adapted BLP model by a random coefficients nested logit model, where each product constitutes a nest and the nesting parameter is set close to 1. Our approach relates to the pure characteristics model of Berry and Pakes (2007), but is easier to estimate.

We apply our analysis to the market for portable PCs in 10 EU countries during 2012-2015, a period where there was strong national market segmentation because of the geo-blocking practices. Our preliminary evidence shows that international price differences for identical products were large, leaving substantial scope for online consumer arbitrage. In fact, in more than half of the cases the lowest online prices prevailed in the low-income countries Poland and Slovakia, while the highest prices were more frequent in high-income countries such as Belgium, Denmark and the Netherlands. We then estimate our adapted BLP demand model in the presence of national market segmentation. Consistent with the documented price differences, we find that consumers are most price sensitive in the group of low-income countries, while consumers are the least price sensitive in the high-income countries. We also find that there is substantial heterogeneity in the valuation for the online distribution channel. This implies that there is only moderate substitution between both distribution channels (though more than in a standard BLP model that would assume consumer heterogeneity that is specific to both the product and the online channel).

After adding a supply side with oligopolistic price-setting behavior, we evaluate the impact of a ban on the geo-blocking restrictions. The ban makes online markets integrated as it enables consumers to purchase online in other countries (possibly at an extra shipping cost). The ban can have both direct effects on prices in the online channel, and possibly indirect effects on prices in the traditional channel (through the fact that both channels are substitutes). We decompose the policy's impact into two main components: a price convergence and a choice expansion effect. The first effect assumes that consumers can only make purchases abroad for products that they could previously already purchase in their own country. This effect focuses on consumer arbitrage and how it induces product-level price convergence and eliminates third-degree price discrimination (based on the consumer's location or country of residence). The second effect assumes that consumers can also choose products in other countries that they had not available before. This effect incorporates how consumers may benefit from increased product availability.

Our main findings can be summarized as follows. First, we find that the total EU effects of integrating online markets are relatively modest, with limited overall average price decreases and output increases. The total consumer benefits are mainly due to the product choice expansion effect rather than the price convergence effect that reduces or eliminates price discrimination. In our setting, the spillover effects of online market integration to the traditional channel are small, because of our finding of substantial heterogeneity in the valuation of online, which implies that most consumers stay with their own preferred channel. While the overall effects are relatively modest, this may change in the future as e-commerce made up only 20 percent of the market but continues to gain in popularity. Furthermore, the geoblocking ban applies to a wide range of retail categories, so that the total effects can add up to a substantial amount.

Second, we find substantial distributional effects of the policy on consumers and firms. Online prices drop on average by 1.5 percent in high-income countries, while they increase by on average 7.9 percent in medium- and by 12 percent in low-income countries. This indicates a redistribution from consumers in low-income to high-income countries because of the price convergence effect. However, the choice expansion effect counterbalances this because more products become available online. As such, consumer welfare in low-income countries remains roughly unchanged, while consumer welfare in high-income countries increases even more strongly after taking into account expanding product availability. Total firm profits drop slightly by .3 percent with profits lost being greater without the choice expansion effect.

These findings are based on our adapted BLP model, which eliminates the artificial individual taste valuations for the channel and country-of-purchase of every product. We show that a standard BLP model that includes such idiosyncratic valuations would lead to misleading conclusions because it mechanically includes gains from additional variety that is specifc to each product, distribution channel and country-of-purchase. First, this would substantially overestimate the effects on consumer welfare. Second, this would also imply that firms would actually greatly benefit from opening up online markets across the EU. Such a prediction is at odds with a simple revealed preference argument, as firms deliberately chose to impose the geo-blocking restrictions to segment online markets.

Our paper contributes in several ways to the literature on international price differences and the law of one price in imperfectly competitive markets.[2] This literature has made various advances in understanding the sources of international price differences (local costs versus markups), e.g. Goldberg and Verboven (2001) for cars; Kanavos and Font (2005) for pharmaceuticals; Gopinath, Gourinchas, Hsieh and Li (2012) for a grocery chain; and Goldberg and Hellerstein (2013) for beer. However, there has been limited attention to the role of cross-border trade costs in obtaining market integration. Our first contribution to this literature is to provide a framework for empirically evaluating the impact of a reduction (or entire removal) of cross-border trade costs on international price differences and welfare. This is distinct from interesting recent work by Dubois, Gandhi and Vasserman (2019), who consider how direct price regulations may affect price differences between

[2] Early work started at the aggregate level with Frankel and Rose (1996) and Obstfeld and Rogoff (2000).

countries without involving cross-border trade.

The literature on international price differences has largely focused on traditional, brick-and-mortar sales channels.[3] There are only a few contributions describing international price differences in online markets, e.g. Gorodnichenko and Talavera (2017) for a range of elecronic products sold in the US and Canada, Duch-Brown and Martens (2014) for household appliances and Duch-Brown, Grzybowski, Romahn and Verboven (2019) for various electronic products sold in the European Union. Our second contribution is to show how our framework can be used to evaluate the indirect spillover effects of a reduction in cross-border trade costs in one distribution channel on international price differences in another distribution channel.

To model consumer demand across distribution channels and countries in a sensible way, we suggest an adapted BLP model. This closely relates to Song (2015), who imposes the individual taste term to be common to products of the same brand.[4] His setting does not allow him to estimate multiple random coefficients, and appears to be computationally cumbersome (for example requiring a very large number of simulation draws). Our approach overcomes these difficulties by approximating the adapted BLP model with a limiting version of a random coefficients nested logit model. We show that for a nesting parameter that is imposed to be sufficiently close to one, the approximation becomes very accurate (i.e. close to the true adapted BLP model).

The remainder of the paper is organized as follows. Section 2 describes the relevant institutional background on cross-country trade restrictions and geo-blocking, and provides preliminary evidence on the scope for arbitrage in the portable PC market. Section 3 provides a general overview on how to model demand in segmented versus integrated markets. Section 4 presents our model and empirical findings on substitution patterns and competition under segmented markets. Section 5 develops our counterfactual approach and discusses our findings of the impact of introducing online market integration. Section 6 concludes.

2 Institutional Background and Data

We first provide a brief description of policies to integrate markets in Europe, and the recent ban on geo-blocking practices. Next, we describe our dataset on the market for portable PCs. Finally, we provide some key information relevant for our empirical analysis.

[3]The literature on e-commerce has focused almost exclusively on price dispersion at the national level, showing that online markets do not exhibit smaller price dispersion than online markets, see e.g. Pan et al. (2004) for a review of the early literature on this topic.

[4]Grubb and Osborne (2015) provide an application where the individual taste term is common to all tariff plans of the same mobile operator, using a maximum likelihood framework. Thomassen (2017) suggests an approach where not only the individual taste parameter is imposed to be common to all engine variants of the same car model, but also the different variants' unobserved characteristics are forced to be equal.

2.1 Cross-country Trade Restrictions and Geo-blocking

One of the cornerstones of the European single market is the achievement of free movement of goods (the other ones being free movement of capital, services and labor). After removing all import tariff barriers to create a single customs union, the European Union focused on reducing a large number of non-tariff trade barriers. Part of these efforts focused on forcing national countries to take steps to harmonize their national legislations, which often implicitly created obstacles to cross-border trade (e.g. differing national product requirements). At the same time, the European Commission has taken numerous actions against private firms for anti-competitive practices that prevented cross-border sales through export restrictions. This has resulted in large fines in many competition cases, including major companies in a variety of industries, such as automobiles (including the 102 million euro fine to Volkswagen in 1998 and the 72 million euro fine to DaimlerChrysler in 2001, beer (with the 200 million euro fine to AB Inbev in 2019) and card payments (fine of 570 million euro to Mastercard in 2019). With the rise of e-commerce, cases also emerged against companies preventing cross-border online shopping, as illustrated by the 40 million euro fine to clothing company Guess in 2018 for preventing consumers to shop online in other countries. The restrictive trade practices by private companies have often prompted the Commission to conduct sector-wide investigations to arrive at guidelines or binding regulations.

Against this background, the European Commission (2017) published a report on the e-commerce sector inquiry in 2017, as part of its broader goal of achieving a Digital Single Market. The investigation highlights that manufacturers increasingly make use of: (i) own online shops, (ii) selective distribution to control their distributors, and (iii) various contractual resrtrictions to control distributors. The Commission showed a particular concern with the widespread use of geo-blocking practices. Geo-blocking practices are actions taken by manufacturers or retailers to restrict cross-border online trade. Based on the visitor's IP adddress, firms can block consumers from access to foreign websites, they can re-route them to the local version of the same online store, or simply refuse to deliver cross-border or refuse payment from a foreign bank. In a Mystery Shopping Survey carried out by GfK, the European Commission (2016) indeed found that geo-blocking was very common in the markets of consumer electronics. It found that 79 percent of cross-border shopping attempts for consumer electronic products were geoblocked.[5]

[5]The GfK Mystery Shopping Survey collected in total 10,537 observations for cross-border online shopping attempts for 147 different country pairs. From each EU country, between 200 and 600 shopping attempts were tested, depending on the relative importance of the country in total (estimated) on-line cross-border trade in the EU. The country pairs were chosen primarily to represent the major online trade routes within the EU28. Mystery shoppers were assigned a website and two products. First, they tested the website and the availability of the two products as a domestic shopper in the country of establishment of the webshop. Via a VPN network, they accessed the targeted webshop with a domestic IP address of the shop's country and recorded the information on the availability of the assigned products, the price, delivery costs and payment options. Then the IP address was changed to the country of residence of the buyer to test whether a cross-border shopping attempt could be completed successfully. From this foreign IP address, the mystery shoppers put the assigned product into the shopping basket and performed all steps to complete the order.

To overcome these cross-border frictions in the online distribution channel and as part of its Single Digital Market strategy, the European Council adopted a new regulation 2018/302 that bans unjustified geoblocking within the EU internal market; see EU Regulation (2018). The regulation became effective on 3 December 2018 and expressly forbids that a consumer located in one Member State is blocked from ordering a product in an online store located in any other Member State.

2.2 Data

We use a panel dataset from GfK for the market of portable computers, with monthly information for 10 EU countries at the product level. Our monthly data cover the period between January 2012 and March 2015. The 10 EU countries are: Belgium, Denmark, France, Germany, Italy, the Netherlands, Poland, Slovakia, Spain, and the United Kingdom. The product-level data consist of sales, prices and various product characteristics, broken down by two distribution channels: the traditional or "brick-and-mortar" channel and the online channel.[6] GfK collected this information from a comprehensive sample of retailers, covering 87 percent of total portable PC sales in these countries.[7] Each portable PC or "product" is described by three identifiers: (i) the brand, such as Acer or Sony; (ii) the series, such as Aspire or Travelmate in the case of Acer, and (iii) the model, such as 7540G - 304G50MN or 7750G - 32314G75MN in the case of Acer Aspire.

An observation in our panel dataset is thus a product (brand-series-model), distribution channel (traditional or online), country and period (month). The initial data set includes 931,509 observations. We aggregate sales for duplicated products and for products with minor variations in model code, which upon inspection is sometimes caused by different coding conventions between countries and sometimes by minor differences in product attributes, such as the color of the chassis. To reduce the computational burden of the estimation and because we find variation in market shares to be limited within quarters, we limit the months in our sample to February, May, August and November. Moreover, we remove observations with very small market shares, such that 1.5 percent of total units sold are dropped. To exclude netbooks, which are small and low-priced laptops that are primarily designed for web browsing only, we censor the price distribution at 400 euros. During the sample period it is also very unusual for a laptop to be sold at a price of more than 2,000 eu-

[6] GfK uses a "point of sales tracking" technology, which reports which products are sold, when, where and for how much, both at online and offline outlets on monthly (or sometimes weekly) basis. The data was collected directly from the electronic point of sales systems from retailers and resellers. Sales were tracked at the individual stock keeping unit level and coded with a full set of features using a cohesive international methodology to allow for accurate comparison both within and across European markets. Any brand or model which was found to be sold in the covered countries is tracked, unless the brand is exclusive, in which case it is still audited but with a label which hides its exact origin. Sales volumes and turnover per item were gathered at the same time as the model specification information. The price of the item was calculated as turnover divided by units sold.

[7] The available data are aggregated across retailers and cover the following types of retailers: system houses, office equipment retailers, computer shops, consumer electronics stores, mass merchandisers, pure internet players, mail orders/online catalogues. It does not include: duty free shops, gas stations, door to door, street markets, discounter stores and direct sales (to staff, hotels, schools, hospitals, etc.). The sample is representative both for the smaller independent sellers as well as for the large chain-stores.

ros, and we drop these "high-end" observations. The final data set consists of 10,288 observations on products, distribution channels, countries and periods. The number of unique products across channels, countries and months in the entire sample is 186.

For each observation, we observe the quantity sold, price and several observable characteristics: the included CPU's speed, the amount of RAM, the laptop's weight, its outer diagonal and its display resolution. As we do not observe the display diagonal for all the models in our data set, we infer the diagonal from each laptop's outer measurements. Moreover, we compute the display resolution from this inferred display diagonal. This is very close to the reported numbers for observations where we do have information (with a slight overstatement for all products).

Table 1 presents summary statistics for these variables. The average sales of a portable PC is 12,530 in the traditional channel, compared with 3,500 in the online channel. Median sales are considerably lower, indicating a skewed distribution of sales towards a more limited number of top selling models. Minimum sales are zero; these are products sold at only one channel. The average price of a portable PC is comparable at both distribution channels (779 euro at the traditional channel and 767 at the online channel). The large variation in prices stems in part from a considerable variation in the product characteristics, but also from variation across countries and over time. The final rows show the number of products by market (J) and by channel and market (J_{trad} and J_{on}). The majority of products tend to be more widely available in the traditional brick and mortar stores than through the online channel. For example, the median number of products in a given market is 46.5 at the traditional channel, while it is only 31 at the online channel. Approximately 26.4 percent of all observations are accounted for by single-channel products, which means that in the corresponding market the product is only available in one of the two distribution channels.

Table 1: Quantities, Prices and Product Characteristics

	min	10^{th} percentile	median	mean	90^{th} percentile	max
$q_{j,\text{trad}}$	0	83.4	2 262	12 530	39 541	186 603
$q_{j,\text{on}}$	0	0	694	3 500	9 882	47 144
Price (traditional, euro)	400	464	681	779	1 195	1 989
Price (online, euro)	400	461	667	767	1 202	1 999
CPU speed (GHz)	1.03	2.08	2.55	2.57	3.13	3.67
RAM (GB)	2	4	6	6.77	8	16
Weight (kg)	0.58	1.29	2.31	2.17	2.72	4.68
Diagonal (inch)	10.9	15.1	17.5	17	18.4	20.5
Display Resolution (ppi)	74	85	94	101	124	217
J	24	33	49	50	66.5	85
J_{trad}	24	32	46.5	47.1	62.5	76
J_{on}	8	15	31	32.0	50.5	68

Note: Based on 10 288 observations. The distributional information for product-level units sold in the two distribution channels, $q_{j,\text{trad}}$ and $q_{j,\text{on}}$, is based on summing each laptop model's unit sales between countries for each date in the sample. CPU speed and RAM are respectively measured in gigahertz and gigabyte. Weight is measured in kilograms. The diagonal is measured in inches and is based on the outer dimensions of each laptop's body, which gives us measures that are larger than the actual display diagonal. Display resolution is measured in pixels per inch and we use the inferred display diagonal to compute this quantity, so that all resolutions are lower than the actual display resolution.

2.3 The Scope for Cross-Border Arbitrage

While we study demand and pricing in both retail channels, our main interest is in the online channel. The left panel of Figure 1 shows the market shares of the online channel for portable PC sales in the various countries of our analysis, for the first and last month of our data set (January 2012 and March 2015). There are substantial cross-country differences in the popularity of online. The online market share exceeds 25 percent in Denmark, Germany, the Netherlands and Slovakia near the end of our sample period, while it is still relatively limited in Belgium, Italy and Spain. The online share tends to be growing in most countries. Notable exceptions here are Denmark and Germany, which already started at higher online shares.

Figure 1: The Online Channel's Market Share and Average Price Differences Relative to Germany

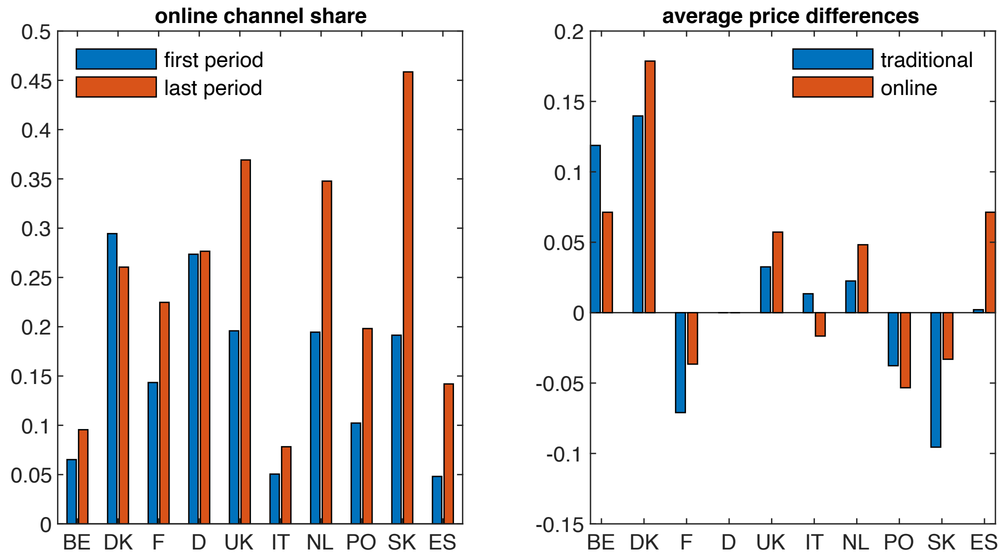

Note: The left panel shows the share of unit sales accounted for by the online distribution channel, where the countries are coded as follows: Belgium (BE), Denmark (DK), France (F), Germany (D), Italy (IT), the Netherlands (NL), Poland (PO), Slovakia (SK), and Spain (ES). The right panel is based on the last period in the sample and Germany is the reference country, so that the German offline and online price differences are both zero by construction. Qualitatively, the right panel is unchanged for other dates in the sample.

The right panel of Figure 1 shows that there are considerable price differences for portable PCs across countries, taking Germany as the base. Interestingly, this is not only the case at the traditional sales channel (where one may expect cross-border shopping to be more difficult), but also at the online channel. Belgium and especially Denmark are on average more expensive, while most notably the Eastern European countries Poland and Slovakia tend to be less expensive. Note that these cross-country price differences tend to be persistent over time, as shown in detail for a larger set of consumer electronics categories by Duch-Brown, Grzybowski, Romahn and Verboven (2019).

These average price differences show a relationship with per capita median income levels. As shown in Figure (A.1) in the Appendix, countries can be divided into three groups: low income (Poland and Slovakia), medium income (Spain and Italy) and high income (the other countries). As we are interested in understanding the sources of price differences before assessing the impact of removing cross-country trade costs, we allow for differences in price sensitivities across these country groups in our empirical analysis.

The cross-country price differences give an indication of the average consumer benefits from shopping abroad. To show the full scope of cross-border arbitrage possibilities, we implement the following exercise. For each product sold in the online channel and each time period, we

Figure 2: Relative Price Difference to Minimum Price Observation at the Product-Level

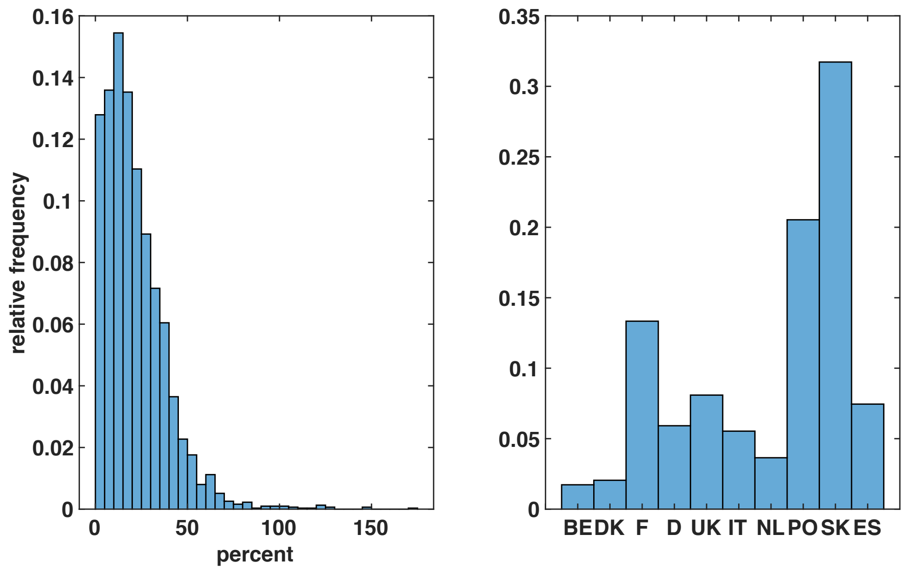

Note: Based on 3 127 online price differences. Models with only one online observation in the cross-section are excluded. The relative price differences are based on model-level online price distributions in each cross-section of the sample. For each model, the bilateral differences are computed relative to the lowest observed online price for that model. Thus, with up to ten online observations for a laptop model, there are up to nine relative price differences per laptop model and cross section. The right panel shows the frequency with which each country's online channel contains that minimum (cross-sectional) price.

determine the lowest available price and compute the percentage price differences between all other online prices for the same product and that minimum price. The left panel of Figure (2) plots the distribution of these relative price differences across products, while the right panel shows the frequency at which each country's online channel contains the lowest price.

The mean and median relative price differences are respectively 21.5 and 17.8 percent. With an average price for a portable PC close to 800 euros, this indicates potential monetary savings of on average 166 euros. As even higher price differences are common, this emphasizes that consumers have a strong incentive to purchase their preferred laptop model abroad. Of course, if a ban on the geoblocking practices would make this feasible, firms may respond by adjusting the cross-country price differences, which is the focus of our counterfactual analysis.

The right panel of Figure 2 shows which countries' online channels tend to offer the lowest prices. Not surprisingly, Poland and Slovakia, which have the lowest average prices and also make up the low-income group, are the lowest price country for more than half of the products, so the incentives to shop online in these countries are strongest. This underlines that these low-income countries also tend to have access to the lowest prices when firms price to market within national borders. With the exception of France, all other high-income countries account for a very small

Figure 3: Fraction of Products Covered at the Online Distribution Channel

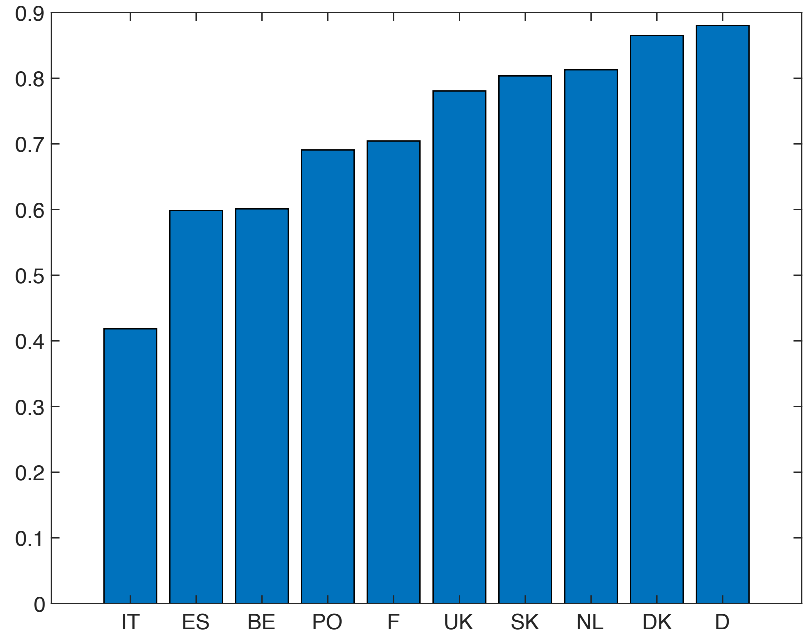

Note: This shows the fraction of products that is covered in each country, relative to the total number of available products across countries, averaged over all periods.

share of the minimum online price observations. This is again consistent with firms implementing price discrimination between populations that have different price sensitivities.

The computed percentage price differences between the cheapest and the most expensive country conceal the fact that some products may be widely available, while other products may be available in only a few countries. From the perspective of consumers, this offers a new reason to shop online: take advantage of an increase in variety in other countries, rather than only a reduced price for products already available in their country of residence. Figure 3 sheds some light on this. It shows that the coverage of online products is indeed limited, and varies across countries. Denmark and Germany have the highest diversity of products at the online channel, covering close to 90 percent of the products available across the ten countries. At the other end of the spectrum are Spain and Belgium (with 60 percent coverage) and especially Italy (only 40 percent coverage). In our counterfactual analysis, we will take into account that a ban of the geo-blocking practices also expands the consumers' product choice set (apart from offering arbitrage opportunities on products consumers can already purchase at home).

3 Modelling Demand in Segmented versus Integrated Markets

Before describing the demand and oligopoly model in detail in the next sections, it is instructive to start with a general outline on how to model consumer demand in segmented and integrated markets. We first show how to implement this through the workhorse differentiated products demand model of BLP. Next, we discuss the critical shortcomings with this approach in our setting. This motivates our adapted BLP demand model, and provides a roadmap of the various parts of our analysis in the next sections.

A consumer i located in a country $c \in C$ faces the choice to buy a certain product j at a distribution channel $k \in \{T, O\}$ in a (possibly different) country $d \in C$. The channel $k = T$ refers to the traditional (or offline) channel; the channel $k = O$ refers to the online channel. The set of available products in country c at distributional channel k is $\mathcal{J}_{c,k}$. A consumer may also choose the outside good, which we define as product $j = 0$ at the offline channel T in the consumer's own country c.

Under segmented markets, a consumer located in country c can buy products only in her own country c, so her choice set consists of $\mathcal{J}_{c,T} \cup \mathcal{J}_{c,O}$. With integrated markets, a consumer in country c can engage in cross-border trade, and hence sees her choice set enlarged to include those of other countries. For example, if integration implies that products become available across all countries through both distribution channels, then the (common) choice set to all consumers becomes $(\cup_{d \in C} \mathcal{J}_{d,T}) \cup (\cup_{d \in C} \mathcal{J}_{d,O})$. If integration implies that products become available across all countries only through the online distribution channel (as with a ban on geo-blocking), then the choice set of a consumer located in country c becomes $\mathcal{J}_{c,T} \cup (\cup_{d \in C} \mathcal{J}_{d,O})$.

The conditional indirect utility of consumer i located in country c for a product j purchased at channel k in country d is:

$$u_{ic,jkd} = \underbrace{x_j \beta_i + \gamma_i \times \mathbf{1}(k = O) - \alpha_c p_{jkd} + \xi_{c,jk}}_{V_{ic,jkd}} + \varepsilon_{ic,jkd}. \tag{1}$$

The vector x_j consists of product characteristics (identical across channels and countries), $\mathbf{1}(k = O)$ is an indicator equal to one for the online channel, p_{jkd} is the price of product j at channel k in country of purchase d, $\xi_{c,jk}$ is the unobserved quality of product j at channel k, as perceived by consumers located in country c. The parameters β_i and γ_i are random taste parameters for the valuation of the product characteristics and the online distribution channel, and α_c is a country-specific price parameter. Finally, the random term $\varepsilon_{ic,jkd}$ is an individual-specific valuation of consumer i located in country c for product j purchased at distribution channel k in country d, i.i.d. distributed according to a type 1 extreme value distribution. We will sometimes write utility excluding the individual-specific valuation term as $V_{ic,jkd} \equiv u_{ic,jkd} - \varepsilon_{ic,jkd}$. We normalize this term to zero for the outside good, $V_{ic,0Tc} = 0$.

Under geographically segmented markets, a consumer i can buy only in her own country c, and not in any other country $d \neq c$. Assuming random utility maximization and integrating over the random taste parameters β_i and γ_i, we can write the market share for product j at channel k in country c as:

$$s_{c,jk} = s_{c,jkc} = \int \frac{\exp\left(V_{c,jkc}(\beta,\gamma)\right)}{1 + \sum_{j' \in \mathcal{J}_{kc}} \sum_{k' \in \{T,O\}} \exp\left(V_{c,j'k'c}(\beta,\gamma)\right)} dF_{\beta\gamma}\left(\beta,\gamma\right). \tag{2}$$

The first equality highlights that a product's market share in country c is just equal to the market share of consumers located in that country, because consumers cannot buy in any other country $d \neq c$, i.e. $s_{c,jkd} = 0$ for $d \neq c$. The second equality is the usual BLP expression for market shares in segmented markets, which averages the logit choice probabilities over unobserved consumer types (β_i, γ_i). Total demand by all L_c consumers located in country c for product j at channel k is $q_{c,jk} = s_{c,jk} L_c$.

With integrated markets, the choice set of a consumer located in country c includes all countries. The market share from consumers located in country c for product j at channel k in country d is then equal to:

$$s_{c,jkd} = \int \frac{\exp\left(V_{c,jkd}(\beta,\gamma)\right)}{1 + \sum_{d' \in C} \sum_{j' \in \mathcal{J}_{k'd'}} \sum_{k' \in \{T,O\}} \exp\left(V_{c,j'k'd'}(\beta,\gamma)\right)} dF_{\beta\gamma}\left(\beta,\gamma\right). \tag{3}$$

Total demand by all consumers L_c located in country c for product j at channel k is $q_{c,jk} = \sum_d s_{c,jkd} L_c$.

Our general goal is to estimate a demand model under segmented markets (as in (2)), add an oligopoly model of price-setting behavior, and perform counterfactuals on how equilibrium changes under integrated markets (as in (3)). However, applying the standard BLP demand models, (2) and (3), is unsatisfying in our context because it is based on the very high dimensional i.i.d. individual taste term $\varepsilon_{ic,jkd}$. This term is not only specific to each product j, but also to each distribution channel k and country of purchase d. This may imply implausible substitution patterns and misleading welfare implications when studying the impact of new goods that become available in other countries or distribution channels (as these would artificially increase the product space). Berry and Pakes (2007) develop an approach to eliminate the individual taste parameter and estimate a "pure characteristics" model by entirely eliminating the term $\varepsilon_{ic,jkd}$. However, their approach involves a considerable increase in computational complexity, and most applications use the standard BLP model (while being cautious to specify a sufficiently rich model to capture heterogeneity in the valuations of the product characteristics).

We instead propose an adapted BLP demand model, and specify utility as

$$u_{ic,jkd} = x_j\beta_i + \gamma_i \times \mathbf{1}(k = O) - \alpha_c p_{jkd} + \xi_{c,jk} + \varepsilon_{ic,j}. \tag{4}$$

This reduces the dimensionality of the individual taste parameter from $\varepsilon_{ic,jkd}$ to $\varepsilon_{ic,j}$: this is still specific to the product j, but no longer to the distribution channel k and country of purchase d. Hence, substitution patterns and welfare gains from increased product availability are not affected by artificial tastes for products at certain channels or countries of purchase. For our empirical demand analysis (section 4), only the elimination of the channel dimension is relevant, because we estimate the model under the assumption of segmented markets. For our counterfactual analysis (section 5), the elimination of the country of purchase dimension also becomes highly relevant.

4 Demand and Oligopoly in Segmented Markets

In this section, we analyze demand and oligopoly under segmented markets, i.e. when firms could use geo-blocking to prevent consumers from shopping online in other countries.

4.1 Adapted BLP Demand Model

Under segmented markets, a consumer located in country c can only purchase in her own country c and not in any other country $d \neq c$. To simplify notation, we suppress the subscripts c in this subsection. The utility of consumer i for product j at channel k, (4), can be simplified to

$$u_{i,jk} = \underbrace{x_j\beta_i + \gamma_i \times \mathbf{1}(k = O) - \alpha p_{jk} + \xi_{jk}}_{V_{i,jk}} + \varepsilon_{i,j}. \tag{5}$$

Specify the online taste parameter as $\gamma_i = \gamma^O + \sigma^O \nu_i^O$, where γ^O is the mean valuation for shopping online (possibly negative), σ^O is the standard deviation, and $\nu_i^O \sim N(0,1)$ is a standard normal random variable. At this point, we do not yet specify the taste parameter for the product characteristics β_i. The key feature of this adapted BLP demand specification is that the individual taste parameter $\varepsilon_{i,j}$ is specific only to the product j, while in the standard BLP model it is specific to every alternative, i.e. every product j at every channel k (with a term $\varepsilon_{i,jk}$).

A consumer can conceptually break her choice problem down in two parts: determine the preferred sales channel for each product j, and then compare the preferred sales channel of every product across all possible products. The first part is simple: a consumer prefers the traditional channel T of product j if

$$u_{i,jT} \geq u_{i,jO},$$

Table 2: Consideration Sets for the Adapted BLP Model

online valuation	$\mathcal{J}_T^j$	$\mathcal{J}_O^j$	$D_{i,j}$
$\nu_i^O \in [-\infty, \Delta_1)$	$\{1, 2, ..., J\}$	$\varnothing$	$\sum_{j' \in \mathcal{J}_T} \exp\left(V_{i,j'T}\right)$
$\nu_i^O \in [\Delta_1, \Delta_2)$	$\{2, ..., J\}$	$\{1\}$	$\sum_{j' \in \mathcal{J}_T^2} \exp\left(V_{i,j'T}\right) + \exp\left(V_{i,1O}\right)$
$\vdots$	$\vdots$	$\vdots$	
$\nu_i^O \in [\Delta_{j-1}, \Delta_j)$	$\{j, ..., J\}$	$\{1, 2, ..., j-1\}$	$\sum_{j' \in \mathcal{J}_T^j} \exp\left(V_{i,j'T}\right) + \sum_{j' \in \mathcal{J}_O^j} \exp\left(V_{i,j'O}\right)$
$\vdots$			
$\nu_i^O \in [\Delta_J, \infty)$	$\varnothing$	$\{1, 2, ..., J\}$	$\sum_{j' \in \mathcal{J}_O^J} \exp\left(V_{i,j'O}\right)$

or equivalently if

$$\nu_i^O \leq \frac{-\alpha\left(p_{jT} - p_{jO}\right) + \xi_{jT} - \xi_{jO} - \gamma^O}{\sigma^O} \equiv \Delta_j$$

(after substituting (5) and making use of $\gamma_i = \gamma^O + \sigma^O \nu_i^O$). Hence, a consumer prefers the traditional channel of product j if and only if her valuation for the online channel is sufficiently low, $\nu_i^O \leq \Delta_j$. Note that the cut-off value Δ_j depends only on the price and unobserved quality difference, and not on the product characteristics, as these are the same on both channels.

The second part of the consumer's choice problem compares the preferred channel of each product across all products. To address this, let us first define the consideration sets of a consumer, depending on her online valuation ν_i^O. Suppose (without loss of generality) that the product cut-off values can be ranked as follows $\Delta_1 \leq ...\Delta_{j-1} \leq \Delta_j \leq \Delta_{j+1} \leq ... \leq \Delta_J$, i.e. product 1 is the least attractive at the traditional channel, whereas product J is the most attractive at the traditional channel. Given this ordering, define the sets $\mathcal{J}_T^j \subseteq \mathcal{J}_T = \{j, ..., J\}$ and $\mathcal{J}_O^j \subseteq \mathcal{J}_O = \{1, ..., j-1\}$. Table 2 uses this notation to show the considerations sets of a consumer for different realizations of her online valuation ν_i^O. For example, if $\nu_i^O \leq \Delta_J$, a consumer only considers products at the traditional sales channel. If $\nu_i^O \in [\Delta_1, \Delta_2)$, she compares product 1 of the online channel with the other products $j = 2, ..., J$ at the traditional channel. If $\nu_i^O \in [\Delta_{j-1}, \Delta_j)$, she compares products $1, 2, ..., j-1$ at the online channel with the remaining products $j, ..., J$ at the traditional channel. Finally, for $\nu_i^O \in [\Delta_J, \infty)$, she compares only products at the online sales channel.

Given these consideration sets, we obtain the following probabilities that a consumer would choose product j at the traditional channel T or online channel O:

$$s_{jT}(\beta_i) = \int_{-\infty}^{\Delta_1} \frac{\exp\left(V_{i,jT}\right)}{1 + D_{i,1}} d\Phi(\nu^O) + \int_{\Delta_1}^{\Delta_2} \frac{\exp\left(V_{i,jT}\right)}{1 + D_{i,2}} d\Phi(\nu^O) + ... + \int_{\Delta_{j-1}}^{\Delta_j} \frac{\exp\left(V_{i,jT}\right)}{1 + D_{i,j}} d\Phi(\nu^O), \quad (6)$$

and

$$s_{jO}(\beta_i) = \int_{\Delta_j}^{\Delta_{j+1}} \frac{\exp\left(V_{i,jO}\right)}{1+D_{i,j+1}} d\Phi(\nu^O) + \int_{\Delta_{j+1}}^{\Delta_{j+2}} \frac{\exp\left(V_{i,jO}\right)}{1+D_{i,j+2}} d\Phi(\nu^O) + ... + \int_{\Delta_J}^{\infty} \frac{\exp\left(V_{i,jO}\right)}{1+D_{i,J+1}} d\Phi(\nu^O), \quad (7)$$

where $\Phi(\nu^O)$ denotes the standard normal distribution, $V_{i,jk} = V_{jk}\left(\beta_i, \nu^O\right)$, and the terms $D_{i,j} = D_j(\beta_i, \nu^O)$ are defined in the final column of Table 2.

To interpret this, consider the expression for $s_{jT}(\beta_i)$. The first term integrates consumers with a very low online valuation ($\nu_i^O \leq \Delta_1$), whose consideration set consists of the traditional channel for every product. The second term integrates over consumers with a higher online valuation ($\nu_i^O \in [\Delta_1, \Delta_2)$), who compare the online channel for product 1 with the traditional channel for all other products. The final term of $s_{jT}(\beta_i)$ integrates over consumers with the highest online valuations for whom the traditional channel may still be chosen ($\nu_i^O \in [\Delta_{j-1}, \Delta_j)$): these consumers compare the online channel for products $1, 2, ..., j-1$ with the traditional channel for products $j, ...J$.

The aggregate market share of product j at channel k is obtained by integrating (6) and (7) over β_i, so $s_{jk} = \int s_{jk}(\beta) dF_\beta\left(\beta\right)$. This adapted BLP model is appealing because of its substitution patterns between the traditional and online channel. To illustrate this, consider the cross-price effect of p_{jO} on s_{jT} (conditional on β_i). It can be verified that this is given by

$$\frac{\partial s_{jT}(\beta_i)}{\partial p_{jO}} = \frac{\alpha}{\sigma^O} \frac{\exp\left(V_{jT}(\beta_i, \Delta_j)\right)}{1+D_j(\beta_i, \Delta_j)}. \quad (8)$$

Intuitively, substitution from the online to the traditional channel of product j stems from the mass of consumers who were close to indifferent between both channels of product j. Substitution between both channels will be strong when there is limited consumer heterogeneity in the valuation for the online channel (low σ^O). In contrast, in a traditional BLP model, the cross-price effect of p_{jO} on s_{jT} is given by

$$\frac{\partial s_{jT}(\beta_i)}{\partial p_{jO}} = \alpha \int_{-\infty}^{\infty} \frac{\exp\left(V_{i,jT}(\beta_i, \nu^O)\right)}{1+D(\nu^O, \beta_i)} \frac{\exp\left(V_{i,jO}(\beta_i, \nu^O)\right)}{1+D(\nu^O, \beta_i)} dF_{\nu^O}(\nu^O), \quad (9)$$

where

$$D(\beta_i, \nu^O) \equiv \sum_{j' \in \mathcal{J}_k} \sum_{k' \in \{T,O\}} \exp\left(V_{j'k'}\left(\beta_i, \nu^O\right)\right).$$

This is the usual cross-price effect from a standard BLP model, which averages the substitution (conditional on β_i) over all online valuation types. Heterogeneity in the valuation for the online channel still plays a role, but it is mixed up with heterogeneity in the tastes for the product/channel alternatives $\varepsilon_{i,jk}$. Hence, even if there would be very limited heterogeneity in the valuation for online, there may still be weak substitution between both channels in the standard BLP model.

4.2 Specification, Estimation and Instruments

We first discuss the utility specification for the adapted BLP model. Next, we discuss how we calculate the market shares and invert the market share system to solve for the error term. Finally, we discuss the instruments used to estimate the model.

Specification We estimate the demand model based on panel data for products sold through both distribution channels across multiple countries and time periods. We reintroduce the subscript for country c (our 10 European countries), but suppress a subscript for the period t (months during December 2012-March 2015).

Similar to the random coefficient for the online channel γ_i, we specify the random coefficients for the product characteristics n to be normally distributed, i.e. $\beta_i^n = \beta^n + \sigma^n \nu_i^n$ where $\nu_i^n \sim N(0,1)$. Furthermore, we decompose the unobserved product quality as perceived by a consumer in country c for product j at channel k into three parts $\xi_{c,jk} = \xi_j + \xi_{c,k} + \widetilde{\xi}_{c,jk}$. The utility of a consumer i located in country c for each alternative can then be written as:

$$u_{ic,jk} = \delta_{c,jk} + \mu_{i,jk} + \varepsilon_{ic,j},$$

where the mean utility part $\delta_{c,jk}$ is

$$\delta_{c,jk} = x_j \beta - \alpha_c p_{jkc} + \xi_j + \xi_{c,k} + \widetilde{\xi}_{c,jk}. \tag{10}$$

and the deviation from this mean is

$$\mu_{i,jk} = \sum\nolimits_{n=1}^{N} \sigma^n \nu_i^n x_j^n + \sigma^O \nu_i^O \times \mathbf{1}(k = O).$$

The product characteristics in the vector x_j include CPU speed, the amount of RAM, weight, the display diagonal and the display resolution. We let the mean price coefficient α_c vary across countries according to the earlier documented three fairly homogenous income groups. We do not allow for heterogeneity in price sensitivity within these groups. As in Petrin (2002), we find that once we control for the different means between income levels, the estimated standard deviation of a random coefficient for price is statistically insignificantly different from zero. We allow for random coefficients for shopping online (through the parameter σ^O) and for two product characteristics (σ^n): the amount of RAM and the display resolution (pixels per inch).

We include a full set of product fixed effects ξ_j, which reflects systematic unobserved product quality common across countries (and time periods). We also include country and channel fixed effects ξ_{ck}, reflecting unobserved valuations for portable PCs that are specific to each country and distribution channel. This flexibility thus also accounts for differences in the popularity of online shopping across countries. We also include month-of-year fixed effects, a general time trend (to account for a gradual substitution out of portable PCs over time), and country-specific trends for the online channel (to account for different evolutions in the popularity of online shopping across countries). Any remaining unobserved quality is captured by the error term $\widetilde{\xi}_{cjk}$.

Estimation Estimating the adapted BLP model requires similar broad steps as the standard BLP model, but the implementation is different. We start from the non-linear market share system

$$s_{c,jk} = \int s_{c,jk}(\beta) dF_\beta(\beta), \tag{11}$$

where $s_{c,jk}(\beta_i)$ is given by (6) and (7) after including a country subscript c. We approximate the integral by simulating over the standard normal random variables ν_i. We then invert the market share system (market by market) to obtain a solution for the mean utilities $\delta_{c,jk}(\mathbf{s}_c, \sigma)$, where $\mathbf{s}_c$ is the market share vector in country c and σ is a vector of the standard deviations of the random coefficients. Using (10), this gives

$$\widetilde{\xi}_{c,jk} = \delta_{c,jk}(\mathbf{s}_c, \sigma) - (x_j\beta - \alpha_c p_{jkc} + \xi_j + \xi_{c,k}). \tag{12}$$

In the standard BLP model, this inversion exists and BLP suggest a contracting mapping to obtain the solution. In our adapted BLP model, several complications arise. First, a solution does not necessarily exist and other methods than BLP's contraction mapping are required (see Berry and Pakes, 2007). Second, the market share integral is complicated by the fact that the consideration sets may change depending on the parameter values, implying discontinuities in the market share function.

As a solution to these problems, we approximate our adapted BLP model through a random coefficients nested logit model, where each product j is a nest containing two alternatives: the traditional and the online channel. The individual-specific taste parameter in such a set up is $\varepsilon_{ic,j} + (1-\rho)\varepsilon_{ic,jk}$ (Berry, 1994), where ρ is a nesting parameter measuring the extent of preference correlation for the two channels within the product nest. If we set $\rho = 0$, we would obtain the standard BLP model. Conversely, as $\rho \to 1$, we obtain the adapted BLP model; see the Appendix for details. Hence, we can approximate the adapted BLP model through a random coefficients nested logit model by setting ρ sufficiently close to 1. The advantage of this approach is that we can rely on properties from the random coefficients nested logit model, for which a contraction mapping solution exists (Grigolon and Verboven, 2014) and which is smooth in the parameters. As $\rho \to 1$, however, the contraction mapping becomes a weak contraction and thereby considerably increases the computational burden of the estimation. To tackle this issue, we use the globally convergent fixed point acceleration approach of Zhang, O'Donoghue and Boyd (2018). On average, it reduces the number of iterations required for the inversion of market shares by a factor of five and the computational runtime of the inversion by a factor of three. For further details, we refer to

Figure 4: Share Deviations - Adapted BLP versus Random Coefficients Nested Logit with $\rho = 0.9$

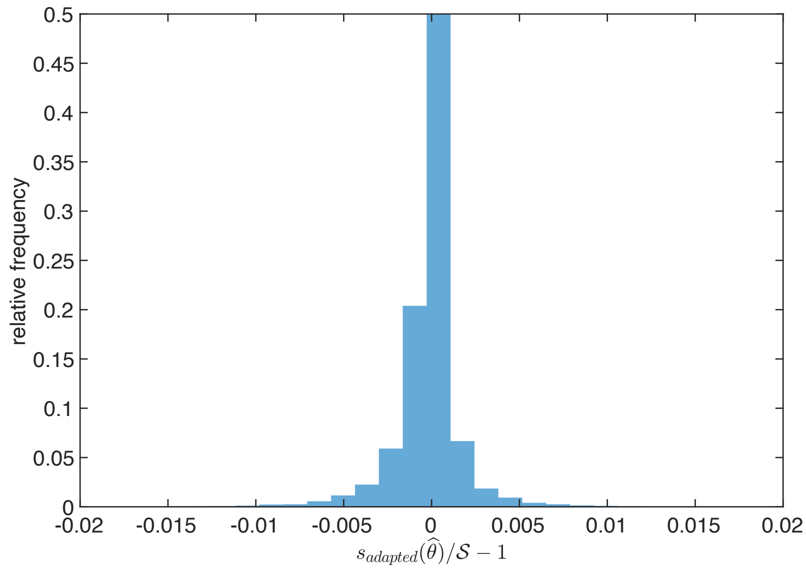

Note: s denotes the aggregate shares obtained from (6) and (7) at the estimated random coefficients nested logit parameter vector, $\widehat{\theta}$. S denotes the observed aggregate share vector, which the estimated approximate adapted BLP model matches very closely.

section A.2.2 in the Appendix and the paper by Zhang, O'Donoghue and Boyd (2018). Note that our approach may be viewed as a "light version" of the scaling approach of Berry and Pakes (2007) to approximate the pure characteristics model: in our notation, they use $(1 - \rho)\, \varepsilon_{ic,jk}$ (without the $\varepsilon_{ic,j}$ term) and let $\rho \to 1$. Unlike Song (2015) it enables the estimation of multiple random coefficients in a broad variety of settings.

A practical question is how close ρ should be to 1 to have a reasonable approximation of the adapted BLP model. Picking large values makes the approximation more accurate, but may also lead to numerical difficulties and slow down the contraction mapping. Specifically, we set $\rho = 0.9$. For higher values of ρ we experienced numerical difficulties, because we obtain numbers that exceed the limits of double precision floating point arithmetic. To assess how well we approximate the adapted BLP model with $\rho = 0.9$, we evaluate the aggregate market share function of our adapted BLP model (11) at the estimated parameter vector of the nested logit random coefficients model with $\rho = 0.9$. Figure 4 shows the distribution of the resulting net relative deviations between the two aggregate market share vectors. The deviations are most often very small. Almost all observations have a relative deviation of less than one percent in absolute value, and most often the deviations are much smaller. We therefore conclude that setting $\rho = 0.9$ is a sufficiently accurate approximation of the adapted BLP model for our purposes.

Instruments A final step consists in constructing instrumental variables that satisfy the orthogonality conditions $E\left(\widetilde{\xi}_{c,jk}|z_{c,jk}\right)=0$, so that they can be interacted with the model's predicted error (12) in a GMM estimator. We need a sufficient number of instruments to estimate both the mean valuations of the product characteristics (β), their standard deviations σ and the price coefficients α_c. We follow BLP and consider that the product characteristics other than price are exogenous, so that functions of the own and rival product characteristics can be used as instruments. There are two concerns with these characteristics-based instruments. First, similar to other markets for durable consumption goods, the market for portable PCs is characterized by product attributes that are improving over time. This may violate the assumption that the directly observable characteristics are fixed. Second, Armstrong (2016) shows that characteristics-based instruments can lose their identifying power when the number of products becomes large. In our entire sample there are between 24 and 85 unique portable PC models available in each market.

With regard to the first concern, our data is observed at a monthly frequency, which makes the assumption that characteristics can be treated as fixed for each individual market more reasonable. Apple, for example, updates its MacBook Pro on average every 301 days.[8] To alleviate these concerns further, we exclude CPU speed and the amount of RAM from the attributes that we use to compute our characteristics-based instruments. These two components of a laptop's design can be adjusted more easily and quickly than its weight, display diagonal and the display's resolution. The latter three attributes determine to a large extent the laptop's form factor and thereby also its overall design. Again, taking the example of Apple, the overall design of a laptop sees much fewer substantial changes over a period of several years than its internal components, such as the CPU, the amount of RAM or the size and type of the hard drive.

Second, with the results of Armstrong (2016) in mind and in the spirit of Gandhi and Houde (2019), we avoid summing over all available rival products in a market to compute our instruments. Instead, we partition the observed characteristics space to delineate groups of products that consumers are likely to perceive as relevant substitutes. Depending on where each laptop is located in this partition, we compute our instruments for this laptop by summing over the characteristics of rival products located in the same bin of characteristics space. As is standard, when forming these sums, we distinguish between observations that are sold by the same firm and observations that are sold by rival firms. Specifically, for each of the three remaining attributes in x_j (weight, diagonal and display resolution), we partition the marginal distribution into two segments: observations above and below the median. The partition of the characteristics space is then based on $2^3=8$ possible bins of each product's possible location in this characteristics space grid. For example, laptop j offering a less than median weight, display diagonal and display resolution has an address of $(0,0,0)$ in this space. We compute characteristics sums within and between firms, to obtain a total of six excluded instruments. We then interact these instruments with the three country group

[8]See the Buyer's Guide on `https://buyersguide.macrumors.com/#Mac`.

dummy variables (since we have a separate price coefficient per country group), so that we have eighteen excluded instruments in total.[9]

Table (A.1) in the Appendix presents the results of the first-stage regression of price on the instruments. The crucial outcome is the F-statistic, which we compute only for the set of excluded instruments. Thus, our null is that the excluded instruments jointly have no explanatory power for prices. Our concern in this context is that as we add fixed effects, an increasing fraction of the excluded instruments' variation is explained by these dummy regressors. This can potentially yield weak instruments and thereby invalidate our identification approach. Column (IV) of Table (A.1) shows, however, that this is not the case. With an F-statistic that exceeds 22, we can comfortably reject the null that the characteristics space instruments are not relevant.

4.3 Empirical Results

We first discuss the demand parameter estimates and then the resulting price elasticities.

Parameter Estimates We compare the demand parameter estimates for four models: logit, adapted logit, BLP and adapted BLP. Table (3) shows the main parameter estimates for these four demand demand models. Figure (5) plots the estimated country-specific online valuations for our preferred model, the adapted BLP model (with complete results on the mean online valuations for all four models in Table A.2 in the Appendix). Broadly speaking the common parameters (i.e. the mean valuations) are broadly comparable across models, though the substitution patterns may differ drastically as we discuss further below. The BLP and adapted BLP models contain as additional parameters the standard deviations of three characteristics: the online channel, RAM and display resolution.[10] We explicitly test the BLP and adapted BLP models against their logit counterparts, and report the associated Wald statistics. The null is that all estimated standard deviations of the random coefficients are statistically insignificantly different from zero, $H_0 : (\widehat{\sigma}_{on}, \widehat{\sigma}_{ram}, \widehat{\sigma}_{ppi}) = \mathbf{0}$. With a value of roughly 50, the Wald statistic exceeds the 99 percent confidence level critical threshold of about 11 and we can therefore reject the corresponding logit models at any reasonable confidence level.

The price coefficients have the expected sign and are precisely estimated for all models. Moreover, all three specifications deliver the intuitively appealing result that the consumers' price sensitivity is highest in the low income country group, and lowest in the high income country group.[11]

[9] We considered implementing efficient or "optimal" instruments, as discussed in Reynaert and Verboven (2014) and Conlon and Gortmaker (2020). However, these are more tedious to compute in our approximation to the adapted BLP model, and since we obtain relatively precise estimates for the random coefficients we did not pursue this further.

[10] Extended specifications with additional random coefficients including for price do not yield additional significant standard deviation estimates.

[11] We also estimated the logit model using ordinary least squares, i.e. without instruments. As expected, this results in a substantial underestimation of the price coefficients. The OLS estimates are between 4 and 9 times smaller than the estimates when we instrument for price. Moreover, without instruments almost 84 percent of all observations are

Table 3: Demand Estimates - Price and Characteristics

	Logit	Adapted Logit	BLP		Adapted BLP	
			mean	std. dev.	mean	std. dev.
α_L	**.0061**	**.007**	**.0068**		**.0069**	
	(.0008)	(.0008)	(.0011)		(.0011)	
α_M	**.0054**	**.0063**	**.0058**		**.0058**	
	(.0007)	(.0007)	(.0011)		(.0010)	
α_H	**.0039**	**.0047**	**.0043**		**.0043**	
	(.0007)	(.0007)	(.0009)		(.0010)	
Online				**7.562**		**8.959**
				(2.541)		(2.384)
CPU Speed	**.5455**	**.6027**	**.8222**		**.8769**	
	(.1231)	(.1273)	(.1922)		(.2021)	
RAM	**.0425**	**.0523**	**-.2091**	**.7568**	**-.2091**	**.7639**
	(.0089)	(.0092)	(.1265)	(.2024)	(.1053)	(.2094)
Weight	**-.0962**	**-.1392**	**-.2569**		**-.3012**	
	(.1456)	(.1506)	(.1710)		(.1806)	
Diagonal	**.1152**	**.1312**	**.1518**		**.1591**	
	(.0228)	(.0236)	(.0314)		(.0324)	
Resolution	**.9356**	**1.117**	**1.262**	**.0011**	**1.290**	**.0045**
	(.2553)	(.2640)	(0.8266)	(1.766)	(.4438)	(1.652)
Constant	**-7.505**	**-7.819**	**-8.247**		**-8.523**	
	(.8350)	(.8633)	(1.369)		(1.137)	
Trend	**-.0877**	**-.0787**	**-.0839**		**-.0921**	
	(.0071)	(.0074)	(.0123)		(.0150)	
Wald Stat.	-	-	41.60		49.99	
Crit. Value			11.34		11.34	
$\overline{\eta}_{jj}$	-3.46	-3.87	-3.75		-3.78	
$\# \ \eta_{jj} > -1$	0	0	0		0	
$\overline{\eta}_{j,\text{trad,on}}$	.0114	9.739	.0041		.0954	
$\overline{\eta}_{j,\text{on,trad}}$	.0410	23.46	.0147		.2712	

Note: Based on 10 288 observations. Standard errors are shown in parentheses. The critical value for the Wald statistic applies to a 99 percent confidence interval and three degrees of freedom. 1 000 modified latin hypercube sampling (MLHS) draws and 30 different starting values for the nonlinearly entering coefficients were used during the estimation. The price coefficients vary between three country groups that are color coded in Figure (A.1). $\overline{\eta}_{j,trad,on}$ and $\overline{\eta}_{j,on,trad}$ are the average product-level cross-price elasticity. Each value is based on the average of 3 787 demand derivatives of all laptops that are available in both the traditional and online distribution channels.

We focus our discussion of the parameters for the product characteristics on the adapted BLP model, but also relate it to the other models where notable differences occur. As expected, consumers have a higher mean valuation for machines with a faster CPU, a larger display size (diagonal) and a larger display resolution. Furthermore, consumers have a lower though not precisely estimated mean valuation for portables with a higher weight. The mean valuation for RAM is negative and statistically significant, but the standard deviation for the RAM valuation is large and significant, showing there is a lot of unobserved consumer heterogeneity for this attribute. Note that in

estimated to be price inelastic, while with instruments all observations are price elastic.

Figure 5: Estimated Online Means and Trends - Adapted BLP Model

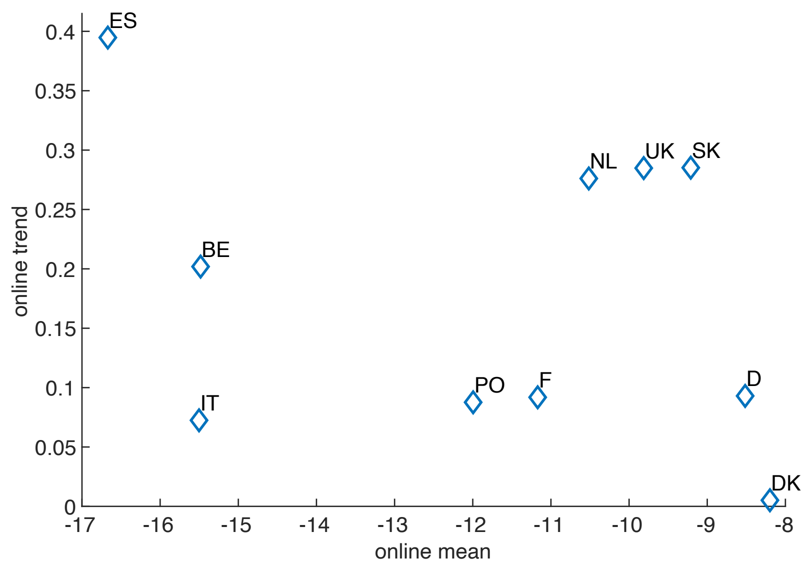

Note: Based on 10 228 sample observations and the adapted BLP demand estimates reported in Tables (3) and (A.2).

the logit specifications without the standard deviation for RAM, the mean valuation is positive and precisely estimated but quantitatively small. We also include a trend to capture general changes in the demand for portable PCs, and estimate this trend to be negative and highly significant. This corresponds to falling sales for portable PCs during our sample period, while demand for smartphones and tablets has been growing.

Finally consider the consumers' valuations for the online distribution channel. We summarize the country-specific mean valuations, and trends in these mean valuations in Figure (5) for the adapted BLP model (and show the complete results for all four models in Table A.2). Figure (5) shows that there are considerable differences in the mean valuations for the online channel across countries, consistent with the patterns reported in Table 1. The countries where we observe low online sales also tend to be the countries with the lowest mean valuations, and vice versa. Similarly, countries that show increasing online sales tend to be the countries with more positive trends. Figure (5) also shows some negative correlation between the country intercepts and trends, indicating that the late coming countries are catching up.[12]

In both BLP models, the standard deviation of the online coefficients (σ^O) is precisely estimated and large when compared with the country-specific mean valuations from Figure (5). This shows the presence of substantial within-country heterogeneity in the valuation of the online distribution

[12]These findings do not hold for the logit models without a random coefficient for the online channel, where we obtain counterintuitive estimates for the country-specific mean valuations and trends inconsistent with the patterns reported in Table 1.

channel. While the magnitude of σ^O is comparable for both BLP models, this does not need to be the case in general, and we will see that it translates in drastically different cross-price elasticities between the two channels.

In sum, our adapted BLP model yields intuitive results consistent with the preliminary evidence reported in section 2. The price sensitivity differs across countries according to their income levels. Consumers show a significant valuation for several product characteristics. There is also important consumer heterogeneity, in particular regarding the valuation of the online distribution channel. Part of this heterogeneity refers to cross-country differences that line up with our earlier evidence on online market shares. But there is also significant heterogeneity within a country.

Based on these estimates, we now discuss the implied price elasticities and firms' markups by adding an oligopoly model of price-setting behavior.

Price Elasticities Table 3 also shows the average own-price elasticities of demand over all observations for the four demand models. These averages are fairly similar, varying between -3.46 and -3.87. The left panel of Figure 6 shows the entire distribution of the own-price elasticities, conditional on the three country income groups. Price elasticities are on average higher in the low income group, and lower for the medium and especially the high income group. This is reflected in a shift to the right of the elasticity distribution for the medium and high income groups.

While the own-price elasticities in Table 3 are similar across the four demand models, this is not the case for the cross-price elasticities. To illustrate the most interesting differences, we calculate the cross-price elasticities between the traditional and online variant of every product, i.e. $\eta_{jO,jT}$ and $\eta_{jT,jO}$ for every product j, and we average this over products. The underlying demand derivatives are given by (8) and (9) for the adapted BLP and standard BLP, respectively. The bottom rows of Table 3 show the results from this comparison. The logit model generates low cross-price elasticities between both channels, which is due to the artificial individual taste shock included for every product and channel combination ($\varepsilon_{i,jk}$). In contrast, the adapted logit model generates extremely high cross-price elasticities between both channels. This is because the individual taste shock is now specific only to the product and no longer to the distribution channel ($\varepsilon_{i,j}$), and there is no other heterogeneity in the valuation of the online channel. As a result, the two channels of the same product are by construction perfect substitutes, implying very high cross-price elasticities.

The BLP model yields even lower cross-price elasticities between both distribution channels than the logit model. But just as in the logit model this is difficult to interpret: it stems from a mixture of heterogeneity in the valuation of the online channel (σ^O) and the individual taste shock for every product and channel combination ($\varepsilon_{i,jk}$). In contrast, in the adapted BLP model substitution is not driven by the individual taste shock (as it specifies $\varepsilon_{i,j}$), but only by the estimated heterogeneity in the valuation for the online channel (σ^O) (and indirectly also by heterogeneity in the valuation of

Figure 6: Distribution of Own-Price Elasticities and Markups Implied by Adapted BLP Model

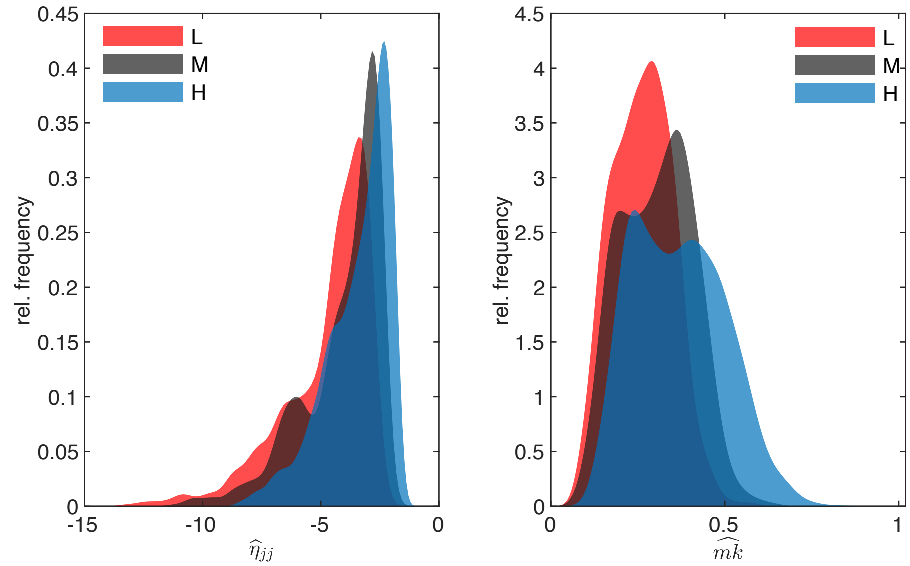

Note: Based on 10 228 sample observations and the adapted BLP demand estimates reported in Tables (3) and (A.2). The average markups in the low-, medium- and high-income country groups are .263, .305 and .368, respectively.

other product characterstics). As a result, we find that the cross-channel substitution is considerably higher in the adapted BLP model than in the standard BLP model, by a factor of at least 10. It is important to stress that this is mainly driven by our empirical estimate of the standard deviation for the valuation of the online channel (σ^O). If we would have obtained a lower estimate of σ^O, then the cross-elasticities would have been even higher, and vice versa (as can be seen from the cross-demand derivative (8)).

4.4 Oligopoly and Markups

To compute markups and back out marginal costs, we use a standard oligopoly model of multi-product price-setting firms, similar to for example Sovinsky Goeree (2008), Eizenberg (2014) and Song (2015) for the PC industry. This approach can be justified under a competitive retail sector, or more generally under an imperfectly competitive retail sector with efficient contracting between producers and retailers (no double marginalization effects). For example, Rey and Vergé (2010) have a model of "interlocking relationships" between producers and retailers that yields this outcome. As such, the markups should be interpreted as the combined markups of the producers and retailers. Similarly, the marginal costs can be interpreted as the sum of the producers' marginal costs (all Asian manufacturers) and the retailers' local distribution costs, which may vary across countries. We do not attempt to explicitly model more complicated relationships between producers and retailers, because we observe only retail prices and not wholesale prices, and only the total sales per product, and not the sales broken down by retailer (as in e.g. Bonnet and Dubois (2010)).

More specifically, let mc_{jkc} be the (constant) marginal cost of product j at distribution channel k in country c and let $\mathcal{F}_{fk}$ be the set of products sold by firm f at channel k. The profits of firm f in country c are the sum of the profits over all its products sold through both distribution channels:

$$\pi_{cf} = \sum_{j' \in \mathcal{F}_{fk'}} \sum_{k' \in \{T,O\}} \left(p_{j'k'c} - mc_{j'k'c}\right) s_{c,j'k'} L_c. \tag{13}$$

The market share $s_{c,jk}$ in (13) is given by (11). Because markets are nationally segmented, this depends only on the prices of alternatives in the same country.

Firms are multi-product Bertrand price-setting firms. They choose the prices of their products to maximize profits, taking as given the prices of the other firms. For each market c, this gives a system of first-order conditions for the optimal prices of every product j and channel k. Let $q_{c,jk} = s_{c,jk}L_c$ and let $\mathbf{p}_c$, $\mathbf{q}_c$ and $\mathbf{mc}_c$ be vectors with elements $p_{c,jk}$, $q_{c,jk}$ and $mc_{c,jk}$. Furthermore, let $\mathbf{\Omega_c}$ be a matrix for country c with own- and cross-demand derivatives $\Omega_{jkc,j'k'c} = \partial s_{c,jk}/\partial p_{j'k'c}$, and define the ownership or holding matrix $\mathbf{H_c}$ with entries $H_{jkc,j'kc} = 1$ if $j, j' \in \mathcal{F}_{fk}$ and zero otherwise. We can then write the system of first-order conditions in matrix notation, to calculate the marginal cost vector in country c as the difference between the price and equilibrium markup:

$$\widehat{\mathbf{mc}_c} = \mathbf{p}_c + \left[\mathcal{H}_c \odot \Omega_c\right]^{-1} \mathbf{q}_c \tag{14}$$

The implied percentage markups for product j at channel k are defined as $(p_{jkc} - \widehat{mc}_{jkc})/p_{jkc}$.

The right panel of Figure (6) plots the markup distributions by median income groups. Country groups with higher median per-capita incomes have markups with distributions that are shifted to the right. The average markups in the low-, medium- and high-income country groups are 26 percent, 29 percent, and 35 percent, respectively. The average markup levels fall somewhere in the middle of other studies for the US PC market: lower markup estimates by Sovinsky Goeree (2008) and Eizenberg (2014) and higher estimates by Song (2015). The markups are also in line with accounting information, with annual gross margins of 37 percent and Ebitda margins of 23 percent in 2019 according to CSI market.

To have an idea on the extent to which local distribution costs differ across countries, we regress marginal cost on the vector of product attributes x_j and a set of fixed effects ω_{kc} to account for systematic differences in local costs between countries for both distribution channels:

$$\widehat{mc}_{jkc} = x_j\gamma + \omega_j + \omega_{kc} + \widetilde{\omega}_{jkc} \tag{15}$$

Table A.3 in the Appendix reports the results and shows that local costs indeed vary across countries, both for the traditional and for the online channel. As a comparison, the regression results for the BLP model, which are quantitatively very close to those for the adapted BLP model, are reported in Table A.4. Higher CPU speeds, more RAM, larger display diagonals and higher display resolutions have an increasing effect on marginal costs. A higher weight, on the other hand, has a negative impact on the margin. Belgium and Denmark are estimated to be the high-cost countries in both the traditional and online channels, while France, Germany and the UK are estimated to have the lowest marginal costs in both channels on average. For the counterfactuals that we discuss below, we constrain the product-level marginal costs to be equal between countries in the online distribution channel, because all laptops are actually produced in Asia. Our findings are robust to allowing for country-specific online channel marginal cost shifters, however.

5 The Impact of Reducing Cross-Border Trade Restrictions

Our goal is to assess the impact of removing cross-border trade restrictions in online markets, following the ban of geo-blocking practices. This event essentially alters the demand of consumers, because they can buy products online across all countries (possibly at the expense of extra shipping costs). This, in turn, leads to a new integrated market equilibrium, where firms take into account the impact of their pricing decisions on the sales in other countries. As a result, online prices may adjust and converge across countries. Furthermore, there may be indirect price effects on the traditional channel, depending on the extent of substitution between both distribution channels.

To assess these effects, we make use of the demand estimates and the backed out marginal costs from the pre-ban situation with nationally segmented markets, as analyzed in section 4. We first provide an overview on how we model online market integration. Next, we show more formally how to obtain the post-integration market equilibrium. Finally, we discuss the results from our counterfactuals.

5.1 Overview of the Approach

We first discuss how the demand system is altered by online markets becoming integrated, and then turn to the role of shipping costs.

Post-integration Demand With nationally segmented markets, consumers can buy products only in their own country. The market shares thus depend only on the utilities for the alternatives available in the consumers' own country. In (11), we use the simplified notation $s_{c,jk}$, but in section 3 we stress that we can write this more explicitly as $s_{c,jk} = s_{c,jkc}$, because $s_{c,jkd} = 0$ for $d \neq c$ (i.e. consumers located in c only buy alternatives in c and not in any other country $d \neq c$).

With integrated markets, consumers in country c face an increased choice set because they can also buy in other countries d, so the market shares $s_{c,jkd}$ will no longer necessarily be zero for $d \neq c$.

In section 3, we obtain the expression (3) for $s_{c,jkd}$ in a standard BLP demand model, where the taste parameter $\varepsilon_{ic,jkd}$ is specific to both the product j, the channel k and the country d. For an adequate analysis it is highly desirable to reduce the dimensionality of the taste parameter to $\varepsilon_{ic,j}$. Otherwise, we obtain misleading conclusions on substitution effects and especially on the welfare effects from newly available products. Our adapted BLP model (11) reduces the dimensionality of the individual taste parameter for segmented markets, by eliminating the channel dimension k in $\varepsilon_{ic,jkd}$. We apply a similar logic in our counterfactual analysis with integrated online markets, where we also eliminate the country dimension d in $\varepsilon_{ic,jkd}$. Intuitively, a consumer's choice problem can again be broken down in two steps. First, for each distribution channel and (in the case of online) each country-of purchase, she determines the preferred product j. Second, given her preferred product, she determines in which distribution channel k and (in the case of online) in which of the ten countries d to place her order. As the multiple online observations of her preferred product share the same idiosyncratic match value ($\varepsilon_{ic,j}$), she only retains the online observation of j offering the highest utility in her consideration set. In this way, consumers impose arbitrage in the online distribution channels between countries.

To understand the economic effects of opening up borders in online markets, we consider two scenarios.

> *Pre-integration availability (PIA):* We begin with the hypothetical case where consumers only obtain access to the products in other countries that were already available in their own country. We refer to this as the scenario with pre-integration availability (PIA). This scenario is helpful to understand the price convergence effect of market integration, because this reduces the possibility to engage in cross-country price discrimination.
>
> *Full availability (FA):* We then consider the case where consumers can also obtain access to other products abroad that were previously not available in their own country. We refer to this as the scenario with full availability (FA). This scenario is what we are ultimately interested in. It combines the effects from removing price discrimination and obtaining more choice. The first part may mainly involve distributional effects (between consumers in different countries, or between consumers and firms). The second part may potentially entail market expansion effects.

Formally, the difference between the two scenarios comes from how we construct the choice sets in the demand equation. Before integration, a consumer from country c has a choice set for online products $\mathcal{J}_{Oc}$. Under integration with FA, consumers have the choice sets $\mathcal{J}_{Od}$ for all countries $d \in C$. Under integration with PIA, a consumer in country c has more limited choice sets $\mathcal{J}_{c,Od} = \mathcal{J}_{Oc} \cap \mathcal{J}_{Od}$ for all countries $d \in C$ (where obviously $\mathcal{J}_{c,Oc} = \mathcal{J}_{Oc}$ is the local choice set already available before integration).

Incorporating Shipping Costs between Countries The European Commission distinguishes between justified and unjustified geoblocking practices. While access to online channels in other countries must not be blocked, it may be justified to charge foreign consumers additional fees for the shipping costs involved in serving them. To incorporate this, we also compute counterfactual equilibria that account for the presence of shipping costs between countries. Such counterfactuals may be interpreted as moves to partial integration as opposed to full integration in the absence of shipping costs.

We calculate the bilateral shipping costs between the ten countries of our analysis. Our measure is based on the parcel postage rates for weight categories between two and five kilograms, which we match to the portable PCs based on their weight. Table A.5 of the Appendix shows the values for these shipping costs.

5.2 Post-integration Equilibrium

Under segmented markets, consumers can only buy products in their own country, so that the market shares (11) depend only on the price vector in the consumers' own country. The profits of a firm f are the sum of (13) across all countries c, $\sum_c \pi_{cf}$, which we write here with separate components for the traditional and the online channel:

$$\begin{aligned}\pi_f \quad &= \sum_{c \in C} \sum_{j' \in \mathcal{F}_{fk'}} \left(p_{j'Tc} - mc_{j'Tc}\right) s_{c,j'Tc} L_c \\ &\quad + \sum_{c \in C} \sum_{j' \in \mathcal{F}_{fk'}} \left(p_{j'Oc} - mc_{j'Oc}\right) s_{c,j'Oc} L_c.\end{aligned}$$

Because markets are segmented, the demands $s_{c,jkc}$ depend only on the local prices in country c, so the first-order conditions for profit maximizing prices can be solved for each country separately.

After integration of the online distribution channel, consumers in each country c face an increased choice set because they can also purchase in other countries $d \neq c$. To purchase these products abroad, firms may face a shipping cost τ_{cd} to ship products from the country of purchase d to the consumers' country c (where we normalize $\tau_{cc} = 0$) The profit of a firm f after integration therefore becomes:

$$\begin{aligned}\pi_f = &\sum_{c \in C} \sum_{j' \in \mathcal{F}_{fT}} \left(p_{j'Tc} - mc_{j'Tc}\right) s_{c,j'Tc} L_c \\ &+ \sum_{c \in C} \sum_{j' \in \mathcal{F}_{fO}} \sum_{d \in C} \left(p_{j'Od} - mc_{j'Od} - \tau_{cd}\right) s_{c,j'Od} L_c.\end{aligned} \tag{16}$$

The first term captures the profits from selling in the traditional channel. This is the same as before: as this channel is still segmented, consumers do not buy in the traditional channel of other countries (i.e., $s_{c,jTd} = 0$ for $d \neq c$). The second term captures the profits from selling online. The demands by consumers in country c for online products in other countries $d \neq c$, $s_{c,jOd}$, may now be positive, but firms need to incur shipping costs to serve these consumers. As a result, it is no longer posible to solve the first-order conditions for each country separately.

More precisely, after integration the demands in the traditional channel $s_{c,jTc}(\mathbf{p}_{Tc}, \mathbf{p}_O)$ and in the online channel $s_{c,jOd}(\mathbf{p}_{Tc}, \mathbf{p}_O)$ now depend on the domestic price vector of the traditional channel $\mathbf{p}_{Tc}$ and on the price vector across all countries of the online channel $\mathbf{p}_O$. Firms choose prices to maximize total profits across countries (16), taking into account that consumers may consider to also buy products abroad. In the traditional channel T, each price p_{jTc} should satisfy the following necessary first-order condition (for each j and c):

$$\begin{aligned}\frac{\partial \pi_f}{\partial p_{jTc}} &= s_{c,jTc} L_c + \sum_{j' \in \mathcal{F}_{fT}} \left(p_{j'Tc} - mc_{j'Tc}\right) \frac{\partial s_{c,j'Tc}}{\partial p_{jTc}} L_c \\ &+ \sum_{d \in C} \sum_{j' \in \mathcal{F}_{fO}} \left(p_{j'Od} - mc_{j'Od} - \tau_{cd}\right) \frac{\partial s_{c,j'Od}}{\partial p_{jTc}} L_c = 0.\end{aligned} \tag{17}$$

The first row of (17) captures the impact of an increase in the price p_{jTc} on profits in the traditional channel, which does not involve any other countries than c because demand in the traditional channel is segmented. The second row captures the impact of an increase in the price p_{jTc} on the online channel. This also involves other countries $d \neq c$, because consumers who substitute out of product j of the traditional channel may choose to buy online abroad.

In the online channel O, each price p_{jOc} should satisfy the following first-order condition (again, for each j and c):

$$\begin{aligned}\frac{\partial \pi_f}{\partial p_{jOc}} &= \sum_{c' \in C} \sum_{j' \in \mathcal{F}_{fT}} \left(p_{j'Tc'} - mc_{j'Tc'}\right) \frac{\partial s_{c',j'Tc'}}{\partial p_{jOc}} L_{c'} \\ &+ \sum_{c' \in C} \sum_{j' \in \mathcal{F}_{fO}} s_{c',j'Oc} L_{c'} + \sum_{c' \in C} \sum_{j' \in \mathcal{F}_{fO}} \sum_{d \in C} \left(p_{j'Od} - mc_{j'Od} - \tau_{c'd}\right) \frac{\partial s_{c',j'Od}}{\partial p_{jOc}} L_{c'} = 0.\end{aligned} \tag{18}$$

The second row of (18) captures the impact of an increase in the price p_{jOc} on profits in the online channel: this raises profits proportional to the demands from all countries (first term on second row), but it also reduces profits proportional to the online margins by affecting bilateral sales flows across all country pairs (second term on the second row). The first row captures the impact of p_{jOc} on profits in the traditional channel of all countries.

To write these first-order conditions in matrix form, we use the following notation. Let $q_{jkd} = \sum_c s_{c,jkd} L_c$ be the total demand for product j in channel k and country d. Furthermore, let $\mathbf{p}$, $\mathbf{q}$ and $\mathbf{mc}$ be vectors with elements p_{jkd}, q_{jkd} and mc_{jkd}, and τ_c be a vector of shipping costs from country c to all other countries. Use $\mathbf{H}$ to denote the holding or ownership matrix across all alternatives (j, k and d), and use $\mathbf{\Omega}$ to denote the matrix with demand derivatives across all alternatives. In contrast to the case of segmented markets (where we had a matrix $\mathbf{\Omega_c}$ per country c), $\mathbf{\Omega}$ is now a matrix across all countries, and it includes non-zeros for products sold online in other countries. We can then write the first-order conditions (17) and (18) after integration as[13]

$$\begin{aligned} \mathbf{p} \;\; = \;\; & \mathbf{mc} - \left[\mathbf{H} \odot \mathbf{\Omega}\left(\mathbf{p}\right)\right]^{-1} \mathbf{q}\left(\mathbf{p}\right) \\ & + \left[\mathbf{H} \odot \mathbf{\Omega}\left(\mathbf{p}\right)\right]^{-1} \left(\sum_{c \in C} \left(\mathbf{H} \odot \mathbf{\Omega}_c\left(\mathbf{p}\right)\right) \tau_c \right). \end{aligned}$$

The first row describes the pricing condition in the absence of shipping costs, showing a uniform markup term capturing consumer price sensitivities across countries. The second row takes into account the pass-through of shipping costs, which gives rise to non-uniform markups with a higher weight to consumer price sensitivities in domestic countries.

To solve for the post-integration equilibrium, we iterate over firms' best response functions until a rest point of the system is reached. At each iteration consumers update their consideration sets, which we find leads to non-monotonic convergence for a few iterations. Nevertheless, we never encounter convergence problems, so that this simple iterative approach turns out to be sufficient to obtain post-integration market outcomes.

5.3 Results

We perform our counterfactuals based on the adapted BLP model, and for comparison purposes we also consider counterfactuals for the standard BLP model. The parameter estimates of both models are shown in Table 3. We compute the counterfactual equilibria (including the predicted status quo) for each month in our sample, after setting the unobserved quality and marginal cost error terms $\widetilde{\xi}_{jkm}$ and $\widetilde{\omega}_{jkm}$ to zero. We then arrive at a weighted annual average across periods by using the pre-ban total number of units sold at each month as weights.

As discussed, we consider two types of counterfactuals to evaluate the impact of reducing cross-border trade restrictions after a geo-blocking ban. In our first scenario, pre-integration availability (PIA), consumers can purchase abroad only those products that they could previously already purchase at home. In the second scenario, full availability (FA), consumers can buy all products online abroad, even those that were not previously available in their own country. We consider both the possibility of no shipping costs after the ban, and remaining shipping costs after the ban.

We first consider the total effects across all EU countries, and then discuss the distributional effects on consumers across the different countries.

[13]To guarantee that all matrix dimensions are conformable, the dimensions of the country-specific matrices of market share derivatives, the $\mathbf{\Omega}_c$'s, are inflated to match those of their cross-country counterpart, $\mathbf{\Omega}$, by filling in the rows and columns corresponding to the cross-sectional observations that are not covered by the traditional and online choice sets of country c (see Figure 3) with zeros.

Total effects Table 4 shows the total effects across the EU from integrating online markets following the geo-blocking ban. We begin our discussion with our main results for the adapted BLP model (top panel of Table 4), but afterwards compare this with results from a standard BLP model (bottom panel). To put the predicted annual consumer and producer surplus changes (in million euro) in perspective, note that the actual annual revenues in the EU amount to 1042 million euro for the traditional channel and 283 million for the online channel.

We find that integrating online markets results in small total EU effects if consumers do not have the possibility to purchase new products after the ban (PIA scenario). Total output and average EU prices remain essentially unchanged, although this hides a shift out of online (-1.1 percent and 2.2 percent before and after accounting for shipping costs) into the traditional channel. Consumer surplus drops by 19 million or 60 million euros annually. The reason for the small effects is that this scenario purely captures the price convergence effect. This may shift benefits between countries (as we will see later on), but does not have important effects at the EU level.

In contrast, integrating online markets implies more sizeable total EU effects if consumers can purchase products that were not available in their own country (FA scenario). Total output increases by between 0.8 and 1 percent (before and after accounting for shipping costs), and this implies a sizeable increase in online sales (by 6.2 percent and 7.6 percent). Consumer surplus now increases by between 227 and 277 million euros with and without shipping costs. The larger consumer benefits in this scenario are of course due to the product choice expansion effect after the ban on geo-blocking, which more than compensates for the slightly negative impact from the pure price convergence effect. Finally, the ban leads to roughly unchanged firm profits, implying the overall welfare impact of the geo-blocking ban is positive.

In sum, integrating online markets implies sizeable effects because of product choice expansion, especially when compared with the level of e-commerce during the period of our study (which was not larger than 20 percent in most countries). The benefits may increase further over time, as e-commerce will likely become more important, also in many other consumer electronics sectors. This is already apparent when comparing the output effects in the traditional and online distribution channels for both counterfactual scenarios. Because of substantial heterogeneity in the taste for shopping online, the effect of online integration on quantities in the traditional sales channel is small relative that in the online channel.

These conclusions are based on the adapted BLP model, with an idiosyncratic taste parameter at the level of the product ($\varepsilon_{ic,j}$) instead of the product, channel and country-of-purchase ($\varepsilon_{ic,jkd}$). This is important in our setting, not only for uncovering reasonable substitution patterns but also for adequately measuring the welfare effects without including artificial gains from making the same products available in other distribution channels or countries. The bottom panel of Table 4 shows the results from a standard BLP model that includes such mechanical gains. This demand model predicts a very large increase of total output by between 7.2 and 8.4 percent, and of output in the

Table 4: Counterfactual Equilibrium Outcomes: Total Effects across Countries

	adapted BLP			
	PIA		FA	
	shipping costs	no shipping costs	shipping costs	no shipping costs
ΔCS	-60.1	-18.7	227.1	276.8
$\Delta\Pi$	-16.9	-12.0	3.6	9.6
ΔQ (%)	-.17	-.03	.82	1.00
ΔQ_{trad} (%)	.16	.02	-.62	-.79
ΔQ_{on} (%)	-2.23	-1.16	6.17	7.57
	BLP			
	PIA		FA	
	shipping costs	no shipping costs	shipping costs	no shipping costs
ΔCS	2 157	2 288	2 358	2 493
$\Delta\Pi$	311.0	324.2	320.3	335.3
ΔQ (%)	7.23	7.65	7.98	8.41
ΔQ_{trad} (%)	-6.48	-6.85	-7.07	-7.45
ΔQ_{on} (%)	52.1	55.4	59.7	63.2

Note: Changes in consumer surplus and changes in profits are measured in millions of euros per year. These may be compared with actual annual revenues in the EU of 1042 million euro for the traditional channel and 283 million euro for the online channel. All relative changes are reported in percentage points; -0.5 represents a drop by 0.5 percent.

online channel by more than 50 percent in all scenarios. Intuitively, the standard BLP model implies a strong outward shift in demand because the same products create new variety at different channels and countries. Similarly, the consumer surplus increase from the ban is greatly overestimated. In both scenarios, consumer surplus gains would exceed 2 billion euros annually, which is roughly 10 times larger than the estimated gains in our adapted BLP model (FA scenario). Finally, the standard BLP model implies a large positive producer surplus increase by more than 300 million euro, or more than 8 percent. This follows again from the mechanically created additional product variety. But this is at odds with a simple revealed preference argument. The firms themselves prefer to restrict cross-border trade through geo-blocking practices, so it seems difficult to rationalize these practices if firms become more profitable when they are banned from using them.

Distribution of Consumer Gains across Countries We now discuss how a ban on geo-blocking practices may differentially affect consumers across countries. We focus our discussion on the results from the adapted BLP model.[14] To interpret the results, recall our earlier finding that consumers from the high income countries are less price sensitive than consumers from the middle and low income countries, which is reflected in higher markups and higher prices in the high income

[14]The results from the standard BLP model again show that consumer welfare gains are substantially overestimated, but do not give other new insights so we do not elaborate on them.

Figure 7: Counterfactual Price Changes in the Online Distribution Channel

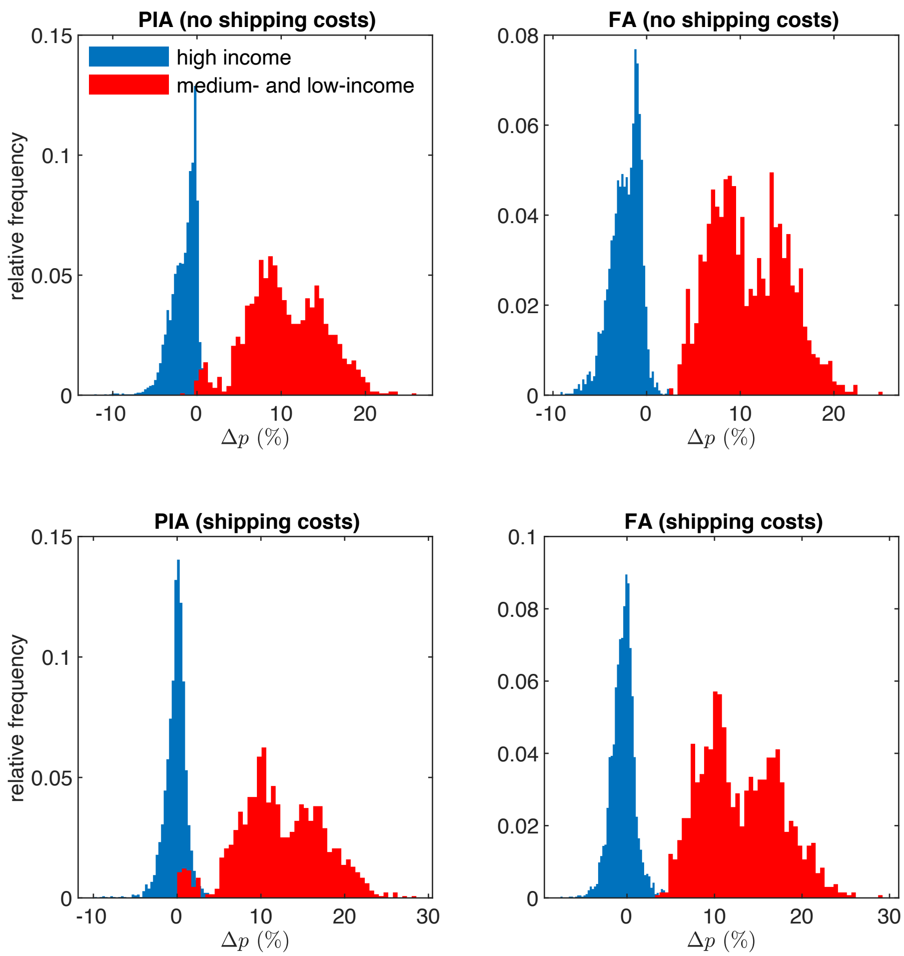

Note: PIA and FA denote the pre-ban online access and full cross-sectional online access counterfactual equilibria. The group of high-income countries are Belgium, Denmark, France, Germany, the UK and the Netherlands, while the low- and middle-income countries are Italy, Spain, Poland and Slovakia.

countries.

Figure 7 shows the distribution of percentage price changes after a geo-blocking ban, separately for the high income countries (blue) and the medium and low income countries (red). The two top graphs show the percentage price changes if there are no shipping costs after the geo-blocking ban. It turns out that almost all products become less expensive in the high income countries, and almost all products become more expensive in the medium and low income countries. If consumers only obtain foreign access for products they could previously purchase at home (PIA scenario), the average price decrease in the high-income countries is 1.5 percent, while the average price increase in the medium income countries is 7.9 percent and online prices in the low-income countries increase by more than 12 percent. If consumers obtain full access to all products abroad (FA scenario), the average effects become somewhat larger. Online prices in high income countries drop by 2.2 percent, while they increase by 7.6 and 12.8 percent in the medium-income and low-income countries, respectively. As such, these findings show how banning geoblocking and thereby achieving online market integration is equivalent to banning third-degree price discrimination between more and less price sensitive consumers. In this case, this actually implies a transfer from the low and medium to the higher income countries.

It is interesting to compare this finding to Dubois, Gandhi and Wasserman (2019). As discussed in the introduction, in their setting (pharmaceuticals) price convergence between countries obtains because of price regulations, whereas in our setting it occurs because of cross-border trade opportunities. But some of their conclusions show interesting parallels. They find that price constraints between the US and Canada would reduce US prices only slightly, and instead mainly raise prices in the smaller country Canada. Hence, both direct price constraints or indirect constraints through cross-border trade may result in comparable price convergence effects, though other implications such as welfare effects may differ.

The two bottom graphs of Figure 7 show that the distribution of percentage price changes is shifted to the right if we account for the presence of shipping costs. Intuitively, firms optimally spread the increase in their marginal costs between the different consumer populations. As consumers impose arbitrage in the online channel, firms cannot differentiate online price by the location from where consumers are shopping. Thus, shipping costs simply become a second source of marginal costs that are spread between all consumers and prices everywhere slightly increase relative to the scenario without shipping costs.

While a look at price changes is intuitive, it does not give a complete picture of the distribution of consumer gains across countries. We therefore consider the impact on consumer surplus in the different countries in Figure 8. First consider the case where there are no shipping costs after the geo-blocking ban (top part). If the ban opens foreign access only to products that were already available at home (PIA scenario, top left), consumers in the high income countries gain at the expense of consumers in the low and medium income countries. In contrast, if the ban opens foreign access to all consumers (FA scenario, top right), the low and medium income countries also gain or lose only slightly, though the gains are again much higher for the high income countries. Intuitively, the PIA scenario mainly involves a transfer of benefits because it purely captures the price convergence effect, while the FA scenario implies gains to all countries because it also captures the product choice expansion effect. If we account for the presence of shipping costs, the same broad picture emerges. The main difference is that the consumer gains are lower (in all countries) because firms partly pass through the shipping costs in consumer prices.

Table 5 takes a further look at the country effects (for the case without shipping costs). In addition to consumer surplus changes, it also shows price and output changes broken down by the two distribution channels. This shows the extent to which there are spillover effects to the

Figure 8: Consumer Surplus Changes

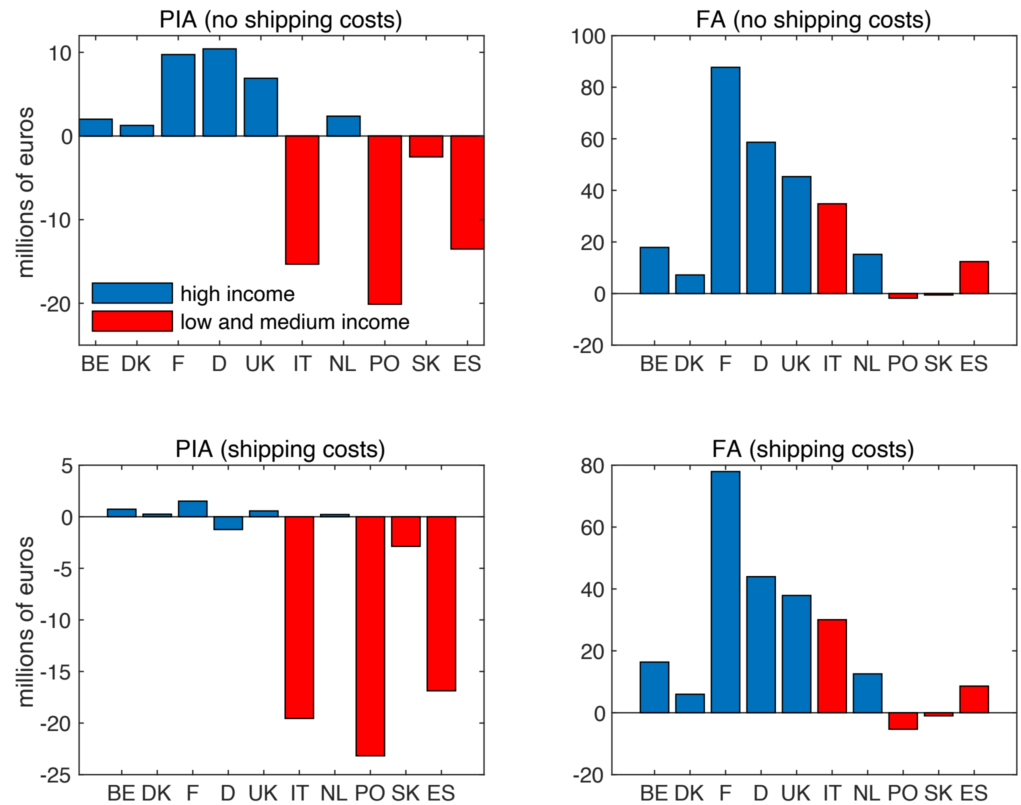

Note: Changes in consumer surplus are scaled up to annual changes. PIA and FA denote the pre-ban online access and full cross-sectional online access counterfactual equilibria. The group of high-income countries are plotted in blue and are Belgium, Denmark, France, Germany, the UK and the Netherlands, while the low- and middle-income countries are shown in red and are Italy, Spain, Poland and Slovakia. All y-axes are fixed at the same range.

traditional channel. The top panel (pre-integration access) gives the sharpest conclusions (because it abstracts from the product choice expansion effect). Countries with high consumer surplus gains (in per capita terms), also see the highest online price drops. Furthermore, they experience some modest price drops in the traditional channel. For example, in Denmark online prices drop by 1.73 percent, inducing a price drop on the traditional channel by 0.05 percent. The extent of substitution is lower in Belgium, because the online channel is less important there. Similarly, in countries with online price decreases there are also increases in online sales (e.g. +1 percent in Denmark), while traditional sales drop because the price drops on the traditional channel are too modest. The reverse findings hold for countries with consumer surplus losses (i.e. price increases online, modest price increases offline, and drops in online sales with a modest shift to traditional sales).

The bottom panel of Table 5 (full access) gives broadly comparable conclusions regarding the

Table 5: Counterfactual Outcomes by Country - Adapted BLP Model (no shipping costs)

	Pre-Integration Access				
		Δp (%)		ΔQ (%)	
	ΔCS (mln euros)	traditional	online	traditional	online
BE	2.011	-.014	-2.175	-.064	2.033
DK	1.264	-.047	-1.727	-.145	1.014
F	9.738	-.023	-1.606	-.066	1.144
D	10.42	-.039	-1.186	-.092	.735
UK	6.901	-.049	-1.457	-.122	.937
IT	-15.32	.021	7.604	.124	-5.707
NL	2.373	-.039	-1.481	-.105	.953
PO	-20.11	.147	12.39	.665	-10.40
SK	-2.500	.374	12.09	1.232	-8.757
ES	-13.52	.050	7.848	.220	-6.389
All	-18.74	-.007	1.279	.019	-1.160
	Full Access				
		Δp (%)		ΔQ (%)	
	ΔCS (mln euros)	traditional	online	traditional	online
BE	17.88	-.021	-2.618	-.823	19.35
DK	7.205	-.073	-2.444	-1.166	6.123
F	87.71	-.028	-2.188	-1.029	11.34
D	58.65	-.063	-1.909	-.747	4.393
UK	45.33	-.079	-2.301	-1.162	6.725
IT	34.80	.047	7.284	-.499	15.50
NL	15.18	-.065	-2.248	-.954	6.582
PO	-1.802	.202	13.10	-.170	-.376
SK	-.570	.404	12.63	-.032	-1.966
ES	12.38	.075	7.628	-.393	6.793
All	276.8	-.013	.765	-.788	7.568

Note: Output changes are computed for the ten populations of consumers, while price changes are computed for the products available for sale in each of the ten countries.

percentage price and output changes: countries with the highest consumer surplus gains show the highest drops in online prices, traditional prices only slightly increase, and there is a shift from traditional to online sales. And the reverse is true for countries with the lowest consumer surplus gains.

6 Conclusion

Governments have taken various measures to remove non-tariff trade barriers and promote market integration. These measures often involve interventions against restrictive distribution practices set up by firms. The European Commission's ban on geo-blocking practices is a recent example, aiming to integrate online markets as part of the Single Digital Market program. In this paper, we develop a framework to evaluate the impact of such online market integration, taking into account possible spillover effects to traditional distribution channels. We adapt the standard random coefficients logit demand model to allow for substitution between multiple distribution channels and to incorporate consumer arbitrage between multiple countries. We also show how to account for the presence of remaining shipping costs after integration.

We apply our framework to the European portable PC market, where geo-blocking restrictions were prevalent during our sample period and have recently been banned. The total consumer and welfare gains from reducing cross-border arbitrage costs are modest, and entirely due to an expanded product choice set rather than reduced third-degree price discrimination possibilities. At the same time, there are considerable distributional effects. Consumers in high income countries gain much more, potentially at the expense of consumers in medium and low income countries.

From a methodological perspective, we show that a straightforward application of the standard BLP demand model is not warranted in our setting. This entails a high-dimensional idiosyncratic taste valuation that is specific to the distribution channel and country of purchase for each product. Such a model generates unreasonable substitution patterns. Furthermore, such a model creates artificial product differentiation as foreign markets open up. This implies implausibly high consumer welfare benefits and even profit gains from the opening up of foreign markets. The latter is inconsistent with the firms' revealed preference for deliberately keeping markets segmented before the geo-blocking ban. We show how our adapted BLP model addresses these issues, and can be estimated in a computationally feasible way.

We hope that this framework can be fruitfully applied in future work to evaluate the impact of increased market integration (or the absence of it) in a variety of other settings.

7 References

Armstrong, Timothy B., (2016), "Large Market Asymptotics for Differentiated Product Demand Estimators with Economic Models of Supply," Econometrica, Vol. 84, No. 5, pp. 1961-1980.

Berry, Steven T., James Levinsohn and Ariel Pakes, (1995), "Automobile Prices in Market Equilibrium," Econometrica, Vol. 63, No. 4, pp. 841-890.

Berry, Steven T. and Ariel Pakes, (2007), "The Pure Characteristics Demand Model," International Economic Review, Vol. 48, No. 4, pp. 1193-1225.

Bonnet, Céline and Pierre Dubois, (2010), "Inference on vertical contracts between manufacturers and retailers allowing for nonlinear pricing and resale price maintenance," RAND Journal of Economics, 41 (1), 139–164.

Brunner, Daniel, Florian Heiss, André Romahn and Constantin Weiser, (2017), "Reliable Estimation of Random Coefficient Logit Demand Models," DICE Discussion Paper No. 267.

Conlon, Christopher and Jeff Gortmaker, (2020), "Best Practices for the Differentiated Products Demand System with Pyblp", forthcoming RAND Journal of Economics.

Dubois, Pierre, Ashvin Gandhi and Shoshana Wasserman, (2019), "Bargaining and International Reference Pricing in the Pharmaceutical Industry", working paper.

Duch-Brown, Néstor, Lukasz Grzybowski, André Romahn and Frank Verboven, (2019), "Are Online Markets More Integrated than Traditional Markets? Evidence from Consumer Electronics", working paper.

Duch-Brown, Néstor and Bertin Martens, (2014), "Consumer Benefits from the EU Digital Single Market: Evidence from Household Appliances Markets," JRC-IPTS Digital Economy Working Paper 2014/04, European Commission.

Eizenberg, Alon, (2014), "Upstream Innovation and Product Variety in the U.S. Home PC Market", Review of Economic Studies, 81, 1003-1045.

EU Regulation, (2018), "Regulation 2018/302 of February 2018 on addressing unjustified geoblocking and other forms of distrimination based on customers' nationality, place of residence or place of establishment within the internal market and amending Regulations (EC) no 2006/2004 and (EU) 2017/2394 and Directive 2009/22/EC", available at `https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=uriserv:OJ.LI.2018.060.01.0001.01.ENG&toc=OJ:L:2018:060I:TOC`.

European Commission, (2016), "Mystery Shopping Survey on Territorial Restrictions and Geoblocking in the European Digital Single Market", Report from GfK prepared for the European Commision, 194pp.

European Commission, (2017), "Final report on the E-commerce Sector Inquiry", Report from the Commission to the Council and the European Parliament, 16pp.

Frankel, Jeffrey and Andrew Rose, (1996), "Currency Crashes in Emerging Markets: an Empirical Treatment", Journal of International Economics, 41(3-4), pp. 351-366.

Gandhi, Amit and Jean-François Houde, (2019), "Measuring Substitution Patterns in Differentiated Products Industries," NBER Working Paper No. 26375, available at `https://www.nber.org/papers/w26375`.

Goldberg, Pinelopi and Frank Verboven, (2001), "The Evolution of Price Dispersion in the European Car Market," Review of Economic Studies, Vol. 68, No. 4, pp. 811-848.

Goldberg, Pinelopi and Rebecca Hellerstein, (2013), "A Structural Approach to Identifying the Sources of Local Currency Price Stability," Review of Economic Studies, Vol. 80, No. 1, pp. 175-210.

Gopinath, Gita, Pierre-Olivier Gourinchas, Chang-Tai Hsieh and Nicholas Li, (2012), "International Prices, Costs and Markup Differences," American Economic Review, 101(6), pp. 2450-86.

Gorodnichenko, Yuriy and Oleksandr Talavera, (2017), "Price Setting in Online Markets: Basic Facts, International Comparisons, and Cross-Border Integration," American Economic Review, Vol. 107, No. 1, pp. 249-282.

Grigolon, Laura and Frank Verboven, (2014), "Nested Logit or Random Coefficients Logit? A Comparison of Alternative Discrete Choice Models of Product Differentiation," Review of Economics and Statistics, Vol. 96, No. 5, pp. 916 - 935.

Grubb, Michael D. and Matthew Osborne, (2015), "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock," American Economic Review, Vol. 105, No. 1, pp. 234-271.

Kanavos, Panos and Joan Costa-Font, (2005), "Pharmaceutical Parallel Trade in Europe: Stakeholder and Competition Effects," Economic Policy, Vol. 20, No. 44, pp. 751-798.

Meschi, M., Irving, T. and Gillespie, M., (2013), "Intra-Community Cross-Border Parcel Delivery. FTI Consulting. A Study for the European Commission," Catalogue number KM-31-13-718-EN-N.

Obstfeld, Maurice and Kenneth Rogoff, (2000), "The Six Major Puzzles in International Macroeconomics: Is There a Common Cause?," NBER/Macroeconomics Annual, 15(1), pp. 339-390.

Pan, Xing, Ratchford, Brian T. and Venkatesh Shankar, (2004), "Price Dispersion on the Internet: A Review and Directions for Future Research," Journal of Interactive Marketing, Vol. 18, No. 4, pp.116-135.

Petrin, Amil, (2002), "Quantifying the Benefits of New Products: The Case of the Minivan," Journal of Political Economy, Vol. 110, No. 4, pp. 705-729.

Rey, Patrick and Thibaud Vergé, "Resale Price Maintenance and Interlocking Relationships," Journal of Industrial Economics, Vol. 58(4), 928-961.

Reynaert, Mathias and Frank Verboven, (2014), "Improving the Performance of Random Coefficients Demand Models: The Role of Optimal Instruments," Journal of Econometrics, Vol. 179, No. 1, pp. 83-98.

Song, Minjae, (2015), "A Hybrid Discrete Choice Model of Differentiated Product Demand with an Application to Personal Computers," International Economic Review, Vol. 56, Issue 1, pp. 265-301.

Sovinsky Goeree, Michelle, (2008), "Limited Information and Advertising in the US Personal Computer Industry", Econometrica, 76, 1017-1074.

Thomassen, Øyvind, (2017), "An Empirical Model of Automobile Engine Variant Pricing," International Journal of the Economics of Business, Vol. 24, Issue 3, pp. 275-293.

Zhang, Junzi, Brendan O'Donoghue and Stephen Boyd, (2018), "Globally Convergent Type-I Anderson Acceleration for Non-Smooth Fixed-Point Iterations," unpublished manuscript, `http://web.stanford.edu/~boyd/papers/nonexp_global_aa1.html`.

A Appendix

A.1 Additional Tables and Figures

Figure A.1: Median income levels

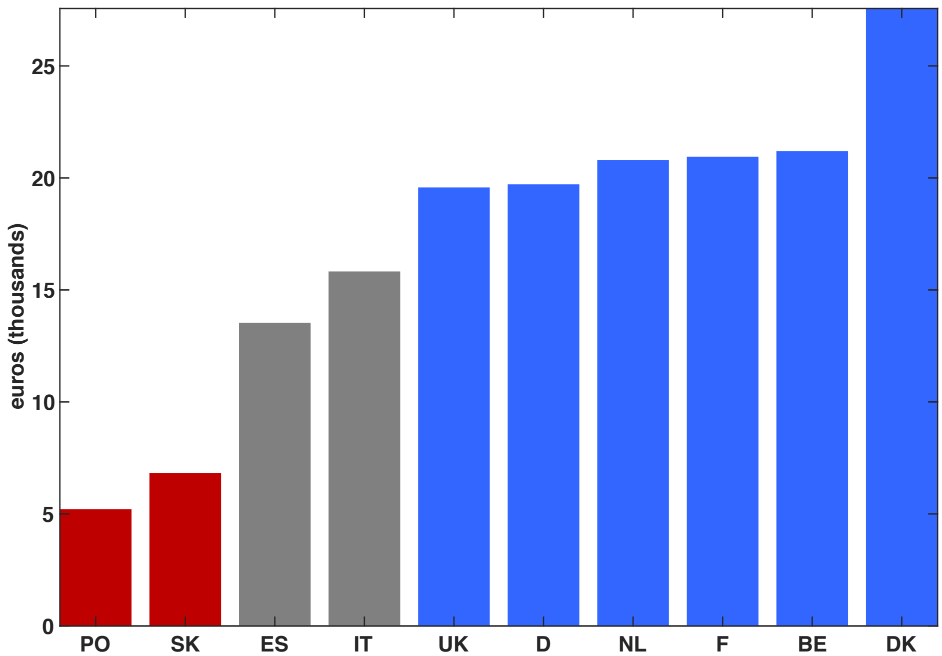

Note: Based on average median incomes over the sample period. The countries are coded as follows: Belgium (BE), Denmark (DK), France (F), Germany (D), Italy (IT), the Netherlands (NL), Poland (PO), Slovakia (SK), and Spain (ES).

Table A.1: First-Stage Regressions of the Demand Model with Alternative Instruments

	(I)	(II)	(III)	(IV)
$z_{Weight,same}$	**-.1507**	**-0.1433**	**-.1458**	**-.1309**
	(.0531)	(.0519)	(.0518)	(.0392)
$z_{Diagonal,same}$	**.1342**	**.1030**	**.1089**	**.1797**
	(.0820)	(.0801)	(.0800)	(.0597)
$z_{Resolution,same}$	**-.0141**	**-.0007**	**-.0029**	**-.0687**
	(.0372)	(.0364)	(.0363)	(.0267)
$z_{Weight,other}$	**-1.325**	**-1.376**	**-1.377**	**-.4074**
	(.0621)	(.0608)	(.0608)	(.0439)
$z_{Diagonal,other}$	**1.631**	**1.717**	**1.719**	**.6075**
	(.0954)	(.0934)	(.0934)	(.0653)
$z_{Resolution,other}$	**-.3804**	**-.4395**	**-.4398**	**-.2614**
	(.0440)	(.0432)	(.0431)	(.0290)
Common Trend	x	x	x	x
Country FE	-	x	x	x
Month FE	-	-	x	x
Product FE	-	-	-	x
Online Trends	-	-	-	x
R^2	.683	.698	.699	.904
F-statistic	173	174	182	22.1

Note: Based on 10 288 sample observations. Standard errors are shown in parentheses. A constant and the matrix of observed characteristics x are included in all regressions. This includes the country-specific means for price and online. The F-statistic is computed for the subset of excluded instruments.

Table A.2: Demand Estimates - Online Means and Trends

	Logit		adapted Logit		BLP		adapted BLP	
	mean	trend	mean	trend	mean	trend	mean	trend
BE	**-2.083**	**.0548**	**.2248**	**-.0094**	**-13.24**	**.1944**	**-15.48**	**.2012**
	(.1618)	(.0172)	(.1673)	(.0178)	(4.515)	(.0406)	(4.417)	(.0530)
DK	**-.7621**	**.0054**	**.1000**	**-.0073**	**-6.928**	**.0123**	**-8.198**	**.0051**
	(.1112)	(.0123)	(.1149)	(.0127)	(2.488)	(.0253)	(2.371)	(.0263)
F	**-1.135**	**.0200**	**.4151**	**-.0494**	**-9.408**	**.0852**	**-11.17**	**.0918**
	(.1316)	(.0144)	(.1361)	(.0149)	(3.556)	(.0257)	(3.438)	(.0251)
D	**-.6450**	**.0229**	**.0410**	**.0014**	**-7.127**	**.0867**	**-8.515**	**.0930**
	(.1024)	(.0115)	(.1059)	(.0119)	(2.679)	(.0195)	(2.549)	(.0203)
UK	**-.9061**	**.0521**	**.2254**	**-.0151**	**-8.266**	**.2519**	**-9.813**	**.2848**
	(.1121)	(.0125)	(.1159)	(.0129)	(3.103)	(.0778)	(3.033)	(.0852)
IT	**-1.653**	**.0077**	**.9494**	**-.0517**	**-12.97**	**.0567**	**-15.50**	**.0725**
	(.1793)	(.0197)	(.1854)	(.0204)	(4.998)	(.0363)	(4.824)	(.0355)
NL	**-1.243**	**.0608**	**.0533**	**-.0003**	**-8.942**	**.2478**	**-10.52**	**.2761**
	(.1166)	(.0127)	(.1205)	(.0131)	(3.128)	(.0620)	(3.072)	(.0701)
PO	**-1.368**	**.0090**	**.3973**	**-.0471**	**-10.13**	**.0790**	**-12.00**	**.0876**
	(.1398)	(.0148)	(.1445)	(.0153)	(3.688)	(.0286)	(3.504)	(.0414)
SK	**-.8221**	**.0824**	**-.1434**	**.0363**	**-7.693**	**.2560**	**-9.211**	**.2852**
	(.1189)	(.0132)	(.1229)	(.0136)	(3.016)	(.0706)	(2.897)	(.0716)
ES	**-2.496**	**.1169**	**-.0057**	**.0330**	**-14.233**	**.3571**	**-16.6695**	**.3948**
	(.1490)	(.0169)	(.1540)	(.0174)	(4.876)	(.0834)	(4.698)	(.0840)

Note: This is a continuation of Table 3 in the main text. Based on 10 288 observations. Standard errors are shown in parentheses.

Table A.3: Marginal Cost Regressions - Adapted BLP Model

	Common Online Cost			Country-Specific Online Cost		
		Traditional	Online		Traditional	Online
CPU speed	**.3712**			**.3732**		
RAM	**.0138**			**.0127**		
Weight	**-.1121**			**-.1016**		
Diagonal	**.0355**			**.0350**		
Resolution	**.6372**			**.6472**		
Constant	**-1.619**			**-1.636**		
Online Trend	**.0002**			**-.0002**		
Trend	**-.0227**			-.0223		
BE		**0**	**-.1766**		**0**	**-.1046**
DK		**.0749**	**-.1766**		**.0773**	**.0099**
F		**-.3297**	**-.1766**		**-.3300**	**-.3982**
D		**-.1813**	**-.1766**		**-.1819**	**-.2358**
UK		**-.2651**	**-.1766**		**-.2647**	**-.2376**
IT		**-.1483**	**-.1766**		**-.1478**	**-.1367**
NL		**-.1320**	**-.1766**		**-.1310**	**-.1571**
PO		**-.1068**	**-.1766**		**-.1074**	**-.1968**
SK		**-.1376**	**-.1766**		**-.1373**	**-.1671**
ES		**-.1322**	**-.1766**		**-.1306**	**-.0868**
R^2	**.8943**			**.8988**		

Note: For the traditional and online marginal cost intercepts the Belgian traditional observations are the base category and the corresponding intercept is therefore normalized to zero. To preserve space, we do not report the standard errors, but summarize the significance pattern as follows. For the common online cost specification, except for the online trend, all estimated coefficients are statistically significant at the 95 percent confidence level. For the country-specific online cost specification, except for the online trend and the online intercept for Denmark, all coefficients are statistically significant at the 95 percent confidence level. The specifications also include product and month-of-year fixed effects, which we do not report.

Table A.4: Marginal Cost Regressions - BLP Model

	Common Online Cost			Country-Specific Online Cost		
		Traditional	Online		Traditional	Online
CPU speed	**.3710**			**.3730**		
RAM	**.0138**			**.0127**		
Weight	**-.1116**			**-.1014**		
Diagonal	**.0354**			**.0349**		
Resolution	**.6371**			**.6475**		
Constant	**-1.623**			**-1.641**		
Online Trend	**.0002**			**-.0002**		
Trend	**-.0227**			**-.0223**		
BE		**0**	**-.1754**		**0**	**-.1043**
DK		**.0752**	**-.1754**		**.0775**	**.0095**
F		**-.3295**	**-.1754**		**-.3298**	**-.3961**
D		**-.1808**	**-.1754**		**-.1815**	**-.2361**
UK		**-.2647**	**-.1754**		**-.2643**	**-.2381**
IT		**-.1447**	**-.1754**		**-.1442**	**-.1307**
NL		**-.1317**	**-.1754**		**-.1307**	**-.1573**
PO		**-.1037**	**-.1754**		**-.1043**	**-.1928**
SK		**-.1342**	**-.1754**		**-.1339**	**-.1638**
ES		**-.1286**	**-.1754**		**-.1270**	**-.0821**
R^2	**.8943**			**.8988**		

Note: For the traditional and online marginal cost intercepts the Belgian traditional observations are the base category and the corresponding intercept is therefore normalized to zero. To preserve space, we do not report the standard errors, but summarize the significance pattern as follows. For the common online cost specification, except for the offline intercept for Denmark, all estimated coefficients are statistically significant at the 95 percent confidence level. For the country-specific online cost specification, except for the online trend and the online intercept for Denmark, all coefficients are statistically significant at the 95 percent confidence level. The specifications also include product and month-of-year fixed effects, which we do not report.

Table A.5: Bilateral Parcel Shipping Costs

	BE	DK	F	D	UK	IT	NL	PO	SK	ES
BE	5.70	30	15	15	15	30	15	30	30	30
DK	30.75	10.07	30.75	28.60	30.75	30.75	30.75	30.75	30.75	30.75
F	12	15.45	7.51	12	14.50	14.50	12	19.50	19.50	14.50
D	17	17	17	6.90	17	17	17	17	17	17
UK	42.28	46.85	46.85	46.85	12.63	49.98	42.28	57.60	57.60	49.98
IT	34	34	34	34	34	11.40	34	34	34	34
NL	14.30	14.30	14.30	14.30	14.30	14.30	8.05	19.80	14.30	14.30
PO	20.61	20.61	20.61	20.61	20.61	20.61	20.61	4.27	20.61	20.61
SK	19	19	19	19	19	19	19	19	2.80	19
ES	30.26	30.26	30.26	30.26	30.26	30.26	30.26	30.85	30.26	6.60

Note: The rates are measured in euros by Meschi et al. (2013) and apply to all parcels with weights between two and five kilograms. Where both economy and priority shipping rates are available, we use the priority rates. Express shipping rates are not used.

Figure A.2: Counterfactual Price Changes in the Traditional Distribution Channel

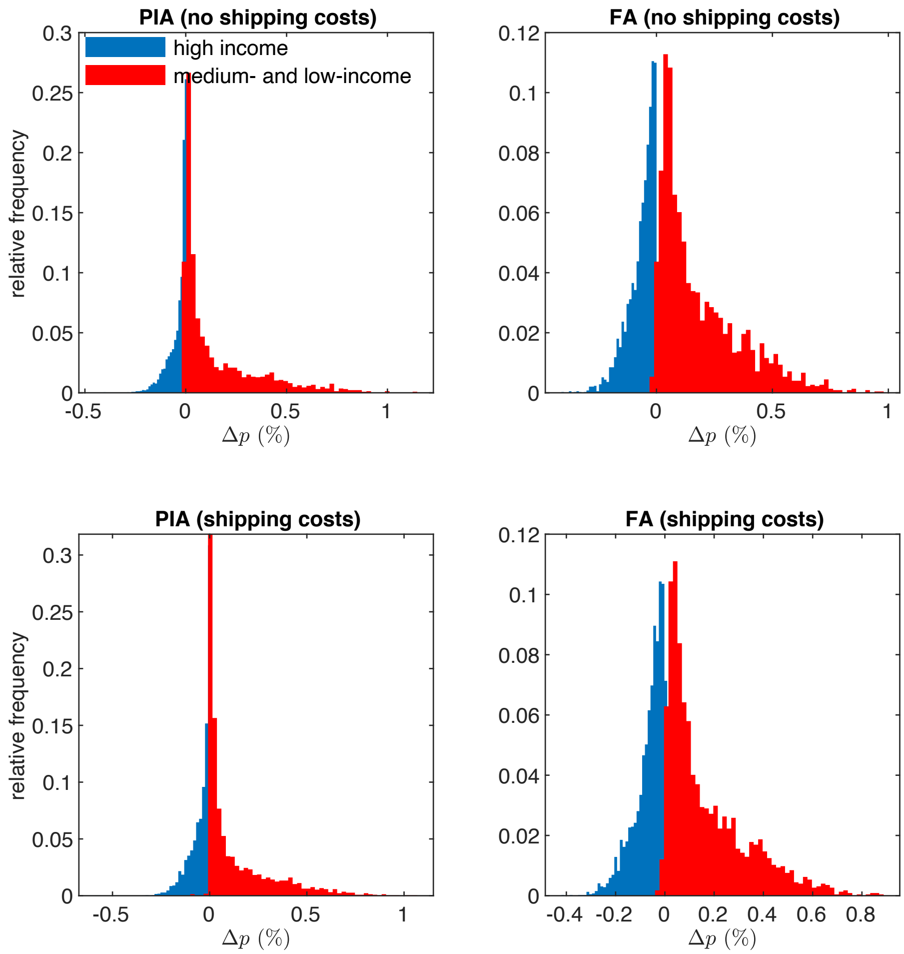

Note: PIA and FA denote the pre-ban online access and full cross-sectional online access counterfactual equilibria. The group of high-income countries are Belgium, Denmark, France, Germany, the UK and the Netherlands, while the low- and middle-income countries are Italy, Spain, Poland and Slovakia.

A.2 Computational Appendix

We provide further computational details in the following three subsections. First, we derive the adapted BLP model as a limiting case of the random coefficients nested logit model. Second, we discuss the inversion of aggregate shares in our estimation. The inversion is considerably slowed down by setting the nesting parameter close to 1. We effectively reduce the increase in computational cost by using a globally convergent Anderson Type-I fixed point acceleration scheme. Third, we provide diagnostics on both our BLP and adapted BLP model estimations.

A.2.1 Approximation of the Adapted BLP Model

As discussed in the main text, we approximate the adapted BLP model with a random coefficients nested logit model, where each product j is a nest containing two alternatives: the traditional and online sales channel. The individual-specific taste parameter in such a set-up is $\varepsilon_{i,j} + (1-\rho)\varepsilon_{i,jk}$, where $\rho \in (0,1)$ is a nesting parameter covering both the standard BLP model ($\rho = 0$) and the adapted BLP model ($\rho \to 1$) as special cases. Hence, to estimate the adapted BLP model we can estimate the random coefficients nested logit model by imposing ρ sufficiently high.

The random coefficient nested logit choice probability (conditional on β_i) for a product j and channel k is equal to

$$s_{jk}(\beta_i) = \int_{-\infty}^{\infty} s_{k|j}(\beta_i, \nu^O) s_j(\beta_i, \nu^O) dF(\nu^O),$$

where

$$\begin{aligned} s_{k|j}(\beta_i, \nu^O) &= \frac{\exp\left(V_{i,jk}/\left(1-\rho\right)\right)}{\sum\limits_{k' \in \{T,O\}} \exp\left(V_{i,jk'}/\left(1-\rho\right)\right)} \\ s_j(\beta_i, \nu^O) &= \frac{\exp\left(I_{i,j}\right)}{1 + \sum\limits_{j' \in \mathcal{J}_k} \exp\left(I_{i,j'}\right)}, \end{aligned}$$

$V_{i,jk} = V_{jk}(\beta_i, \nu^O)$ and $I_{i,j} = I_j(\beta_i, \nu^O)$ is the so-called "inclusive value" defined as

$$I_{i,j} = (1-\rho)\ln\left(\sum_{k' \in \{T,O\}} \exp\left(V_{i,jk'}/\left(1-\rho\right)\right)\right).$$

As $\rho \to 1$, we have $I_{i,j} \to \max\{V_{i,jT}, V_{i,jO}\}$ and $s_{k|j}(\mu_{ij}, \nu^O) \to \mathbf{1}(V_{i,jk} = \max\{V_{i,jT}, V_{i,jO}\})$. We can then write the probability as

$$s_{jk}(\beta_i) = \int_{-\infty}^{\infty} \mathbf{1}(V_{i,jk} = \max\{V_{i,jT}, V_{i,jO}\}) \frac{\exp\left(\max\{V_{i,jT}, V_{i,jO}\}\right)}{1 + \sum\limits_{j' \in \mathcal{J}_k} \exp\left(\max\{V_{i,j'T}, V_{i,j'O}\}\right)} dF(\nu^O).$$

For channel $k = T, O$, we can write this as

$$\begin{aligned} s_{jT}(\beta_i) &= \int_{-\infty}^{\infty} \mathbf{1}(\nu_i^O \leq \Delta_j) \frac{\exp\left(\max\{V_{i,jT}, V_{i,jO}\}\right)}{1 + \sum\limits_{j' \in \mathcal{J}_k} \exp\left(\max\{V_{i,j'T}, V_{i,j'O}\}\right)} dF(\nu^O) \\ s_{jO}(\beta_i) &= \int_{-\infty}^{\infty} \mathbf{1}(\nu_i^O > \Delta_j) \frac{\exp\left(\max\{V_{i,jT}, V_{i,jO}\}\right)}{1 + \sum\limits_{j' \in \mathcal{J}_k} \exp\left(\max\{V_{i,j'T}, V_{i,j'O}\}\right)} dF(\nu^O), \end{aligned}$$

Using the ordering $\Delta_1 \leq ...\Delta_{j-1} \leq \Delta_j \leq \Delta_{j+1} \leq ... \leq \Delta_J$, we can break up the integral in parts to obtain the expressons in the main text, namely

$$
\begin{aligned}
s_{jT}(\mu_{ij}) &= \int_{-\infty}^{\Delta_1} \frac{\exp(V_{i,jT})}{1+D_{i,1}} dF(\nu^O) + \int_{\Delta_1}^{\Delta_2} \frac{\exp(V_{i,jT})}{1+D_{i,2}} dF(\nu^O) + ... + \int_{\Delta_{j-1}}^{\Delta_j} \frac{\exp(V_{i,jT})}{1+D_{i,j}} dF(\nu^O) \\
s_{jO}(\mu_{ij}) &= \int_{\Delta_j}^{\Delta_{j+1}} \frac{\exp(V_{i,jO})}{1+D_{i,j+1}} dF(\nu^O) + ... + \int_{\Delta_{J-1}}^{\Delta_J} \frac{\exp(V_{i,jO})}{1+D_{i,J}} dF(\nu^O) + \int_{\Delta_J}^{\infty} \frac{\exp(V_{i,jO})}{1+D_{i,J+1}} dF(\nu^O).
\end{aligned}
$$

where the terms $D_{i,j}$ (for $j = 1, ...J+1$) are given by the expressions in the last column of Table 2.

A.2.2 Aggregate Share Inversion in the Random Coefficient Nested Logit Model

Grigolon and Verboven (2014) show that the BLP contraction mapping can be applied to the random coefficient nested logit model, provided that the update to mean utilities at each iteration of the fixed point is damped by $(1-\rho)$.

$$
\delta_{jk}^{iter+1} = \delta_{jk}^{iter} + (1-\rho)\ln\left(\frac{\mathcal{S}_{jk}}{s_{jk}(\delta^{iter};\widehat{\theta})}\right) \tag{A.1}
$$

Approximating the adapted BLP model closely requires us to set ρ as high as possible. Specifically, we set $\rho = 0.9$. This slows down the convergence of the fixed point, so that compared to a BLP model, many more iterations are required to obtain the vector of mean utilities that matches the observed and model-implied aggregate market shares. To counteract this increase in computational burden, we apply the globally convergent fixed point acceleration scheme of Zhang, O'Donoghue and Boyd (2018). The approach preserves the global contraction property of the BLP fixed point, while convergence of the accelerated fixed point is no longer guaranteed to be monotonic. The method stores the outcomes of a fixed number of iterations and uses these outcomes to approximate the Jacobian of the nonlinear equation system at low computational cost. If the quality of the approximation is sufficiently good, the iteration takes an approximate Newton step. Otherwise, the damped iteration, (A.1), is used. In practice, we find that the acceleration scheme is highly effective and reduces the required iterations to convergence by a factor of roughly 5. Figure (A.3) plots the convergence path for the damped fixed point (GV) and the accelerated scheme (AA-I).

Figure A.3: Iterations Until Convergence: Damped BLP Contraction versus Globally Convergent Type-I Anderson Acceleration

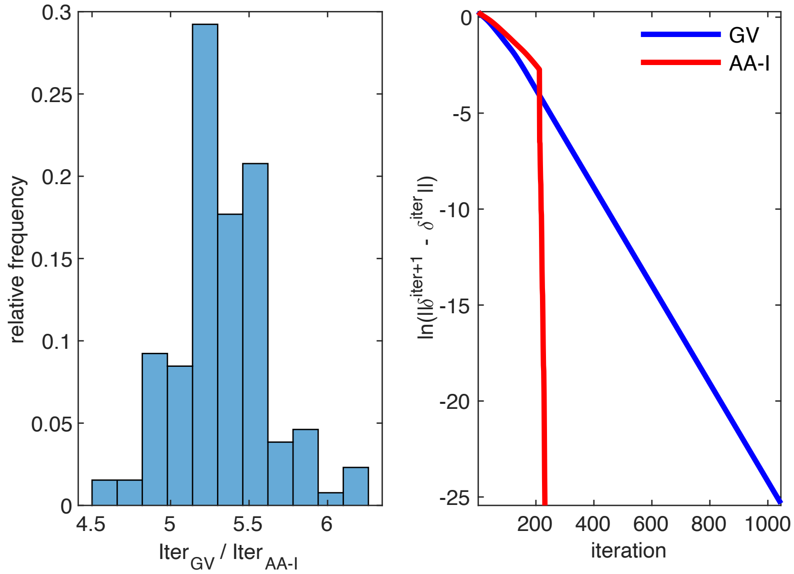

Note: The left panel shows the relative frequency histogram for the two approaches' ratio of iterations until convergence, while the right panel plots the convergence path for the two share inversion schemes evaluated at the coefficient vector that corresponds with the each model's global minimum candidate. GV stands for Grigolon-Verboven and the iteration is given by (A.1). AA-I denotes the globally convergent Anderson Acceleration Type-I scheme of Zhang, O'Donoghue and Boyd (2018). The drop off around 200 iterations for AA-I actually contains several iterations, which is visually imperceptible due to the scale of the x-axis.

A.2.3 Diagnostics

The estimation of the BLP and adapted BLP (or parameterized random coefficient nested logit) models are based on 1 000 modified latin hypercube sampling (MLHS) draws and 30 randomly drawn initial iterates for the nonlinearly entering parameters, $\theta_2^{blp/rcnl} = (\sigma_{on}, \sigma_{RAM}, \sigma_{ppi})'$. To approximate the adapted BLP model closely, we parameterize the nesting coefficient in the random coefficient nested logit model to $\rho = 0.9$. Our BLP estimation routine returns either the positive or negative square root of the squared entries in θ_2. We restrict the estimates of the entries in θ_2^{rcnl} to be positive, because allowing for $-\sqrt{\sigma_{on}^2}$ changes the sign of the cutoffs in the adapted BLP model, which unnecessarily complicates the computation of the model-implied aggregate shares.[15]

The inner convergence tolerance is set to 10^{-11} for inverting the aggregate market shares and we use a trust region optimizer with analytical gradients to minimize the nonlinear GMM-IV objective functions for both models. An extreme value of the objective function is classified as a local minimum if the norm of the gradient is close to zero and the objective function's Hessian is

[15]In the BLP model, $-\sqrt{\theta_{2,k}^2}$ is equivalent to $\sqrt{\theta_{2,k}^2}$ as long as the distribution of ν is symmetric around zero, which holds for $\nu_{ik} \sim N(0,1)$, and the number of simulation draws is large.

positive definite. For the BLP and adapted BLP estimations, the coefficients of variation of the local minima are 1.21 and 1.42 percent, respectively. The tight clustering of the local minima is evidence that the propagation of simulation error in the objective functions is bounded, so that the estimators yield consistent and asymptotically normal estimates (see Berry, Linton and Pakes (2004)).

As Brunner et al. (2017) show, variation in local minima that is due to simulation error can yield substantial variation in model-implied economic outcomes. To evaluate whether the remaining variation between local minima is economically important, we compute the own-price elasticities for all observations in the sample for all local minima. Table (A.6) reports the outcomes. Pooling the model-implied own-price elasticities for all local minima, the bold figures represent the average value of the own-price elasticity at the given percentile. The figures in square brackets are the corresponding minimum and maximum values. Clearly, there is very little variation in the elasticities across all the minima. This holds along the entire distribution of elasticities and for both the BLP and adapted BLP estimations. We conclude that our estimates are not affected by the propagation of simulation error in the GMM-IV objectives in any economically meaningful way.

Table A.6: Model-Implied Own-Price Elasticity Distributions of All Local Minima

			percentiles		
	1^{st}	25^{th}	median	75^{th}	99^{th}
BLP	**-9.15**	**-4.52**	**-3.33**	**-2.56**	**-1.74**
	[-10.6, -8.50]	[-4.68, -4.40]	[-3.43, -3.26]	[-2.64, -2.50]	[-1.81, -1.70]
RCNL	**-9.02**	**-4.69**	**-3.49**	**-2.69**	**-1.82**
	[-9.18, -8.87]	[-4.72, -4.65]	[-3.51, -3.46]	[-2.70, -2.68]	[-1.83, -1.76]

Note: The average values of the own-price elasticity between all local minima is reported in bold. The minimum and maximum values of the own-price elasticity at the respective percentiles is reported in square brackets.