WORKING paper

BANQUE DE FRANCE
EUROSYSTÈME

# Gravity beyond CES

**Paul Piveteau[1] and Gabriel Smagghue[2]**

December 2024, WP #976


## Abstract

We derive a linear structural gravity equation that allows for rich substitution patterns based on observable characteristics. To achieve this, we take advantage of recent econometric work to linearize an import demand system with mixed CES (Constant Elasticity of Substitution) preferences. Compared to traditional gravity models, the resulting equation features additional regressors that capture heterogeneity in the patterns of substitution across exporters. Importantly, this equation can be easily estimated through two stage least squares (2SLS) and without additional data requirements relative to traditional gravity. We implement this method using bilateral trade data and find that the data strongly rejects the Independence of Irrelevant Alternative (IIA) assumption implied by standard trade models: we find an important role for vertical and geographical differentiation so that exporters with similar prices, or originating from similar regions, are closer substitutes. We show that this pattern has important implications in the context of the recent (2018-2019) US-China trade war, in which our model can correctly predict which countries benefitted the most from the reallocation of trade flows due to US tariffs on Chinese imports.


Keywords: Gravity Equation; Trade Wars; Substitution Patterns; Mixed Preferences.

JEL classification: F14, F13.

[1] Cornerstone Research, paul.piveteau@cornerstone.com

[2] Banque de France, gabriel.smagghue@banque-france.fr. We are grateful to seminar audiences at Banque de France, CEPII, University of Strasbourg and GEN Workshop. We particularly wish to thank Thierry Mayer, Vincent Vicard and Houssein Guimbard for comments and suggestions, and Yuyin Cheng for research assistance. Part of this paper was written while Piveteau was visiting CEPII: we thank them for their hospitality and their help with the data.

# NON-TECHNICAL SUMMARY

In this paper, we address a critical limitation of the gravity equations of international trade—the assumption of the Independence of Irrelevant Alternatives (IIA), which posits that all varieties of goods are equally substitutable. Gravity equations are widely used empirical models in trade economics that explain bilateral trade flows based on factors such as economic size, distance, and trade costs between countries. These equations, grounded in analogy to Newtonian gravity, predict that trade between two countries increases with their economic size and decreases with the distance between them. While they are celebrated for their predictive power and theoretical elegance, traditional gravity models rely on simplifying assumptions, including the IIA assumption, which posits that all varieties of goods are equally substitutable. This constraint neglects the nuanced ways competition unfolds across countries, particularly among exporters with similar characteristics such as price or geographical origin. As a result, conventional models fail to capture the heterogeneous effects of trade shocks across competing exporters.

We propose an alternative approach by deriving a linearized gravity equation that incorporates observable characteristics to capture realistic substitution patterns. Building on recent econometric developments, our method rejects the restrictive IIA assumption and introduces artificial regressors that quantify the role of price and regional differentiation in trade competition. Our model allows for richer substitution dynamics, demonstrating that exporters with similar prices or shared regional traits are closer substitutes. Importantly, our framework retains the simplicity of estimation associated with traditional models, requiring no additional data and leveraging on two-stage least squares (2SLS) for implementation.

We empirically validate our approach using trade data in two significant contexts: the "China shock" (the massive surge in Chinese exports following its entry into the WTO in 2001) and the U.S.-China trade war in 2018-2019. Our findings reveal that countries with similar prices to China were disproportionately impacted by the rise of Chinese exports during the China shock. Conversely, during the U.S.-China trade war, countries such as Vietnam, India, and Turkey (on the left- hand side of the chart below), which offer goods similar to those from China, benefitted the most from the reallocation of trade flows. These outcomes stand in sharp contrast to the predictions of standard CES-based models, which suggest uniform effects across competitors, regardless of their characteristics.

Our method offers a practical and tractable framework for analyzing trade dynamics. By introducing heterogeneity in substitution patterns, we enhance the explanatory and predictive power of gravity models. This improvement is particularly relevant for policymakers, as it provides a more detailed understanding of how trade policies redistribute market shares and affect global competition. Furthermore, our results underscore the significance of vertical and geographical differentiation, highlighting their role in shaping trade outcomes in response to shocks.

**Relationship between Cross-Price Elasticity and Price Distance from China**

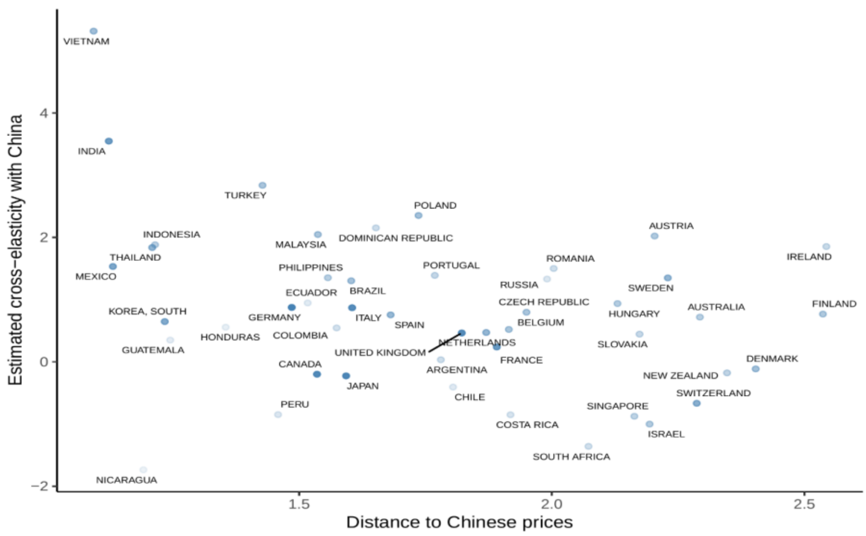


Note: This scatterplot illustrates the relationship between cross-price elasticity with China and the distance to Chinese prices, highlighting a key finding of this study. Countries with goods priced similarly to Chinese exports exhibit higher cross-price elasticities, meaning they benefit more from trade shocks affecting China negatively, such as tariff increases. Conversely, countries with greater price differentiation experience weaker substitution effects. This pattern underscores the limitations of traditional CES-based gravity models, which assume uniform substitutability, and highlights the importance of accounting for heterogeneity in substitution patterns as proposed in our model.

---

# La gravité au-delà de la CES

## RÉSUMÉ


Nous dérivons une équation de gravité structurelle linéaire permettant des schémas de substitution riches basés sur des caractéristiques observables. Pour ce faire, nous exploitons des travaux économétriques récents pour linéariser un système de demande d'importation avec des préférences CES mixtes. Comparée aux modèles de gravité traditionnels, l'équation obtenue intègre des variables explicatives supplémentaires qui capturent l'hétérogénéité des schémas de substitution entre exportateurs. Il est important de noter que cette équation peut être facilement estimée par la méthode des doubles moindres carrés (2SLS) et sans nécessiter de données supplémentaires par rapport aux modèles de gravité traditionnels. Nous appliquons cette méthode aux données de commerce bilatéral et constatons que les données rejettent fortement l'hypothèse d'indépendance des alternatives non pertinentes (IIA) implicite dans les modèles de commerce standard : nous mettons en évidence un rôle significatif pour la différenciation verticale et géographique, indiquant que les exportateurs aux prix similaires ou provenant de régions proches sont des substituts plus proches. Nous montrons que ce schéma a des implications importantes dans le contexte de la récente guerre commerciale entre les États-Unis et la Chine (2018-2019), où notre modèle peut prédire correctement les pays qui bénéficient le plus de la réallocation des flux commerciaux due aux droits de douane américains sur les importations chinoises.



Mots-clés : équation de gravité, guerres commerciales, préférences mixtes.




---

# 1 Introduction

The gravity equation of international trade is one of the most successful empirical tools in economics. It is easy to estimate, has strong predictive power and is theoretically grounded. Owing to these strengths, the gravity equation has become a major tool to quantify the impact of trade frictions on the patterns of trade and on welfare. However, the gravity model features important limitations regarding the patterns of competition across countries. Specifically, it imposes the Independence of Irrelevant Alternative assumption (IIA), which rules out any role for product differentiation and implies that all varieties served to a market are equally substitutable. As a result, the gravity equation is silent about the redistributive effects of bilateral trade policy across third countries. For instance, a standard gravity equation implies that an increase in US tariffs on Chinese goods would have the same impact on Canadian and Vietnamese exports to the US. However, we expect countries with similar characteristics to China to benefit significantly more from the reallocation of US imports.

In this paper, we derive a linear structural gravity equation that allows for rich substitution patterns based on observable characteristics. We take advantage of recent work by Salanié and Wolak (2022) to linearize an import demand system with mixed CES preferences. The resulting gravity equation features additional regressors that capture heterogeneity in the patterns of substitution across exporters. Importantly, this model can be easily estimated through two stage least squares (2SLS) and without additional data requirements relative to traditional gravity models. We implement this method using bilateral trade data and find that the data strongly rejects the IIA assumption implied by standard trade models: we find an important role for vertical and geographical differentiation so that exporters with similar prices, or originating from similar regions, are closer substitutes. We show that this pattern has important implications in the context of the recent US-China trade war, in which the model can predict which countries benefit most from the reduction of Chinese exports to the US.

We start this paper by providing evidence of violation of the IIA assumption implied by traditional gravity models. We show that trade data exhibit a role for vertical differentiation that allows some exporters to be relatively more protected from competition shocks taking place in segments of the market far from them. We provide evidence of this phenomenon in two contexts. During the China shock first, we find that the growth of China was mostly at the expense of exporting countries located close to them in the price distribution. By contrast, countries with much higher or lower prices were less affected by this competition shock. Importantly, we estimate this heterogeneous response within 6-digit product categories. Therefore the fact that some countries are more exposed to Chinese cometition is not driven by their industry composition but rather by their positioning along the price distribution within products. Second, in the context of the recent US-China trade war, we show that countries with similar prices to China benefit the most from the imposition of US tariffs on imports from China. These two pieces of evidence demonstrate that patterns of substitution are far from symmetric across exporters, and that the export performance of a country responds disproportionally to trade shocks that affect countries located closely in the price distribution.

Based on this observation, we develop a tractable estimation method that can capture these patterns of substitution across exporters. The starting point of our approach is to introduce heterogeneity in consumer preferences in a standard CES demand system. This heterogeneity generates stronger substitution patterns between similar varieties as they compete over a common subset of consumers. Moreover, while this mixed-CES model is notoriously difficult to estimate, we employ recent techniques developed in Salanié and Wolak (2022) to linearize the model and facilitate its estimation. The intuition is to perform a "small-sigma" expansion to obtain a linear estimator in which the degree of consumer heterogeneity can be identified through the addition of new regressors to the gravity equation. These so-called "artificial" regressors – that result from the approximation of the model – are measures of differentiation of varieties along some attributes and identify the extent of consumers' heterogeneity in preferences for these characteristics. For instance, the variance in price-elasticity across consumers can be estimated by including as additional regressor the distance of a variety's price to the average price in the market.

While generating the same realistic patterns of substitution, this approximation method has several advantages over the estimation of the full model. First of all, it is significantly easier and transparent to estimate: the model can be directly estimated through 2SLS using regressors and instruments that can be easily computed from the data. Second, the method does not require to specify the full distribution of preferences across consumers. Our simulations indicate that the method is more robust to non-gaussian distributions in preferences than the standard estimation method. Finally, the linearity of the model also facilitates the post-estimation computation of objects of interest, such as cross or own price-elasticities. One of the contributions of this paper is to show how to derive these objects so that they can be directly calculated from the data, without relying on the integration of the full distribution of preferences across consumers.

We estimate our augmented gravity model on bilateral country-level trade data. We are able to strongly reject the IIA assumption as the coefficient on the artificial regressor for prices is significant, both statistically and economically. All things equal, this means that an exporting country sells more to a market when it is located far away from the average price of its competitors. The gain from price differentiation is important economically: a variety that is positioned 0.5 log-point away from the average price (either above or below) enjoys an additional 7 percent export value relative to a variety with a price equal to the mean. Moreover, we find that the region of origin is another dimension of product differentiation that matters for substitution patterns. To obtain this result, we bin source countries into 14 regions and construct the artificial regressor measuring the role of regional differentiation. Including this artificial regressor in the gravity equation captures the possibility that varieties from a same region are closer substitutes. We find that a country belonging to a region with a market share of 20 percents exports 3.5 percent less than a country from a region with a market share of 10 percents. Overall, these results illustrate how our method can flexibly and easily accommodate for deviations from the IIA assumption in trade patterns, and that these deviations are economically significant.

In order to assess the out-of-sample performance of our method, we circle back to the motivating evidence and study the ability of our estimated model to predict the reallocation patterns across countries that resulted from the trade war. Importantly, this prediction is made out of sample, in the sense that it only relies on information available before the actual start of the trade war. It is therefore a legitimate assessment of the improved predictive power of our model, compared to standard gravity models. We find that the model does a good job at predicting which countries won the most from the reduction in Chinese exports to the United States. In particular, the model identifies correctly identifies countries like Vietnam, India or Turkey as the biggest winners due large cross-price elasticities with China. When we run a country-level regression between the actual and the predicted effect of the trade war, we find an estimated coefficient of 0.95 (p-value $< 0.001$) and a $R^2$ equal to 0.27. As a comparison, we also report the prediction in the case of a CES model. In this scenario, the model does not predict any heterogeneity in the reallocation of Chinese market shares across countries: all countries feature the same cross-price elasticity which is equal to the estimated elasticity of substitution of the CES model, 2.29. In conclusion, our model that allows for heterogeneous patterns of substitutions across countries is a significant improvement to predict the winners and losers from the reallocation effects of a large trade shock.

This paper adds to the literature on demand estimation for differentiated product markets, recently surveyed by Berry and Haile (2021) and Gandhi and Nevo (2021). This paper relies on an established tradition in Industrial Organization, started with Berry (1994) and Berry, Levinsohn, and Pakes (1995), which introduces heterogeneity in preferences to generate realistic patterns of substitution across varieties. More recently, Salanié and Wolak (2022) shows how to linearize the model to avoid the well-known challenges associated with its estimation. We extend this estimation strategy to the specific needs of trade data. We notably allow for the presence of demographic drivers of consumer heterogeneity and we consider mixed CES instead of mixed logit.

Our paper also follows an immense literature estimating gravity equations using international trade flows.[1] A series of recent papers have shown how to extend this gravity estimation to abandon the IIA assumption and allow for more realistic substitution patterns. Adao, Costinot, and Donaldson (2017) estimate a mixed-CES model of factor demands to reduce the number of parametric assumptions required to perform counterfactual trade experiment. Lind and Ramondo (2018) extend the standard Ricardian model to allow for more flexible substitution patterns. Relative to these papers, we develop an empirical strategy that captures similar patterns of substitution, but with a much simpler estimation, by 2SLS. Moreover, the existing papers mostly study substitution patterns at the country-level. Therefore, it is not clear whether two countries are highly substitutable because they specialize in the same product categories, or because they specialize in the same market segments, within products. By contrast, we focus on violations of the IIA assumption within detailed product categories. In terms of industrial policy, our results point to the importance of within-product positioning, both for the export performance and for the resilience to competition shocks in foreign markets.

Finally, several recent papers have provided evidence for the importance of capturing these complex patterns of substitution in an international trade context. Fajgelbaum, Goldberg, Kennedy, Khandelwal, and Taglioni (2021) estimates a translog demand system using US-China trade wars to identify the patterns of substitution that follows the reallocation of trade flows between China and the US. Piveteau and Smagghue (2022) shows that the competition shock emanating from the China shock predominantly affected French firms with low-prices, consistent with heterogeneity in substitution patterns along the quality ladder. By contrast, Head and Mayer (2021) shows that omitting mixed preferences when estimating the gravity equation can lead to satisfactory predictions when focusing on the aggregate impact of counterfactual experiments. However, they recognize that ignoring this heterogeneity in preferences can lead to minimize the heterogeneous impact of trade shocks across countries. We follow their approach by comparing our predictions between mixed and simple CES models, but our paper strongly reduces the cost of incorporating mixed preferences by providing a simple estimation strategy.

We start this paper by providing evidence of deviations from the IIA assumption in the next section. We then develop a model in section 3 that can capture more realistic patterns of substitution. Section 4 shows the estimation results of the model and section 5 performs out-of-sample predictions to assess its performance.

## 2 Motivating Evidence

We start this paper by providing empirical evidence that the patterns of substitution across exporting countries deviate from the Independence of Irrelevant Alternatives (IIA) assumption implied by the CES demand system. We study two recent episodes – the China shock and the US-China trade war – to show that countries with similar prices display stronger substitution patterns. This evidence is at odds with a CES demand system that assumes that all varieties are equally substitutable.

In the next subsections, we present the data, specifications and results that lead to this conclusion, first from the China shock, and then from the 2018 US-China trade war.

### 2.1 Evidence from the China shock

In this section, we show that the growth of China in export markets during the last 20 years was at the expense of exporting countries that displayed similar prices to Chinese exporters.

[1] See Head and Mayer (2014) and Yotov, Piermartini, Larch et al. (2016) for recent surveys of this literature.

**Data** We employ trade data from BACI produced by the CEPII,[2] to obtain export data at the country×product level from 1996 to 2018. We focus on exports to the US market and restrict our sample to the top 100 exporting countries to limit the risk of small countries driving some results with noisy data. We define a product as a HS6 product category, which leads to a dataset of more than 2.8 millions observations defined as an exporting country × HS6 product × year triplet. In figure 1, we report the aggregate Chinese market share in total imports from 1996 to 2018. The figure documents the striking growth of China during the last decades, growing from 5 percents of imports in 1996 to more than 20 percents in 2018.


Figure 1: Growth of Chinese market share among US imports

The specification presented in the next paragraph aims at identifying the role of prices to understand which exporters suffered the most from this impressive growth. In the context of trade data, we use unit values – the ratio of export values to quantities in weight – to proxy prices, and we will use these two terms interchangeably in this paper.

**Specification** In this section, we investigate how the price positioning of exporting countries shapes the patterns of substitution with their competitors. In the context of the China shock, our hypothesis is that within a destination×product market, varieties whose prices are more similar to Chinese prices should suffer more from the China shock because they are most substitutable to Chinese varieties.

To perform this exercise, we start by defining how we measure the proximity to Chinese prices. For each origin country×product pair $ok$, we identify the log-price at initial time $t_0$, which is the first time this export flow appears in our dataset.[3] Then, we compare this initial price – $\ln p_{okt_0}$ – to the price of Chinese varieties the same year. As a result, each country×product pair is characterized by its relative log-price defined as $\widetilde{\ln p}_{ok} \equiv \ln p_{okt_0} - \ln p_{Ckt_0}$.

Our specification then interacts the market share of China for a specific product and year – $ms_{Ckt}$ – with the quadratic distance between the price of that variety and the Chinese price. Formally, our

[2] see Gaulier and Zignago (2010) for a full description of the data

[3] Most export flows appear in 1996, the first year of the dataset, but we also include export flows that start in later years.

estimating equation is

$$\ln export_{okt} = \alpha\, ms_{Ckt} + \beta\, \widetilde{\ln p}^2_{ok} \times ms_{Ckt} + \delta_{ok} + (\delta_{kt}) + \varepsilon_{okt}$$

with $export_{okt}$ the total export value from $o$ to the US for product $k$ at time $t$, and $\delta$ a set of two fixed effects: a first fixed effects controls for time-invariant exporter characteristics and a second controls for market-specific unobservables. This specification ensures that parameter $\beta$ measures the relative change in export performance in response to a change in the Chinese market share. We also run this specification without the market fixed effects $\delta_{kt}$ to measure the average effect of the change in Chinese market share through the parameter $\alpha$.

As a less parametric alternative to the quadratic distance, we also divide the distribution of relative prices to measure the heterogeneous effect of Chinese competition along the price distribution. Figure 2 reports the distribution of the log-prices relative to China. We see a large variance in this distribution and an average relative price above zero, as more countries have higher prices than China. This figure also reports the thresholds that we use as alternative to the quadratic distance to categorize different varieties according to their distance to Chinese prices. We set these thresholds at -3, -1, -0.2, 0.2, 1, 3 and 5, to create eight groups of products for which we will separately measure the impact of Chinese competition. This specification also allows us to test whether moving away from Chinese prices on either side of the price distribution reduces the degree of substitution with Chinese varieties.

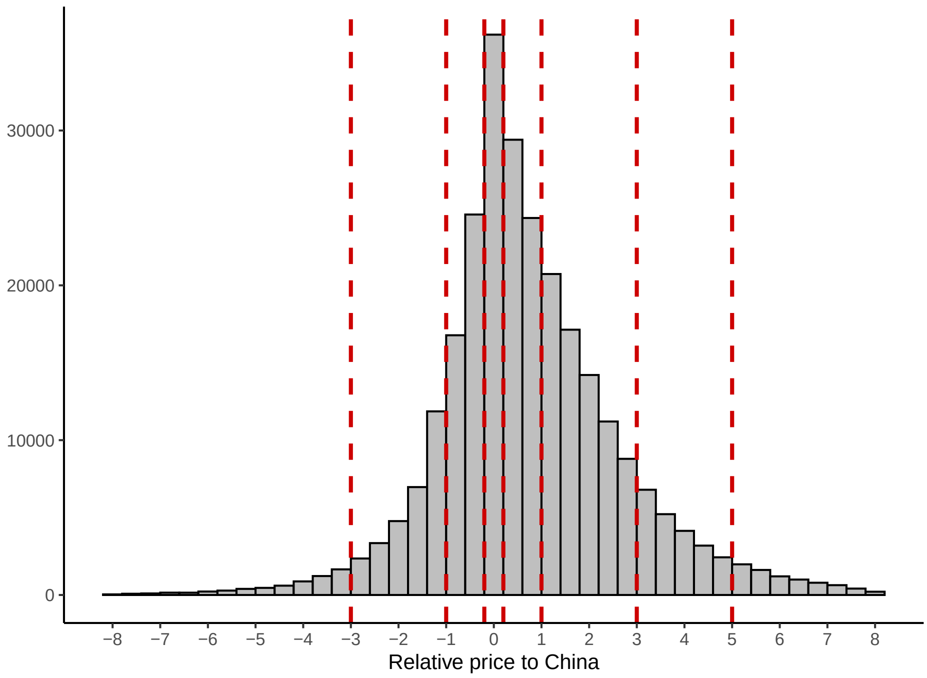


FIGURE 2: Distribution of initial export log-prices relative to China

Using the Chinese market share as independent variable in the regression could generate some concerns regarding the identification of the estimated relationship. For instance, a reduction in the export performance of a country could cause the growth of China in a specific market. Alternatively, a demand shock for a specific level of quality could generate a positive correlation of the export performance of varieties with similar prices. In order to limit these issues, we follow a similar strategy to ADH, by instrumenting the market share of China with its market share in eight other countries, selected due to their ressemblance with the US.[4] We use the Chinese market share in these markets to ensure that our

[4]Similar to ADH, we use the following countries: Australia, Denmark, Finland, Germany, Japan, New Zealand,

Table 1: Competitive Pressure by China decreases with the Distance to Chinese prices

| | OLS | | | | 2SLS | |
|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) |
| $ms_{Ckt}$ | −0.84*** | −1.11*** | | | | |
| | (0.08) | (0.12) | | | | |
| $ms_{Ckt} \times \widetilde{\ln p}_{ok}^{2}$ | 0.03*** | | 0.03*** | | 0.04*** | |
| | (0.00) | | (0.00) | | (0.01) | |
| $ms_{Ckt} \times \widetilde{\ln p}_{ok} \in [-\infty, -3]$ | | 0.40** | | 0.32** | | 0.34 |
| | | (0.17) | | (0.13) | | (0.32) |
| $ms_{Ckt} \times \widetilde{\ln p}_{ok} \in [-3, -1]$ | | 0.36*** | | 0.07 | | 0.32 |
| | | (0.12) | | (0.09) | | (0.20) |
| $ms_{Ckt} \times \widetilde{\ln p}_{ok} \in [-1, -0.2]$ | | 0.30*** | | 0.10 | | 0.17 |
| | | (0.08) | | (0.07) | | (0.12) |
| $ms_{Ckt} \times \widetilde{\ln p}_{ok} \in [0.2, 1]$ | | −0.02 | | 0.02 | | −0.04 |
| | | (0.11) | | (0.09) | | (0.13) |
| $ms_{Ckt} \times \widetilde{\ln p}_{ok} \in [1, 3]$ | | 0.61*** | | 0.58*** | | 0.61*** |
| | | (0.11) | | (0.10) | | (0.13) |
| $ms_{Ckt} \times \widetilde{\ln p}_{ok} \in [3, 5]$ | | 1.08*** | | 0.89*** | | 1.13*** |
| | | (0.14) | | (0.13) | | (0.17) |
| $ms_{Ckt} \times \widetilde{\ln p}_{ok} \in [5, \infty]$ | | 1.24*** | | 1.01*** | | 1.26*** |
| | | (0.19) | | (0.14) | | (0.20) |
| Product × year FE | | | X | X | X | X |
| First stage F-stat | | | | | 384.5 | 107.3 |
| $R^2$ | 0.795 | 0.795 | 0.814 | 0.814 | 0.814 | 0.814 |

Number of observations: 2 850 065. Standard errors are clustered at the country and product levels. * p < 0.1, ** p < 0.05, *** p < 0.01

results are driven by the supply shock that generates the growth of China in global markets, rather than local market conditions.

**Results** We present the results of our regressions in table 1. The table shows that varieties located far away from Chinese prices experience a smaller reduction in export performance when the Chinese market share increases. This relationship is true across all our specifications: in columns (1) and (2) without product×year fixed effects, we find that a 10 point increase in the Chinese market share leads to a reduction of 8 to 11 percent in exports for firms with prices similar to China. These columns also show that this effect decreases as we look at varieties with much lower or higher prices. In columns (3) and (4) that control for product-year unobservables and therefore do not identify the average effect, we find similar results, although the relief for varieties with lower prices than China appears to be limited. Finally, the specification that instruments the Chinese market shares with other destinations confirms our results. The identification strategy leads to much larger standard errors, which makes some results insignificant, but we still find that varieties are less affected from Chinese competition when their prices are much higher than Chinese prices.

Overall, these results confirm that not all exporters are impacted similarly by the growth of China in foreign markets. Exported varieties with similar prices, and probably similar quality, are much closer substitutes to Chinese varieties. As a result, they experience a larger reduction in their export perfor-

Spain and Switzerland.

mance when China gains market share in a foreign market. In the next subsection, we confirm similar patterns of substitution in the context of the US-China 2018 trade war.

## 2.2 Reallocation patterns during the 2018 trade war

Unlike our previous example, the episode of the 2018 trade war between the US and China offers a more experimental framework. During the summer of 2018, the United States increased their tariffs on hundreds of products originating from China. In this section, we use this episode to quantify the negative impact of these tariff increases on Chinese exports, and identify the patterns of reallocation that took place across other exporters to the US.

**Data** For this section, we use data from the US Census to obtain values and quantities imported by the US, disaggregated by origin countries and HS10 product categories. We use monthly data from 2017 to 2019, which allows us to closely track the impact of tariffs on US imports, but also the reallocation of these imports toward other countries, as US consumers face higher prices from China. To this end, we take advantage of the data provided by Fajgelbaum et al. (2021), which provides the complete description of the tariff increases imposed by the US administration: this allows us to identify the affected HS10 products, the day of application of the policy, and the amount of the increase in tariff.

From this comprehensive dataset of US imports, we eliminate products or countries that could constitute a challenge to the clean identification of the effects of the tariff hikes. First, we eliminate product codes that were subject to tariff increase applied to many countries. For instance, steel and aluminium products, washer and dryers were also subjects to tariff hikes during this time. Because these policies applied to many countries, and not only China, we decide to remove them from our analysis. As a result, only products that saw a tariff increase on Chinese exports or did not see a tariff increase at all are included in the dataset. Moreover, we eliminate product codes for which there is no Chinese export to the US in 2017. Finally, to quantify the reallocation effects of the rise of US tariffs, we restrict our attention to the top 50 exporters to the United States, excluding oil-exporting countries.[5] This allow us to limit the role of small exporting countries driving our results.

**Specification** Our identification strategy compares the evolution of trade flows of products that were affected by the tariff hikes relative to unaffected product categories. Given the level of disaggregation of our data, our dependent variable $\ln export_{okmy}$ is the logarithm of the amount of export from origin country $o$, HS10 product category $k$, month $m$ and year $y$. We define $T_{kmy}$ as a post-treatment dummy for product categories that were affected by the policy.

Given this identification strategy, we include origin-product dummies to control for levels, month-year dummies to control for aggregate trends in exports and product-month dummies to control for product-specific seasonality. The resulting specification is the following:

$$\ln export_{okmy} = \beta_o T_{kmy} + \gamma_{ok} + \gamma_{my} + \gamma_{km} + \varepsilon_{okmy} \tag{1}$$

where $\gamma$ are the different sets of fixed effects. Importantly, we will allow the effect of the treatment $\beta_o$ to vary by origin countries, in order to capture the heterogeneous effects of the policy across countries.

The resulting estimates of $\beta_o$ can be interpreted as an average treatment effect across all product categories that saw an increase in tariffs. However, there are two reasons why we might find heterogeneous effects across categories. First, the increase in tariffs is not the same for all products. Second, the change in export values is directly related to the market share of China on this market: tariff increase in a

[5] We follow Fajgelbaum, Goldberg, Kennedy, Khandelwal, and Taglioni (2021) and eliminate Algeria, Angola, Irak, Kuwait, Libya, Nigeria, Norway, Trinidad and Tobago, Saudi Arabia, United Arab Emirates and Venezuela from the analysis.

Table 2: Effects of tariffs on Chinese exports to the US

| Dependent Variable | (1) log $V_{Ckt}$ | (2) log $Q_{Ckt}$ | (3) log $V_{Ckt}$ |
|---|---|---|---|
| Treatment × Post | $-0.22^{***}$ | $-0.26^{***}$ | $-0.31^{***}$ |
| | (0.01) | (0.02) | (0.02) |
| Years included | 2017 – 2019 | 2017 – 2019 | 2017, 2019 |
| N | 292 862 | 258 459 | 167 067 |
| $R^2$ | 0.942 | 0.950 | 0.973 |

All specifications include hs10 × month and year × month fixed effects. Standard errors are clustered at the hs10 level.
* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

market where China has a very large market share will generate a larger change in the export values of other countries. For instance in a CES world, the resulting change in trade flows from an increase in Chinese price can be written

$$\mathrm{d}\ln s_{ok} = \frac{\partial \ln s_{ok}}{\partial \ln p_{Ck}} \mathrm{d}\ln p_{Ck} = \sigma s_{Ck} \, \mathrm{d}\ln p_{Ck}$$

where $s_{Ck}$ is the market share from China in product $k$ and $\sigma$ is the elasticity of substitution. Therefore, in an attempt to estimate the origin-specific elasticity of substitution, we run a second specification in which we scale the treatment by the actual change in tariffs, and the market share of China in this product category. To avoid any endogeneity issue, we compute the Chinese market share based on the year 2017 alone, aggregating all monthly export flows to obtain the market share for the year. The specification becomes

$$\ln export_{okmy} = \sigma_o \, d\tau_{kmy} \times s_{Ck} + \gamma_{ok} + \gamma_{my} + \gamma_{km} + \varepsilon_{okmy} \tag{2}$$

where $d\tau_{kmy}$ is the change in tariffs relative to the year 2017.

**Result** We start by running the regression using imports from China only, to show the negative impact of the tariff increases on the import flows directly targeted by the policy. In table 2, we show the negative impact of tariffs using three different specifications: specifications (1) and (2) estimates the impact of the policy using equation (1), respectively on import values and import quantities.[6] Specification (3) looks at the impact on import values but restrict the data to the years 2017 and 2019, which allows us to eliminate the possible anticipatory or lagged effects that are likely to occur in the months surrounding the implementation of the policy in the summer of 2018.

All three specifications highlight the negative impact of the tariffs on Chinese imports. Using our preferred specification, specification (3), we find a 31 percent reduction in the imports of products that were targeted by tariffs, relative to products that did not experience a rise in tariffs. Specifications (1) and (2) show similar results which indicate little adjustments in prices.

Having demonstrated the negative impact on Chinese imports, we now look at the impact on other countries, to identify which countries picked up the imports that were not originating from China anymore. To do so, we run equations (1) and (2) including import flows from all top 50 exporters to the US. We use the import values as dependent variable and restrict our analysis to the years 2017 and 2019 to avoid the dynamic effects that take place in 2018 during the months surrounding the implementation of

[6] The regression with quantities has fewer observations due to the presence of missing quantities in the data.

the policy. Table 3 report the results for all 50 exporters:[7] specification (1) estimates the reduced form impact when the treatment is a dummy for products that were impacted by the tariffs, while specification (2) estimates the cross-elasticity by normalizing the treatment as in equation (2) using the change in the tariff rate and the market share of Chinese imports for this product category.

Table 3 identifies the winners and losers from the tariffs hikes. Naturally, we find that China sees a sharp reduction in their exports to the US, similar to the previous table. We also find a strong reduction in imports from Hong Kong, probably due to the nature of the global value chains for products manufactured in China. In terms of winners, we find that the countries with the larger cross-elasticities are Vietnam (5.32), India (3.55) and Turkey (2.84). Countries like Thailand, the Dominican Republic, Malaysia or Indonesia also appear to benefit from this increasing US trade protection toward Chinese goods.

In order to show that these cross-elasticities are related to the proximity to China in the product space, we construct a country-level measure of price distance to China that we correlate to the cross-elasticity. For each origin country $o$, we construct the average distance in price from Chinese goods by averaging across products the absolute log price difference to China:

$$Dist_{oC} = \frac{1}{N_k} \sum_k |\ln p_{ok} - \ln p_{Ck}|. \tag{3}$$

In this distance, the products $k$ are the hs10 product categories that are subject to tariffs.

Using this distance to quantify the proximity in the price space, we display in figure 3 the correlation with the cross elasticity estimated above. This scatter plot shows that countries that are located closer to Chinese prices also have a larger cross-price elasticity with China. For instance, Vietnam and India, which we estimated to have the larger cross-price elasticity with China, are also the countries that are the most similar to China in terms of prices. Combining all countries, this negative relationship between cross-price elasticity and distance is statistically significant, as displayed by the red line.[8]

This result reflects the importance of accounting for realistic patterns of substitution across countries. Assuming a CES demand function implies that the trade war would have generated uniform gains in market shares across third countries. Instead, we find that countries that are similar to China in the product space benefit much more from the reallocation of Chinese exports induced by the trade war. In the next section, we provide a simple estimation framework that can capture these important aspects of the substitution across countries.

# 3 Theory and Estimation

In this section, we derive a gravity equation of international trade which is (i) micro-founded, (ii) captures rich substitution patterns between source countries and (iii) can be easily estimated by 2SLS.

## 3.1 Mixed Preferences and Aggregate Demand

We first present the demand system which underlies the gravity equation. The global economy is a collection of markets $m$, each populated with a unit mass of heterogeneous consumers. A market is a unique combination of a destination $d$, an HS6 product $k$ and a year $t$. Each consumer $i$ has CES

[7] We omit the standard errors for conciseness, but include them in the version of the table in the appendix B.

[8] The degree of opacity of the blue circles describes the precision of the cross-price elasticity estimates. The regression, with a coefficient and p-value of respectively -1.44 and 0.001, is also weighted by the inverse of the standard errors of these estimates. We remove China, Taiwan and Hong Kong from the figure because these observations are directly and negatively affected by the increase in tariffs.

Table 3: Effects of tariffs on exports to the US

| | | (1) | (2) |
|---|---|---|---|
| Treatment | × ARGENTINA | 0.01 | 0.03 |
| | × AUSTRALIA | 0.04 | 0.72 |
| | × AUSTRIA | 0.12*** | 2.02*** |
| | × BELGIUM | 0.05** | 0.52 |
| | × BRAZIL | 0.04 | 1.30** |
| | × CANADA | −0.01 | −0.20 |
| | × CHILE | −0.01 | −0.41 |
| | × CHINA | −0.30*** | −6.39*** |
| | × COLOMBIA | 0.04 | 0.54 |
| | × COSTA RICA | 0.02 | −0.85 |
| | × CZECH REPUBLIC | 0.06** | 0.80 |
| | × DENMARK | 0.03 | −0.12 |
| | × DOMINICAN REPUBLIC | 0.07* | 2.15** |
| | × ECUADOR | 0.06 | 0.95 |
| | × FINLAND | 0.08*** | 0.77 |
| | × FRANCE | 0.01 | 0.24 |
| | × GERMANY | 0.04*** | 0.87*** |
| | × GUATEMALA | 0.01 | 0.35 |
| | × HONDURAS | −0.01 | 0.55 |
| | × HONG KONG | −0.62*** | −10.97*** |
| | × HUNGARY | 0.05 | 0.94 |
| | × INDIA | 0.20*** | 3.55*** |
| | × INDONESIA | 0.07** | 1.88*** |
| | × IRELAND | 0.07* | 1.85** |
| | × ISRAEL | 0.01 | −1.00 |
| | × ITALY | 0.06*** | 0.87*** |
| | × JAPAN | 0.01 | −0.23 |
| | × KOREA, SOUTH | 0.07*** | 0.65 |
| | × MALAYSIA | 0.13*** | 2.05*** |
| | × MEXICO | 0.08*** | 1.53*** |
| | × NETHERLANDS | 0.04* | 0.47 |
| | × NEW ZEALAND | −0.02 | −0.18 |
| | × NICARAGUA | 0.05 | −1.74 |
| | × PERU | −0.01 | −0.85 |
| | × PHILIPPINES | 0.06* | 1.35** |
| | × POLAND | 0.17*** | 2.35*** |
| | × PORTUGAL | 0.08** | 1.39* |
| | × ROMANIA | 0.09** | 1.50 |
| | × RUSSIA | 0.12** | 1.33 |
| | × SINGAPORE | −0.01 | −0.88 |
| | × SLOVAKIA | 0.04 | 0.44 |
| | × SOUTH AFRICA | −0.02 | −1.36 |
| | × SPAIN | 0.06*** | 0.75 |
| | × SWEDEN | 0.09*** | 1.35*** |
| | × SWITZERLAND | 0.00 | −0.67 |
| | × TAIWAN | 0.02 | −0.02 |
| | × THAILAND | 0.09*** | 1.84*** |
| | × TURKEY | 0.16*** | 2.84*** |
| | × UNITED KINGDOM | 0.04*** | 0.46 |
| | × VIETNAM | 0.29*** | 5.32*** |
| Num.Obs. | | 1 948 501 | 1 948 501 |
| R2 | | 0.878 | 0.878 |

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

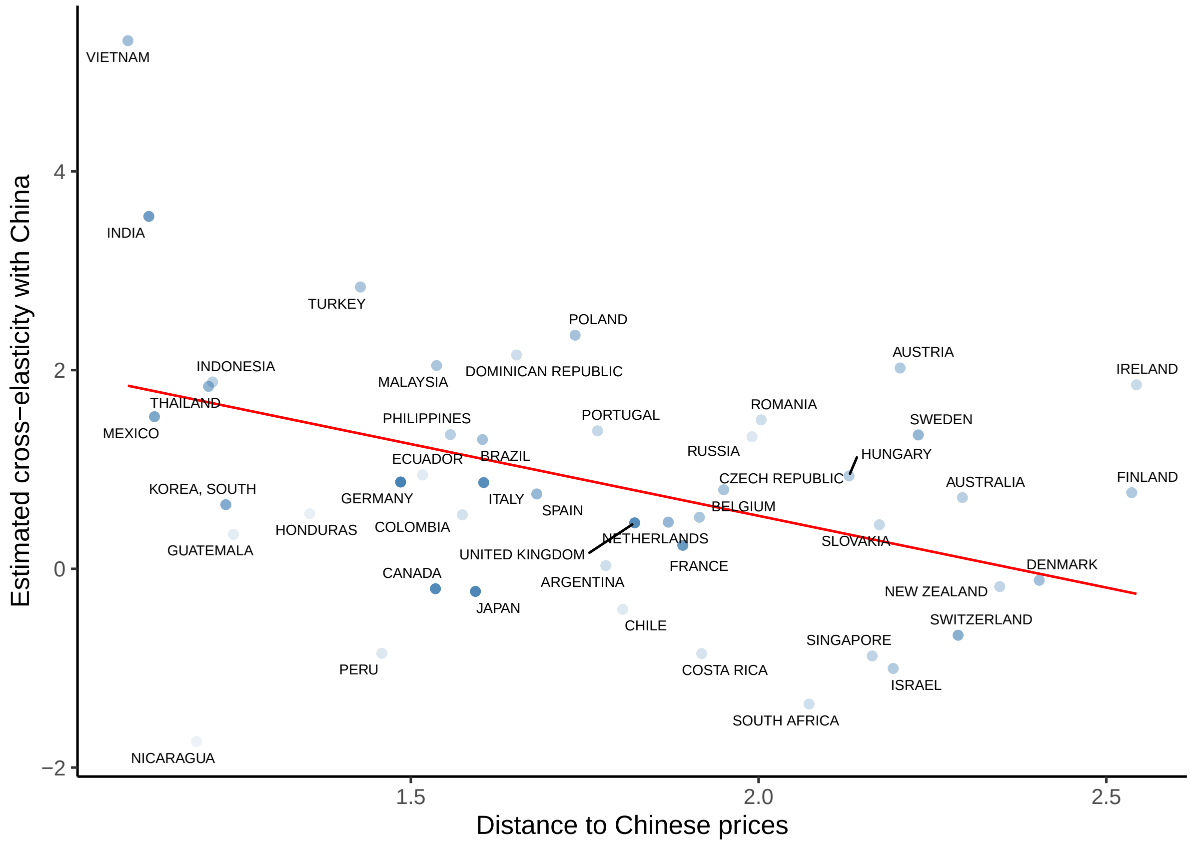


FIGURE 3: Scatterplot between cross-price elasticity with China and distance to Chinese price

preferences over the set $\mathcal{O}_m$ of varieties supplied by different source countries $o$:

$$U_{im} = \left( \sum_{o \in \mathcal{O}_m} \exp(X'_{om}\beta + \xi_{om})^{\frac{1}{\sigma_i}} \, q_{io}^{\frac{\sigma_i}{\sigma_i - 1}} \right)^{\frac{\sigma_i - 1}{\sigma_i}}.$$

$X_{om}$ is a vector of observable characteristics and $\xi_{om}$ is the "quality" of variety $o$. This measure of quality contains all characteristics that enter the utility function – both tangible and intangible – which are observable to consumers and unobservable to the econometrician. $q_{io}$ is the physical quantity of good from country $o$ purchased by consumer $i$. This utility function features mixed-preferences since $\sigma_i$ is the consumer-specific elasticity of substitution across varieties. $\beta$ is a $(n_X \times 1)$ vector of parameters driving the relative preference for the different observable characteristics and we assume that it is constant. However, we show in appendix D that the estimation of the model can be extended to let $\beta$ vary across consumers. Let $\alpha_i \equiv \sigma_i - 1$. The share of variety $o$ in consumer $i$'s expenditure is

$$s_{io} = \frac{\exp\left(-\alpha_i \ln p_{om} + X'_{om}\beta + \xi_{om}\right)}{\sum_{o' \in \mathcal{O}_m} \exp\left(-\alpha_i \ln p_{om} + X'_{om}\beta + \xi_{om}\right)}, \tag{4}$$

with $p_{om}$ the price faced by consumers in market $m$ for goods for origin $o$. $p_{om}$ includes tariffs and transportation costs. The aggregate revenue market share of variety $o$ in market $m$ is

$$s_{om} = \frac{\int s_{io} e_{im} \, di}{\int e_{im} \, di} = \int s_{io} \omega_{im} di, \tag{5}$$

with $e_{im}$ the total budget that consumer $i$ spends on varieties from market $m$ and $\omega_{im} \equiv \frac{e_{im}}{\int e_{im}\, di}$ the share of consumer $i$ in the aggregate sales of market $m$. Assuming that all consumers in a given destination have identical Cobb-Douglas preferences across HS6 products, then $e_{im} \propto y_i$ and $\omega_{im} = \frac{y_i}{\int y_i\, di} = \omega_{id}$: the share of consumer $i$ in the sales of any market $m$ is equal to $i$'s share of the national income.

**Substitution patterns** Introducing mixed-CES preferences in a gravity equation delivers more realistic substitution patterns relative to a simple CES case. To see this, consider the expression of the cross-price elasticity of variety $o' \neq o$ with respect to the price of $o$.

$$\frac{\partial \ln s_{o'}}{\partial \ln p_o} = \frac{1}{s'_o} \int \alpha_i s_{io'} s_{io} \omega_{id}\, di. \tag{6}$$

All things equal, this cross-elasticity is large if varieties $o$ and $o'$ serve similar consumers, which happens if $o$ and $o'$ share similar product characteristics. For instance, varieties with low prices will be purchased by price-sensitive consumers. As a result, an increase in the price of one of these varieties will induce price-sensitive consumers to reallocate their consumption, benefiting predominantly other affordable varieties. As such, mixed-CES preferences relax the IIA condition: when $c$ increases its price and loses market shares, not all competing varieties benefit in the same proportion. Instead, varieties closer to $c$ in the product space experience a larger increase in their market shares. By comparison, with simple CES (i.e. when $\alpha_i = \alpha$ for all consumers) the cross-elasticity collapses to:

$$\left.\frac{\partial \ln s_{o'}}{\partial \ln p_o}\right|_{\alpha_i = \alpha\, \forall i} = \alpha s_o,$$

in which case the IIA restriction holds because all varieties $o' \neq o$ are equally impacted by the change in the price of $o$.

**Random Coefficients** To discipline the distribution of preferences across consumers, we assume that $\alpha_i$ verifies:

$$\alpha_i = \alpha + \pi\mu_d + \lambda V_d^{1/2}\nu_i + \gamma\varepsilon_i, \tag{7}$$

with $\varepsilon_i$ a random shock and $\nu_i$ the standardized demographic shock: for a consumer $i$ from destination $d$, $\ln y_i = \mu_d + V_d^{1/2}\nu_i$. We expect $\pi$ and $\lambda$ to be negative: richer consumers are less price elastic. (7) allows for $\lambda \neq \pi$. Economically speaking, this means that idiosyncratic income shocks $V_d^{1/2}\nu_i$ may impact differently the price elasticity $\alpha_i$ than shocks to the mean income $\mu_d$. It could be for instance that the price-elasticity of a consumer does not depend on her absolute level of income but rather on her distance to the average national income, in which case $\lambda \neq \pi = 0$.

Hereafter we refer to $\theta \equiv \{\lambda, \gamma\}$ as the "non-linear parameters". If $\theta = 0$, there is no dispersion across consumers within a market, and we are back to simple CES preferences. Importantly, the estimation procedure does not impose parametric assumptions on the distribution of these random coefficients. This is in contrast with traditional estimation methods that require to integrate the distribution of random coefficients to evaluate an objective function, and therefore need to assume a specific parametric distribution for these coefficients. By contrast, we only assume that the mean and variance of the log-income $\ln y_i$ – respectively $\mu_d$ and $V_d$ – vary by destination and are known to the econometrician.[9] We also normalize $\varepsilon_i$ to have zero mean and unit variance. Finally, we assume that $\nu_i$ and $\varepsilon_i$ are i.i.d. across consumers.

[9] In appendix D, we show how to extend the estimation to a model with multiple demographics.

### 3.2 Estimation

The presence of heterogeneity in preferences poses challenges for the estimation of the model. It requires the econometrician to evaluate the integral featured in equation (5) which complicates the estimation of the model. By contrast, the model without random coefficient – with $\theta = 0$ – generates a log-linear demand equation that can be estimated with standard linear regressions. Therefore, Salanié and Wolak (2022) proposes to estimate an approximation of the true model around $\theta = 0$, in the context of mixed logit preferences. In this paper, we extend their approach to mixed CES preferences with demographic differences across markets.[10] Both extensions fit the specific needs of international trade data. First, mixed CES preferences imply that the consumer-level market share depends on log-prices, as opposed to prices with mixed logit preferences.[11] Since prices tend to be noisy in trade data, using log-prices help mitigate the role of measurement error and outliers. Second, considering demographics makes it possible to leverage the variation across destination markets from international trade data to identify the impact of demographics on consumer preferences and substitution patterns.

**The "small-$\sigma$" approach** To understand the "small-$\sigma$" approach developed in Salanié and Wolak (2022), it is useful to first sketch the way mixed demand systems are traditionally estimated. Following Berry et al. 1995 (hereafter, BLP), the identification usually comes from orthogonality conditions between structural errors $\xi$ and instruments $Z$:

$$E(\xi Z) = 0. \tag{8}$$

Let $\boldsymbol{\xi}_m(\theta)$ be the vector of "structural errors" of the model. For any candidate parameter value $\theta$, $\boldsymbol{\xi}_m(\theta)$ is such that the market shares are equalized in the data and the model :

$$\boldsymbol{\xi}_m(\theta): \quad \boldsymbol{\xi} \text{ s.t. } s_{om}(\boldsymbol{\xi}, \theta) = S_{om}, \tag{9}$$

with $S_{om}$ the observed market share. $\hat{\theta}$ is obtained by iterating over $\theta$ until minimizing the sample distance between $E(\xi(\theta)Z)$ and 0. This can be computationally challenging, notably because each interation involves inverting market shares and estimating $(\beta_1, \beta_2)$ in order to obtain $\boldsymbol{\xi}_m(\theta)$. By contrast, the solution proposed by Salanié and Wolak (2022) is to estimate the model through a linear regression equation obtained from a Taylor expansion. Let $\tilde{\theta}$ be the vector of non-linear parameters re-scaled by a scalar $\sigma$ : $\theta = \sigma\tilde{\theta}$. They perform a Taylor expansion of the structural errors $\boldsymbol{\xi}(\theta)$ around $\sigma = 0$:

$$\xi(\theta) = \xi(\sigma\tilde{\theta}) = \xi(0) + \left.\frac{\partial \xi(\sigma\tilde{\theta})}{\partial \sigma}\right|_{\sigma=0} \sigma + \left.\frac{\partial^2 \xi(\sigma\tilde{\theta})}{\partial \sigma^2}\right|_{\sigma=0} \frac{\sigma^2}{2} + O(\sigma^2). \tag{10}$$

**The Regression Equation** After some algebra, (10) delivers the following regression equation:

$$\ln s_{om} = X'_{om}\beta - \alpha \ln p_{om} - \pi\mu_d \ln p_{om} + \gamma^2 K_{om} + \lambda^2 V_d K_{om} + \xi_{om} + (\alpha + \pi\mu_d) \ln P_m + O(\sigma^2) \tag{11}$$

$$\text{with} \begin{cases} K_{om} \equiv \frac{1}{2}(\ln p_{om} - \overline{\ln p}_m)^2 \\ \overline{\ln p}_m \equiv \sum_o s_{om} \ln p_{om}. \end{cases}$$

[10] Breinlich, Fadinger, Nocke, and Schutz (2020) also derive a log-linear gravity equation through a first order approximation. However, while we are interested in extending the standard gravity equation to mixed preferences, they focus on a generalization to oligopolistic competition.

[11] Mixed CES preferences are getting increasingly adopted in empirical models (Adao et al., 2017; Dubé et al., 2021; Head and Mayer, 2021). See Birchall and Verboven (2022) for a comparison of mixed logit and mixed CES preferences in terms of own and cross-price elasticity.

and $P_m$ the price index.[12] $K_{om}$ is what Salanié and Wolak (2022) refer to as an "artificial regressor". $K_{om}$ is (half) the squared distance of variety $o$ to the weighted average log price in the market. According to equation (11), $\ln s_{om}$ is increasing in this squared distance. This means that the second order approximation preserves a major feature of the mixed CES model: varieties that are more isolated in the price space have a larger market share, all things equal. The only difference between the exact model and the approximated model is that the approximation only involves the distance to the mean product characteristic. By contrast, in the exact model, sales depend on the distance to each competing variety.

Demographic moments enter two terms of equation (11). First, $\ln s_{om}$ depends on the interaction between mean income $\mu_d$ and $K_{om}$. We expect market shares to be less elastic to prices in richer markets, which corresponds to the case $\pi > 0$. Second, demographics enter the equation through the interaction term $V_d K_{om}$: the dispersion of log-income shapes the relationship between market shares and price differentiation. Interestingly, this interaction always has a positive impact on market shares since $\gamma^2 > 0$. In words, it means that the return to product differentiation is larger in markets with more income dispersion. The intuition is the following: less income dispersion means less preference dispersion. As a result, there is less scope to gain market shares by adopting a niche position in the product space.

Importantly, equation (11) extends naturally to multiple non-linear characteristics and demographics. The expression and derivation of (11) in the general case can be found in appendix (D).

**Bringing the Linearized Model to the Data** Equation (11) can be estimated by 2SLS. It is necessary to instrument $\ln p$ (as well as $\mu_d \ln p$) for obvious simultaneity reasons. Note that it is also necessary to instrument $K_{om}$ because it depends on $\ln p$ and $s$, which are both endogenous variables. If $Z^{lnp}$ is a set of instrumental variables for prices and $\widehat{\ln p}$ and $\widehat{s}$ are respectively the predictions of $\ln p$ and $s$ based on $Z^{lnp}$, then a natural instrument for $K_{om}$ is

$$\hat{K}_{om} = \frac{1}{2}\left[\widehat{\ln p}_{om} - \sum_{o} \hat{s}_{om}\widehat{\ln p}_{om}\right]^2$$

Accordingly, one can obtain an instrument for $V_d K_{om}$ by interacting $V_d$ with $\hat{K}_{om}$.

Hereafter, we use the acronym "FRAC" to refer to the 2SLS estimator of equation (11). This acronym introduced by Salanié and Wolak (2022) stands for Fast, Robust and Approximately Correct. "Fast" refers to the fact that 2SLS is faster to compute than non-linear GMM techniques usually required to estimate mixed preferences. "Robust" emphasizes the fact that equation (11) does not require distributional assumptions on the random coefficients, beside specifying their mean and variance. "Approximately Correct" has to do with the fact that (11) is an approximation of the true model. At the end of this section, we use simulated data to show that in spite of this approximation, FRAC adequately estimates substitution patterns.

## 3.3 Counterfactual Market Shares in the Linearized Model

Once estimated $\theta$, practitioners will likely be interested in computing cross (-price) elasticities between varieties from different source countries. One approach is to plug $\hat{\theta}$ into the full model and evaluate equation (6). We propose a different approach which consists in computing the cross-elasticticies from the

[12] In the exact model, there are as many CES price indices as consumers. However, once linearized, the aggregate demand equation is consistent with a representative consumer having CES preferences and whose price index verifies

$$(\alpha + \pi\mu_d)\ln P_m = -\sum_{o\in\Omega_m} \exp\left(X'_{om}\beta - \alpha \ln p_{om} - \pi\mu_d \ln p_{om} + \gamma^2 K_{om} + \pi V_d K_{om} + \xi_{om} + O(\sigma^2)\right)$$

linearized model. This alternative approach presents two advantages: (i) it does not require to specify the distribution of the random coefficient shocks $\nu_i$ and $\varepsilon_i$ and (ii) it does not require numerical integration techniques because it delivers a closed form expression for the elasticity.

In order to derive the expression of the cross price elasticity $\frac{\partial \ln s_o}{\partial \ln p_c}$ in the linearized model, let us first consider the expression of the market share. Slightly re-arranging (11) gives

$$s_{om} = \frac{\exp(X'_{om}\beta - \bar{\alpha}_d \ln p_{om} + \Sigma_d K_{om} + \xi_{om})}{\sum_k \exp(X'_{om}\beta - \bar{\alpha}_d \ln p_{km} + \Sigma_d K_{km} + \xi_{km})}, \quad \text{with} \begin{pmatrix} \bar{\alpha}_d \\ \Sigma_d \end{pmatrix} \equiv \begin{pmatrix} \alpha + \pi\mu_d \\ \gamma^2 + \lambda^2 V_d \end{pmatrix}. \tag{12}$$

As we show in appendix (F), getting an expression for the price elasticity from (12) is straightforward:

$$\frac{\partial \ln s_{om}}{\partial \ln p_{cm}} = -\Sigma_d \frac{\partial \overline{\ln p}_m}{\partial \ln p_{cm}} \widetilde{\ln p}_{om} + \Sigma_d (\mathbb{1}_{o=c} - s_c) \widetilde{\ln p}_{cm} - \bar{\alpha}_d (\mathbb{1}_{o=c} - s_c) \tag{13}$$

with $\widetilde{\ln p}_{om} \equiv \ln p_{om} - \sum_k s_{om} \ln p_{om}$ the deviation of $o$ to the weighted average log price. Equation (13) reveals that the way varieties respond to a change in the average price depends on their position with respect to this average price. For varieties whose prices are above the average ($\widetilde{\ln p}_{om} > 0$), an increase in the average decreases their market shares because it reduces their price differentiation. Conversely, a variety whose price stands below the average will benefit from an increase in the average price.

Although equation (13) is analytically informative, it raises a problem: $\frac{\partial \ln s_{om}}{\partial \ln p_{cm}}$ depends on the elasticity of the average price $\frac{\partial \overline{\ln p}_m}{\partial \ln p_{cm}}$ which depends itself on the vector of elasticities $\frac{\partial \ln s_{om}}{\partial \ln p_{cm}}$. Therefore, getting an equilibrium expression for the market share elasticities requires finding a fixed point for $\frac{\partial \overline{\ln p}_m}{\partial \ln p_{cm}}$. As we demonstrate in appendix F, this fixed point has a closed form solution which can be easily computed from the data:

$$\frac{\partial \overline{\ln p}_m}{\partial \ln p_{cm}} = \frac{s_c \left\{ \Sigma_d \widetilde{\ln p}^2_{cm} - \bar{\alpha}_d \widetilde{\ln p}_{cm} + 1 \right\}}{1 + \Sigma_d \mathbb{V}^{\ln p}_m}, \tag{14}$$

with $\mathbb{V}^{\ln p}_m = \sum_o s_o \left( \ln p_{om} - \overline{\ln p_m} \right)^2$, the market share weighted variance of log prices.

In this section we described to "small-$\sigma$" approach to estimate and implement counterfactual analysis in the context of mixed-CES preferences. In appendix E, we provide simulation-based evidence of the ability of the "small-$\sigma$" to estimate the cross elasticities.

# 4 Empirical Implementation

We now implement the estimation strategy described in the previous section, which allows us to estimate a gravity equation with rich patterns of substitution across countries. We start by describing the data employed to implement our estimation, before showing our results.

## 4.1 Data

The estimation of our gravity model does not require more data than usually employed in this kind of exercise. It requires data on bilateral trade flows across countries, traditional bilateral gravity variables such as distance, common languages or colonial history, as well as information regarding trade policy such as tariffs, or preferential trade agreement. In the context of this paper, we augment this set of data with moments of the income distribution of the destination countries, to improve the identification of the random coefficients, as well as information regarding the domestic good. Although these two pieces of information are not required, they are useful to document the implementation of the model.

First, we use trade data from BACI that is available at the bilateral level and disaggregated by year and HS6 product categories. We obtain export values in US dollars and export quantities measured in weight. From these exported values and quantities, we construct unit values that we use as proxy for prices in our analysis. We augment these bilateral trade flows with tariffs measures obtained from the MacMap dataset. This dataset, developed with the CEPII, provides measures of applied tariff duties at the bilateral level and disaggregated at the HS6 level. In the context of the estimation of the import demand system, we will use these tariffs as instruments for the prices of exported goods: because they increase the final price of these exports, while being uncorrelated with the demand shifter in the import demand function, this measure is a perfect instrument to consistently estimate the price-elasticity of imports. We also use standard gravity variables, such as geographical distance, common language, colonial links, in our main specification. These variables are obtained from the gravity database, provided by the CEPII.

In addition to these standard datasets, some of the specifications require demographics information about the destination markets, and data regarding the domestic good. We obtain the average income per capita and the Gini index from the World Development Indicators database from the World Bank. This information allows us to obtain the mean and standard deviation of the income distribution in each destination markets. Finally, to obtain information on the domestic good, we rely on trade data and the World Input Output Database (WIOD). We first need the share of the domestic good at the destination-product level, that we obtain from the WIOD.[13] Moreover, we compute the price of the domestic good from the export data of that country. We observe the FOB prices of the good exported from each destination markets at the HS6 level, which we use as proxy for the price of the domestic good in that market.[14]

To facilitate the estimation of the model, we reduce the coverage of this dataset: we restrict our sample to the 39 destinations that are included in the World Input Output Database (WIOD).[15] We focus on this limited set of destinations because we use domestic information from the WIOD that only include these destination markets. Moreover, we restrict our analysis to the exporting performance of the top 100 exporting countries to this set of destinations, to eliminate countries with missing or imprecise observations. Finally, we estimate the model using the years 2001, 2004, 2007 and 2010: these are the years for which tariffs data are available from the MacMap database. Because our estimation mostly relies on cross-sectional variations, four years of data provide sufficient variation to precisely identify the parameters of the model.

## 4.2 Econometric Specification

The main econometric equation comes directly from the model derived in the previous section. From (11), we can write the trade flow from origin country $o$ to destination $d$ in product $k$ at time $t$ as

$$\begin{aligned}\ln export_{odkt} =& X'_{odkt}\beta - \alpha \ln p_{odkt} + \pi\, \mu_d \times \ln p_{odkt} + \gamma^2 K_{odkt} + \lambda^2\, V_d \times K_{odkt} \\ &+ \delta_{okt} + \delta_{dkt} + \varepsilon_{odkt}.\end{aligned}$$

This gravity regression contains six bilateral variables $X_{odkt}$ that are traditionally included in the gravity models. In our specifications, we use the physical distance, whether the two countries share a language, a currency, a border, the same main religion and a history of colonial links. In additional to these gravity variables, we have variables that identify our demand system. The price of the variety (inclusive of tariffs and transportation costs) – $p_{odkt}$ – identifies the average price-elasticity $\alpha$, while the artificial regressor

[13]The WIOD only provides the share of domestic consumption in total consumption at the CPA level, which is more aggregated than the HS6 product categories. We use a conversion table to assign HS6 products to CPA categories.

[14]More details on the construction of these variables appear in appendix A.

[15]The database includes 42 countries but we exclude Cyprus, Malta and Taiwan due to data constraints.

$K_{odkt}$ captures the heterogeneity in price-elasticity $\gamma$, that generates rich substitution patterns. Finally, the interaction of the price with the average income $\mu_d$ and the interaction between the variance of the income $V_d$ and the artificial regressor $K_{odkt}$ capture the role of income in shaping heterogeneity in the price-elasticity. To faciliate the interpretation of the coefficients, both $\mu_d$ and $V_d$ are deviated from their respective sample means.

We add two sets of fixed effects to this equation. First, we include a market-level fixed effect $\delta_{dkt}$ that controls for the price index that is specific to the market. As a result, we do not need to normalize our regression by the the market share of an outside good, or to compute the price index based on the set of varieties available to the consumers. Second, we include a producer fixed effect $\delta_{okt}$ to control for supply shocks that would apply to many destination markets. As a result, our identification relies on variations at the bilateral level, where producers perform differently in different destinations based on their relative position on the market.

Despite this set of fixed effects, the estimating equation is subject to endogeneity concerns due to the presence of prices and artificial regressors. These regressors are constructed using market shares and prices, which are both endogenous variables, and therefore need to be instrumented. To deal with the endogeneity of prices, we use the tariff rate at the bilateral level that can be arguably used as a cost shifter: it is likely to shift the final price of a good while, at the same time, not being correlated with the demand residual in the gravity regression. To instrument the artificial regressor, we first estimate a gravity model using the six exogenous gravity variables, the tariff rates and the two sets of fixed effects. From this gravity model, we obtain a prediction of the market share of each variety, which allows us to construct an exogenous version of the artificial regressor. Formally, we have

$$\hat{K}_{odkt} = \frac{1}{2}(\ln \hat{p}_{odkt} - \sum_{o'} \hat{s}_{o'dkt} \ln \hat{p}_{o'dkt})^2$$

in which $\hat{p}_{odkt}$ and $\hat{s}_{o'dkt}$ are predicted version of the market share and prices using exogenous variables instead. We then use this exogenous artificial regressor as instrument for the true artificial regressor.

### 4.3 Results

We present the results of our estimation in table 4. In columns (1) et (2), we run a simple structural gravity regression that includes the price of the trade flows and the set of bilateral gravity variables. In column (1), prices are not instrumented, which explains the small coefficient on prices at -0.08. In column (2), we show that the instrumentation of prices using tariffs generates a more realistic price elasticity around -2.29, which validates the instrumental strategy. In addition to these consistent results regarding the effect of prices, we also find expected results on the gravity variables. The geographical distance has a negative impact on trade flows, but each variable describing a common feature between countries generates more trade between these countries.

In columns (3) to (5) of table 4, we introduce the artificial regressors that identify the heterogeneity in consumer preferences, which is characteristic of the mixed-CES model. In columns (3) and (4), we add the artificial regressor on prices, to capture the role of price differentiation. Even though this regressor is not instrumented in column (3), we still find a positive coefficient, which shows that varieties with a price that differs from the average price perform better in export markets. This benefit from price differentiation is confirmed in column (4) in which the coefficient on the artificial regressor $K_{\ln p}$ is even larger. This larger coefficient from the instrumentation is expected since the instrument only identifies the patterns of substitution through supply shocks. By contrast, the OLS coefficient in column (4) is also identified from demand shocks, that are likely to be positively correlated between similar varieties. As a result, the correct identification through supply shocks better captures the stronger substitution patterns between varieties that are closer in the price space. This gain from price differentiation is

Table 4: FRAC regressions – estimation results

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| $\ln p$ | $-0.08^{***}$ | $-2.29^{***}$ | $-2.50^{***}$ | $-2.68^{***}$ | $-2.76^{***}$ |
| | (0.00) | (0.07) | (0.08) | (0.08) | (0.09) |
| $K_{\ln p}$ | | | $0.30^{***}$ | $0.54^{***}$ | $0.60^{***}$ |
| | | | (0.01) | (0.02) | (0.02) |
| $\ln p \times \mu_Y$ | | | | | $0.30^{***}$ |
| | | | | | (0.04) |
| $K_{\ln p} \times V_Y$ | | | | | $0.17^{***}$ |
| | | | | | (0.06) |
| log distance | $-0.91^{***}$ | $-0.46^{***}$ | $-0.48^{***}$ | $-0.49^{***}$ | $-0.46^{***}$ |
| | (0.00) | (0.01) | (0.01) | (0.01) | (0.02) |
| Common currency | 0.18*** | 0.15*** | 0.13*** | 0.11*** | 0.11*** |
| | (0.00) | (0.01) | (0.01) | (0.01) | (0.01) |
| Common border | 0.67*** | 0.57*** | 0.57*** | 0.57*** | 0.56*** |
| | (0.00) | (0.01) | (0.01) | (0.01) | (0.01) |
| Common language | 0.24*** | 0.17*** | 0.19*** | 0.20*** | 0.21*** |
| | (0.00) | (0.01) | (0.01) | (0.01) | (0.01) |
| Common religion | 0.12*** | 0.07*** | 0.07*** | 0.08*** | 0.04*** |
| | (0.01) | (0.01) | (0.01) | (0.01) | (0.01) |
| Colonial link | 0.20*** | 0.06*** | 0.06*** | 0.06*** | 0.06*** |
| | (0.00) | (0.01) | (0.01) | (0.01) | (0.01) |
| $\ln p$ instrumented | N | Y | Y | Y | Y |
| $K_{\ln p}$ instrumented | . | . | N | Y | Y |
| F-stat first stage | . | 6 155 | 5 242 | 2 801 | 1 117 |

*Notes:* $N = 11\,660\,255$. Standard errors between parentheses are clustered at the origin-product and destination-product levels. All regressions include origin-product-year and destination-product-year fixed effects. * p < 0.1, ** p < 0.05, *** p < 0.01.

important economically: a variety that is positioned 0.5 log-point above or below the average price enjoys an additional 7 percent export value ($\frac{1}{2}0.5^2 \times 0.54$) relative to a variety with a price equal to the mean.

In column (5), we introduce demographics to better capture the distribution of heterogeneity across consumers. First of all, we find that richer consumers have a lower price-elasticity given the positive coefficient on the interaction between income and price.[16] Moreover, we also find that destinations with larger income inequality features more dispersion in preferences, given the positive coefficient on the interaction between the variance in income and the artificial regressor. As an illustration of the effect of the income distribution, the United States that record higher income on average and more inequality have an average price-elasticity of -2.56 and a parameter on the artificial regressor of 0.63.

Even though the positive role of the variance is consistent with the effect of average income, it is reassuring to see that the identification of a role for income is not only driven by variation in the price-elasticity across destinations, but also by a stronger role for price differentiation within destinations with more income inequality. This result connects our finding to the micro-foundations of the mixed CES model, which generates complex substitution patterns from heterogeneity in consumer preferences.

[16] This result is reminiscent of Bergstrand (1990)'s evidence that bilateral trade decreases with income per capita differences between the origin and destination countries.

Finally, remember that the role of income was over-identified in a regression that features a role for average income across destinations and for the variance in income within a country: in theory, both the average and variance identifies the same effect of income on the price-elasticity, such that the coefficient associated with the variance should be the square of the coefficient associated with the average. In our regression, we cannot reject this relationship with a coefficient on the variance (0.17) that is much smaller than the coefficient on the average (0.30).

Overall, these results confirm that the data strongly reject the CES assumption that implies symmetric patterns of substitution across firms. We find that the performance of a country in export markets depends on its position in the price distribution: countries that are able to occupy a segment of the market that is less dense perform better than countries located closer to the average exporter. As a result, countries or firms are more affected by a change in the export performance of a country that is located near them.

**Alternative specifications** In addition to the results showing the relevance of the augmented gravity model, we now show alternative specifications to document the versatility of the method. In table 5, we show the results of specifications that include the domestic good as outside good, similar to the current practice in the Industrial Organization literature, and use alternative measures of differentiation.

In columns (1) and (2) of table 5, we incorporate the possibility for the consumer to choose the domestic good as an option. The results in the first two columns of table 5 show that adding this domestic variety has a limited impact on the parameter estimates of the model. We do find a smaller role for vertical differentiation once we account for the domestic variety, as witnessed by the smaller value of the estimates on the artificial regressors. However, this might be due to the fact that we have limited information on the different choices among the domestic varieties, which might mis-measure the extent of the vertical differentiation between these varieties. Overall, these results are reassuring on the ability of the model to capture realistic substitution patterns when using trade data only, by focusing essentially on the import demand system.

In column (3), we estimate the model without prices. This specification mimics traditional gravity regressions which aim at estimating the tariff elasticity of trade flows rather than the price-elasticity. In this scenario, we directly use the bilateral tariffs in the specification as a measure of trade barriers and use the gdp per capita to capture possible patterns of substitution across countries: we might expect countries with similar level of developments to have a higher degree of substitution between them. The results of column (3) confirm this intuition: the artificial regressor based on the gdp per capita, rather than price, is highly significant. This implies that countries with similar gdp per capita are closer substitutes. This specification shows that the FRAC method can be applied to trade data without information on prices and yet capture realistic substitution patterns.

Finally, we show in column (4) that we can also include additional artificial regressor to capture other sources of differentiation. In this column, we allow for the patterns of substitution to differ based on the geographical origin of the trade flow. We divide the set of origin countries in 13 regions and construct the artificial regressor to test whether countries from the same region display stronger substitution patterns between them. In the case of a discrete variable, the artificial regressor is

$$K_{region} = \frac{1}{2} - s_{region} \tag{15}$$

where $s_{region}$ is the market share of the region that a country belongs to.[17] Intuitively, a country that

[17] To derive the expression of $K_{region}$, consider that consumer preferences include a full set of regional dummies $\mathbb{1}_{o\in r}$ for each one of the different regions $r$. Consumers differ in the valuation of each one of these dummies. Assuming that the variance of the valuation is the same for each dummy $\Sigma_r = \Sigma_{region} \quad \forall r$, $\Sigma_{region}$ can be identified by including the artificial regressor $\frac{1}{2}\sum_r(\mathbb{1}_{o\in r} - s_{r,m})^2$ in the gravity equation. One can show that when market fixed effects are included in the equation, including $\frac{1}{2}\sum_r(\mathbb{1}_{o\in r} - s_{r,m})^2$ is equivalent to including

TABLE 5: FRAC regressions: alternative specifications

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| $\ln p$ | $-2.43$*** | $-2.58$*** | | $-2.78$*** |
| | (0.07) | (0.08) | | (0.08) |
| $K_{\ln p}$ | 0.22*** | 0.37*** | | 0.55*** |
| | (0.01) | (0.02) | | (0.02) |
| $\ln p \times \mu_Y$ | | 0.38*** | | |
| | | (0.04) | | |
| $K_{\ln p} \times V_Y$ | | 0.26*** | | |
| | | (0.05) | | |
| log tariff | | | $-2.49$*** | |
| | | | (0.06) | |
| $K_{\ln GDPc}$ | | | 0.23*** | |
| | | | (0.00) | |
| $K_{region}$ | | | | 0.35*** |
| | | | | (0.02) |
| log distance | $-0.47$*** | $-0.43$*** | $-0.93$*** | $-0.51$*** |
| | (0.01) | (0.02) | (0.00) | (0.02) |
| Common currency | 0.13*** | 0.12*** | 0.19*** | 0.12*** |
| | (0.01) | (0.01) | (0.01) | (0.01) |
| Common border | 0.57*** | 0.55*** | 0.68*** | 0.57*** |
| | (0.01) | (0.01) | (0.00) | (0.01) |
| Common language | 0.18*** | 0.20*** | 0.24*** | 0.20*** |
| | (0.01) | (0.01) | (0.00) | (0.01) |
| Common religion | 0.07*** | 0.03*** | 0.15*** | 0.09*** |
| | (0.01) | (0.01) | (0.01) | (0.01) |
| Colonial links | 0.06*** | 0.06*** | 0.20*** | 0.04*** |
| | (0.01) | (0.01) | (0.00) | (0.01) |
| N | 11 660 255 | 11 660 255 | 11 619 710 | 11 660 255 |
| F-stat first stage | 2 977 | 1 118 | 3 772 437 | 1 937 |

*Notes:* Standard errors between parentheses are clustered at the origin-product and destination-product levels. All regressions include origin-product-year and destination-product-year fixed effects. * p < 0.1, ** p < 0.05, *** p < 0.01.

belongs to a region that performs well in a specific market indicates more direct competition for that country. As a result, this country does not perform as well in this specific market.[18] The results in table 5 shows that geographical differentiation has an important role on export performance: countries perform relatively better when competing countries from the same region are less present in a destination markets. To put things in perspective, the coefficient equal to 0.35 means that a country belonging to a region with a market share of 20 percents exports 3.5 percent less than a country from a region with a market share of 10 percents. Therefore, this regression shows that both vertical differentiation – through prices – and geographical differentiation – through the region of origin play a role in shaping export performance across countries.

artificial regressor $K_{region} = \frac{1}{2} - s_{region}$.

[18] Note that a similar closed-form solution exists in the nested CES model, that can accommodate heterogeneous substitution patterns based on a discrete grouping. In the nested CES model, the regressor that captures these patterns of substitution is based on the market share of a variety within the group. The FRAC framework identifies these effects in a similar fashion, but can accommodate discrete and continuous variables with linear estimation.

# 5 Out of sample predictions

In section 2 of this manuscript, we studied the reallocation patterns across countries induced by the imposition of tariffs on Chinese goods during the recent US-China trade war. This episode allowed us to estimate the cross-price elasticities of many exporters with respect to Chinese exports, and show that these elasticities are not constant across countries, rejecting the IIA assumption imposed by standard gravity models. Specifically, countries that gained the most from the imposition of tariffs on Chinese exports to the US were countries with similar prices to China. This result motivated the development of the mixed-CES model that can generate a role for vertical differentiation and more realistic patterns of substitution across countries.

In this section, we circle back to the motivating evidence and test the ability of our estimated model to predict the reallocation patterns across countries that resulted from the trade war. We rely on the expression derived in the section 3 for the cross-price elasticity, coupled with the estimated parameters of the model, to obtain a prediction for which countries should gain the most from this episode. Importantly, this prediction is made out of sample, in the sense that it only relies on information available before the actual start of the trade war. Therefore, it is a true test for the ability of the model to improve the predictive power of standard gravity models.

## 5.1 Implementation

The objective is to obtain a prediction for the normalized cross-price elasticity, $e_{o,C,m} \equiv \frac{1}{s_{C,m}} \frac{\partial \ln s_{o,m}}{\partial \ln p_{C,m}}$, that we estimate separately for each origin country in section 2. This object can be estimated using FRAC. Specifically, plugging equation 13 into 14 we have

$$e^{FRAC}_{o,C,m} = \bar{\alpha}_d - \Sigma_d(\ln p_{C,m} - \overline{\ln p}_m) - \Sigma_d(\ln p_{o,m} - \overline{\ln p}_m)\left[\frac{1 - \bar{\alpha}_d(\ln p_{C,m} - \overline{\ln p}_m) + \Sigma_d(\ln p_{C,m} - \overline{\ln p}_m)^2}{1 + \Sigma_d \mathbb{V}^{\ln p}_d}\right]$$

This equation indicates that we can directly compute the cross-price elasticity between two varieties, by knowing the distribution of prices $\ln p$ and market shares $s$ in the market and the value of the parameters $\bar{\alpha}_d$ and $\Sigma_d$.

To compute this object, we rely on observations from the year 2017, preceding the trade war that took place in 2018: for each country $o$ exporting to the US and each hs10 product $k$ that will be targeted by a tariff, we observe the price and market share for the year 2017. Moreover, we compute the US-specific parameters based on the mean and variance of the income distribution in the United States: $\begin{pmatrix} \bar{\alpha}_{US} \\ \Sigma_{US} \end{pmatrix} = \begin{pmatrix} \hat{\alpha} + \hat{\pi}\mu_{US} \\ \hat{\gamma}^2 + \hat{\lambda}^2 V_{US} \end{pmatrix}$. We combine these observations and estimated parameters for the US to obtain a prediction of the cross-price elasticity of country $o$ with respect to tariffs on Chinese products, $\hat{e}_{o,C,k,US}$, for each hs10 category $k$. Finally, we compute the average cross-price elasticity as $\hat{e}^{FRAC}_{o,C,US} = \frac{1}{N_k} \sum_k \hat{e}^{FRAC}_{o,C,k,US}$.

## 5.2 Results

In figure 4, we compare the estimated cross-elasticity in section 2 with the elasticities predicted by our model: the y-axis reports the predictions for each origin country $o$ based on the formula above, while the x-axis is based on the estimation of the country-specific cross-price elasticities in section 2. This figure shows that the model does a good job at predicting which countries won the most from the reduction in Chinese exports to the United States. In particular, it identifies the biggest winners such as Vietnam, India or Turkey with large cross-price elasticities.

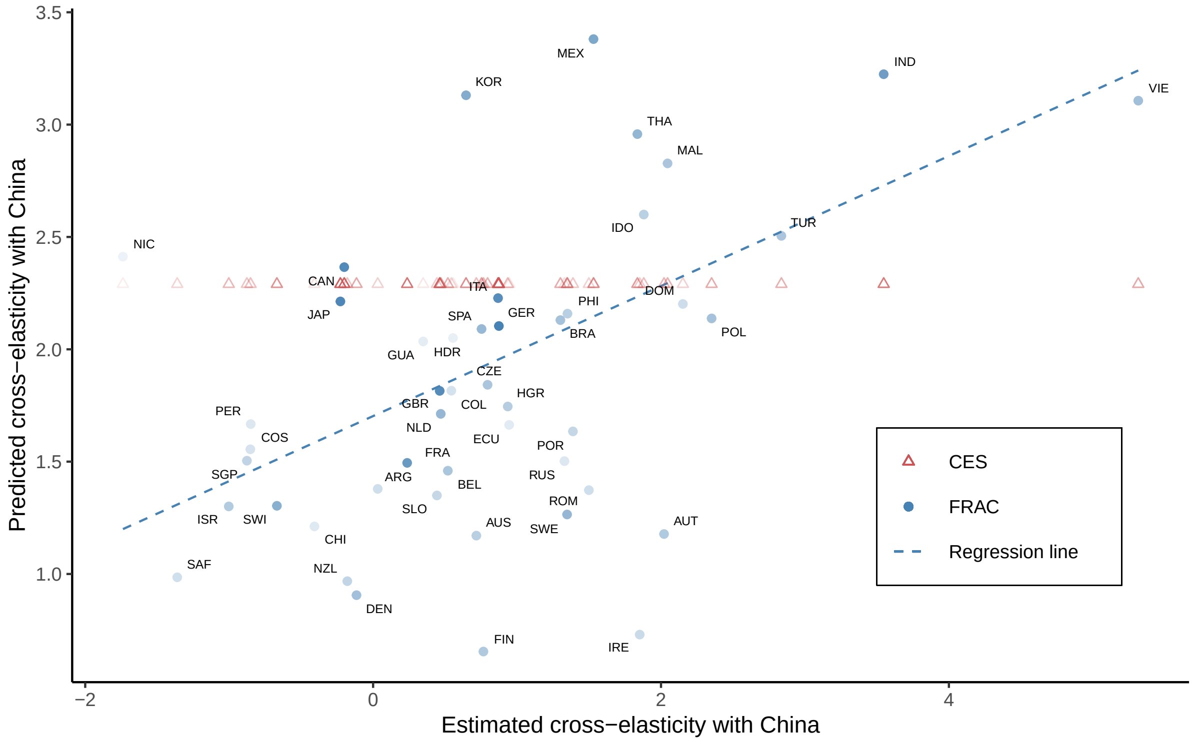


FIGURE 4: Prediction vs realization of trade reallocation from the US-China trade war.

*Notes:* Each point reports the predicted and estimated cross-price elasticity of an origin country. The dots are predictions based on the FRAC model, while the triangles are based on a CES model. The transparency of each marker is related to the standard error of their estimated elasticity. The regression line is based on a regression weighted by the inverse of the standard error of the estimated elasticity.

To quantify the performance of the model at predicting the reallocation of Chinese exports across countries, we run a linear regression between the two measures that is represented with a dashed line on the figure. This relationship is significantly positive with an estimated coefficient of 0.95 (p-value $< 0.001$) and a $R^2$ equal to 0.27. As a comparison, we also report the prediction in the case of a CES model. In this scenario, the model does not predict any heterogeneity in the reallocation of Chinese market shares across countries: all countries feature the same cross-price elasticity which is equal to the estimated elasticity of substitution of the CES model, 2.29. In conclusion, our model that allows for heterogeneous patterns of substitutions across countries is a significant improvement to predict the winner and losers from the reallocation effects of a large trade shock.

## 5.3 The Unequal Impact of Trade Wars on the Cost-of-Living of US consumers

Through the lens of our structural model, the uneven redistribution of Chinese market shares across source countries (figure 4) reveals that US consumers have heterogeneous preferences. This heterogeneity implies that consumers differ in the composition of their import basket, with consumers importing more intensively from China losing more purchasing power due to trade wars.

In this section, we use our structural estimates (table 4, column 5) to quantify the unequal impact of trade wars on the cost-of-living of US consumers. Let $d\tau_k$ be the cumulative change in US tariffs on Chinese imports between 2017 and 2019, for product $k$. We make the following counterfactual experiment: if trade wars had happened in 2017, increasing Chinese prices to some counterfactual value $\breve{p}_{k,C,2017} = p_{k,C,2017}(1 + d\tau_k)$, by how much would have increased the cost of imports?

In order to compute the consumer-specific import price index, we first compute the price index at the consumer $\times$ product level and then aggregate across products. A consumer is fully characterized by her income shock $\nu$ and her random shock $\varepsilon$, which together determine her price elasticity

$$\alpha(\nu,\varepsilon) = \alpha + \pi\mu_{Y,US} + \lambda V_{Y,US}^{1/2}\nu + \gamma\varepsilon.$$

For a consumer with shock realisation $(\nu,\, \varepsilon)$, the CES import price index of product category $k$ at date $t$ is

$$P_{kt}(\nu,\varepsilon) = \left(\sum_{o\in\Omega_{kt}} \exp(X'_{okt}\beta + \xi_{okt})p_{okt}^{-\alpha(\nu,\varepsilon)}\right)^{-\frac{1}{\alpha(\nu,\varepsilon)}}$$

Since consumers have Cobb-Douglas preferences across products, the overall import price index is

$$P_t(\nu,\varepsilon) \propto \prod_k P_{kt}(\nu,\varepsilon)^{\eta_k}$$

with $\eta_k$ the budget share of product $k$ in aggregate US imports. Analogously, we define $\breve{P}_t(\nu,\varepsilon)$ as the import price indexed when Chinese prices are equal to their counterfactual value $\breve{p}_{kCt}$. The percentage impact of trade wars on the import cost is

$$\hat{P}(\nu,\varepsilon) = \frac{\breve{P}_{2017}}{P_{2017}}(\nu,\varepsilon) - 1$$

Following the same logic as for the price index, we can compute the consumer $\times$ product specific import share of China

$$s_{Ckt}(\nu,\varepsilon) = \frac{\exp\left(-\alpha(\nu,\varepsilon)\ln p_{Ckt} + X'_{Ckt}\beta + \xi_{Ckt}\right)}{\sum_{o'\in\mathcal{O}_{Ckt}} \exp\left(-\alpha(\nu,\varepsilon)\ln p_{okt} + X'_{okt}\beta + \xi_{okt}\right)},$$

which we aggregate across products to obtain the consumer-specific Chinese import penetration:

$$s_{Ct}(\nu,\varepsilon) = \sum_k \eta_k s_{Ckt}(\nu,\varepsilon)$$

In figure 5 we plot the import share $s_{C,2017}(\nu,\varepsilon)$ for $\varepsilon = 0$ (the mean value of $\varepsilon$) and for values of $\nu$ corresponding to the percentiles of the standard normal. As depicted in figure 5, our estimates imply that Chinese import penetration is about 2 percentage points larger for US consumers at the bottom of the income distribution than at the top.

The direct implication of these differences in Chinese import penetration is that poor consumers are more sensitive to tariffs on Chinese products. To get a sens of the magnitude of this differential impact, in figure 6 we plot $\hat{P}$, the percentage change impact of trade wars on the import price. Consumers at the bottom of the US income distribution suffered roughly 10% more that consumers at the top. Notice that these results were obtained without giving up on the simplicity of our estimation method: quantifying the price index impact required neither assumptions on the distribution of random coefficient $\varepsilon$ nor numerical integration technique. Our empirical method can therefore be used to easily evaluate the distributional impact of trade policy both across source countries and across consumers.

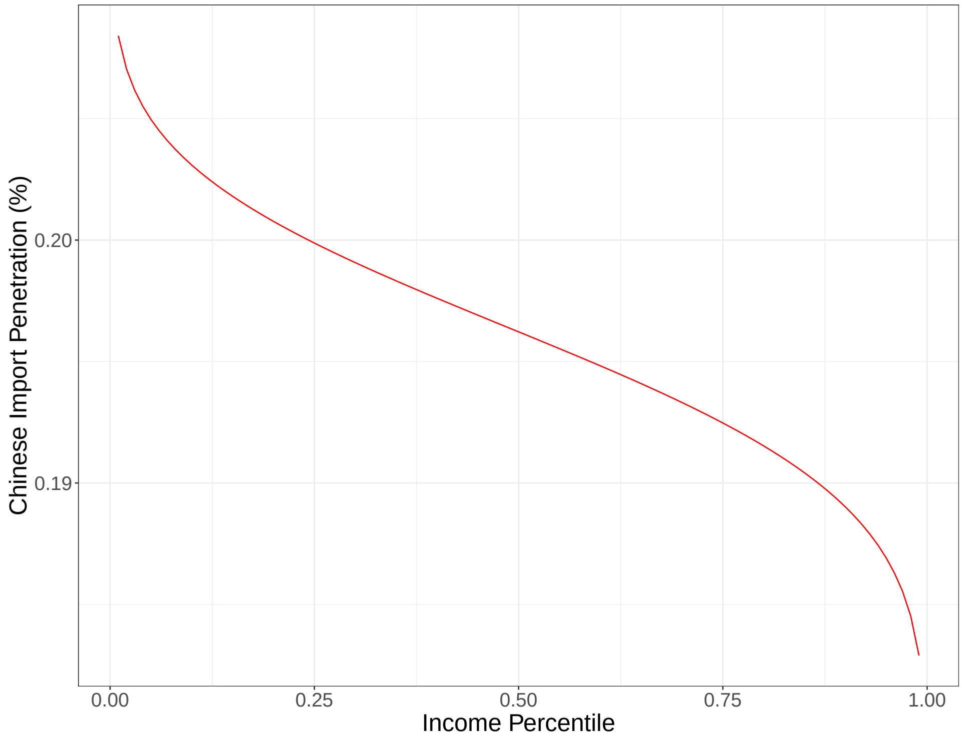


FIGURE 5: Chinese Import Penetration along the US Income Distribution

*Notes:*

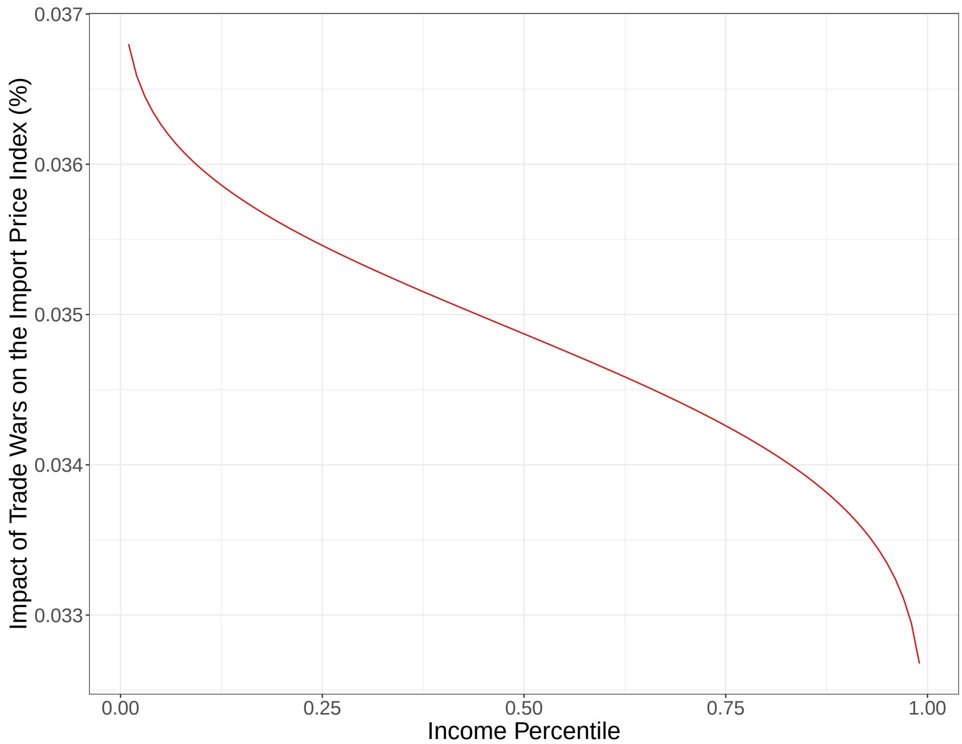


FIGURE 6: Chinese Import Penetration along the US Income Distribution

*Notes:*

# 6 Conclusion

In this paper, we propose a strategy to estimate mixed CES preferences using bilateral country-level trade data. We rely on the "small-$\sigma$" approach from Salanié and Wolak (2022) to linearize the demand system and structurally derive a log-linear gravity equation. Compared to standard models, our gravity equation features "artificial regressors", which measure the distance of a variety to its competitors in terms of product characteristics. These regressors identify the role of product differentiation in substitution patterns. The larger the coefficient on artificial regressors, the more substitution patterns between two varieties depend on their relative positioning in the product space, the starker the violation of the IIA assumption. Importantly, our augmented gravity equation makes it possible to relax the IIA assumption and obtain realistic substitution patterns at virtually no cost compared : estimation relies on linear techniques and has the same data requirement as traditional gravity estimation. Perhaps the only additional cost is to construct the additional regressors, which is straighforward: these regressors are simply quadratic distances of a variety's product characteristics to the its average competitor.

When implemented on bilateral trade data, our estimation reveals important deviations from the IIA assumption. Varieties which are similar in terms of prices or geographical location are significantly more substitutable. We further document the out-of-sample performance of our method in the context of the recent US-China trade war. Using our structural gravity and data prior to the beginning of the war, we are able to predict which source country most benefited from the reduction of Chinese exports to the US. In comparison, CES preferences make the prediction that all source countries equally benefited from the trade war. The data strongly rejects this prediction.

In light of its simplicity, our method provides an important tool to evaluate the impact of product differentiation on export performance and the resilience to competition shocks.

## References


Adao, R., A. Costinot, and D. Donaldson (2017): "Nonparametric counterfactual predictions in neoclassical models of international trade," *American Economic Review*, 107, 633–89.

Bergstrand, J. H. (1990): "The Heckscher-Ohlin-Samuelson model, the Linder hypothesis and the determinants of bilateral intra-industry trade," *The Economic Journal*, 100, 1216–1229.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, 841–890.

Berry, S. T. (1994): "Estimating discrete-choice models of product differentiation," *The RAND Journal of Economics*, 242–262.

Berry, S. T. and P. A. Haile (2021): "Foundations of demand estimation," in *Handbook of Industrial Organization*, Elsevier, vol. 4, 1–62.

Birchall, C. and F. Verboven (2022): "Estimating Substitution Patterns and Demand Curvature in Discrete-Choice Models of Product Differentiation," .

Breinlich, H., H. Fadinger, V. Nocke, and N. Schutz (2020): "Gravity with granularity," .

Dubé, J.-P., A. Hortaçsu, and J. Joo (2021): "Random-coefficients logit demand estimation with zero-valued market shares," *Marketing Science*, 40, 637–660.

Fajgelbaum, P., P. K. Goldberg, P. J. Kennedy, A. Khandelwal, and D. Taglioni (2021): "The US-China trade war and global reallocations," Tech. rep., National Bureau of Economic Research.

Gandhi, A. and J.-F. Houde (2019): "Measuring Substitution Patterns in Differentiated-Products Industries," Tech. rep., National Bureau of Economic Research.

Gandhi, A. and A. Nevo (2021): "Empirical models of demand and supply in differentiated products industries," in *Handbook of Industrial Organization*, Elsevier, vol. 4, 63–139.

Gaulier, G. and S. Zignago (2010): "Baci: international trade database at the product-level (the 1994-2007 version)," .

Head, K. and T. Mayer (2014): "Gravity equations: Workhorse, toolkit, and cookbook," in *Handbook of international economics*, Elsevier, vol. 4, 131–195.

——— (2021): *Poor Substitutes?: Counterfactual Methods in IO and Trade Compared*, Centre for Economic Policy Research.

Lind, N. and N. Ramondo (2018): "Trade with correlation," Tech. rep., National Bureau of Economic Research.

Miller, K. S. (1981): "On the inverse of the sum of matrices," *Mathematics magazine*, 54, 67–72.

Piveteau, P. and G. Smagghue (2022): "Foreign competition along the quality ladder," in *Technical Report, Working Paper*.

Salanié, B. and F. A. Wolak (2022): "Fast, Detail-free, and Approximately Correct: Estimating Mixed Demand Systems," .

Yotov, Y. V., R. Piermartini, M. Larch, et al. (2016): *An Advanced Guid to Trade Policy Analysis: The Structural Gravity Model*, WTO iLibrary.

## Appendices

# A Data Construction

Our estimation mainly relies on BACI. In this dataset, a unit of observation is a combination of a source country, a destination country, a six digit product category of the HS classification ($HS6$) and a year. For each observation, we know the value of the shipment along with the physical quantity shipped. This appendix describes the way we prepare the data for estimation.

**Geographical Coverage** We limit the set of destination countries to the 40 countries present in the WIOD database. [19] Because trade flows involving Luxembourg and Belgium are reported together in the raw trade data, we input all of Luxembourg's trade to Belgium. We drop Taiwan, Malta and Cyprus for which the tariffs are missing. In total, we have 36 destinations. We keep observations for the top 100 exporters to these destinations. We eliminate destination $\times$ $HS6$ $\times$ year markets with less than 5 exporters.

**Tariffs data** Our tariff measures come from the Market Access Map (MAcMap) dataset provided by the CEPII. It provides bilateral information on the applied tariffs rates at the HS6 level for four years: 2001, 2004, 2007 and 2010.

**Market Share and Price of the domestic Good** In order to study the role of the domestic good in our demand system estimation, we need information on the domestic market share and price. At the two-digits level of the CPA classification, we construct the market share of the domestic good by computing the share of domestic consumption in total consumption from the WIOD database. We then convert these domestic shares to HS6 using a correspondence table available on RAMON Eurostat Metadata Server.

The estimation also requires to know the price of the outside good. However, the price of the domestic variety is not available in our international trade data since domestic goods do not cross a border. In order to proxy the price of the domestic good in a given country and year, we use the price of its exports as measured in the BACI dataset. Since we observe this price for many destinations, we infer the domestic unit values by regressing the logarithm of the FOB unit value on a set of fixed effects:

$$\ln p_{ijkt} = \delta_{ikt} + \gamma_j + \varepsilon_{ijkt}.$$

$\delta_{ikt}$ is the average price of variety $ikt$, controlling for the destination market. We use it as a proxy for domestic prices. However, some countries don't export the relevant products $ikt$. When this is the case, I estimate the regressions at a lower level (HS4 or HS2) and extrapolate the predicted price for the good.

**Income Distribution** Our estimation requires information on income distribution. We obtain information on income per capita and the Gini index by destination country from the World Bank. In order to feed this information into the estimation, we assume that income distribution is log-normal. This distribution is convenient because it makes it possible to recover the mean $\mu_{y_d}$ and standard deviation $\sigma_{y_d}$ parameters from the average income per capita $m_{y_d}$ and Gini Index $\Lambda_{y_d}$, through following formula

[19] For countries absent from WIOD, we are unable to construct the domestic market share. Observing the domestic market share is when estimating an import demand system. Nonetheless, we impose this data restriction throughout our analysis to keep the set of destinations fixed across estimations.

$$\begin{aligned}\sigma_{y_d} &= \sqrt{2}\Phi^{-1}\left(\tfrac{1+\Lambda_{y_d}}{2}\right)\\ \mu_{y_d} &= \ln m_{y_d} - \tfrac{1}{2}\sigma_{y_d}\end{aligned}$$

# B Additional Tables

# C Derivation of Equation (11)

In this appendix, we derive equation (11) using the "small-$\sigma$" approach. Let $G_d$ be the such that $\Sigma_d = \sigma^2 G_d^2$, with $\Sigma_d$ the income-weighted variance of $\alpha_i$ in destination $d$. We can re-write equation (7) as

$$\alpha_i(\sigma) = \alpha + \pi\mu_d + \sigma G_d \epsilon_i, \tag{16}$$

with $\epsilon_i$ a scalar shock whose income-weighted mean is zero and whose income-weighted variance is one. For any variable $x$, let $\underset{\cdot}{x}_o \equiv x_o - x_r$ be the deviation of $x_o$ relative to a reference variety $r$. The market share of variety $o$ in consumer $i$'s expenditure is

$$s_{io}(\sigma, \underset{\cdot}{\boldsymbol{\xi}}_m) = \frac{\exp\left(-\alpha_i(\sigma)\ln \underset{\cdot}{p}_{om} + \underset{\cdot}{X}'_{om}\beta + \underset{\cdot}{\xi}_{om}\right)}{1 + \sum_{o'\neq r}\exp\left(-\alpha_i(\sigma)\ln \underset{\cdot}{p}_{o'm} + \underset{\cdot}{X}'_{o'm}\beta + \underset{\cdot}{\xi}_{o'm}\right)},$$

with $\underset{\cdot}{\boldsymbol{\xi}}_m$ the vector of normalized demand quality in market $m$. The aggregate revenue market share of variety $o$ is

$$s_{om}(\sigma, \underset{\cdot}{\boldsymbol{\xi}}_m) = \int s_{io}\omega_{id}di, \tag{17}$$

with $\int \omega_{id} = 1$, $\omega_{id} \propto y_i$. Let $\underset{\cdot}{\boldsymbol{\xi}}_m(\sigma)$ be the vector of structural errors in market $m$, i.e. the vector of errors which, for a given value of the scaling parameter $\sigma$, equalize the market shares in the data and in the model:

$$\underset{\cdot}{\boldsymbol{\xi}}_m(\sigma): \quad \underset{\cdot}{\boldsymbol{\xi}} \text{ s.t. } s_{om}(\underset{\cdot}{\boldsymbol{\xi}}, \sigma) = S_{om} \quad \forall o \neq r, \tag{18}$$

We seek to obtain an analytical expression for the second-order expansion of $\underset{\cdot}{\boldsymbol{\xi}}_m(\sigma)$ around $\sigma = 0$

$$\underset{\cdot}{\boldsymbol{\xi}}_m(\sigma) = \underset{\cdot}{\xi}(0) + \left.\frac{\partial\underset{\cdot}{\xi}(\sigma)}{\partial\sigma}\right|_{\sigma=0}\sigma + \left.\frac{\partial^2\underset{\cdot}{\xi}(\sigma)}{\partial\sigma^2}\right|_{\sigma=0}\frac{\sigma^2}{2} + O(\sigma^2). \tag{19}$$

**First-order derivative of $\underset{\cdot}{\boldsymbol{\xi}}(\sigma)$** Hereafter, we drop subscript $m$ to save on notations. In this paragraph, we derive an expression for $\left.\frac{\partial\underset{\cdot}{\xi}(\sigma)}{\partial\sigma}\right|_{\sigma=0}$. One important derivative that we will use throughout:

$$\frac{\partial s_{io}(\sigma, \underset{\cdot}{\boldsymbol{\xi}})}{\partial\sigma} = s_{io}(\sigma, \underset{\cdot}{\boldsymbol{\xi}})\widetilde{\ln p}_{io}(\sigma)G\epsilon_i \tag{20}$$

| | | (1) | (2) |
|---|---|---|---|
| Treatment | × ARGENTINA | 0.01 | 0.03 |
| | | (0.04) | (0.96) |
| | × AUSTRALIA | 0.04 | 0.72 |
| | | (0.03) | (0.69) |
| | × AUSTRIA | 0.12*** | 2.02*** |
| | | (0.02) | (0.62) |
| | ×BELGIUM | 0.05** | 0.52 |
| | | (0.02) | (0.58) |
| | ×BRAZIL | 0.04 | 1.30** |
| | | (0.02) | (0.54) |
| | ×CANADA | −0.01 | −0.20 |
| | | (0.01) | (0.29) |
| | ×CHILE | −0.01 | −0.41 |
| | | (0.04) | (1.39) |
| | ×CHINA | −0.30*** | −6.39*** |
| | | (0.01) | (0.23) |
| | ×COLOMBIA | 0.04 | 0.54 |
| | | (0.04) | (1.05) |
| | ×COSTA RICA | 0.02 | −0.85 |
| | | (0.05) | (1.11) |
| | ×CZECH REPUBLIC | 0.06** | 0.80 |
| | | (0.03) | (0.58) |
| | ×DENMARK | 0.03 | −0.12 |
| | | (0.03) | (0.53) |
| | ×DOMINICAN REPUBLIC | 0.07* | 2.15** |
| | | (0.04) | (0.94) |
| | ×ECUADOR | 0.06 | 0.95 |
| | | (0.05) | (1.48) |
| | ×FINLAND | 0.08*** | 0.77 |
| | | (0.03) | (0.61) |
| | ×FRANCE | 0.01 | 0.24 |
| | | (0.02) | (0.34) |
| | ×GERMANY | 0.04*** | 0.87*** |
| | | (0.01) | (0.28) |
| | ×GUATEMALA | 0.01 | 0.35 |
| | | (0.06) | (1.58) |
| | ×HONDURAS | −0.01 | 0.55 |
| | | (0.06) | (1.77) |
| | ×HONG KONG | −0.62*** | −10.97*** |
| | | (0.04) | (0.70) |
| | ×HUNGARY | 0.05 | 0.94 |
| | | (0.03) | (0.66) |
| | ×INDIA | 0.20*** | 3.55*** |
| | | (0.02) | (0.35) |
| | ×INDONESIA | 0.07** | 1.88*** |
| | | (0.03) | (0.70) |
| | ×IRELAND | 0.07* | 1.85** |
| | | (0.04) | (0.87) |
| | ×ISRAEL | 0.01 | −1.00 |
| | | (0.03) | (0.63) |

| | | |
|---|---|---|
| ×ITALY | 0.06*** | 0.87*** |
| | (0.01) | (0.30) |
| ×JAPAN | 0.01 | −0.23 |
| | (0.01) | (0.29) |
| ×KOREA, SOUTH | 0.07*** | 0.65 |
| | (0.02) | (0.41) |
| ×MALAYSIA | 0.13*** | 2.05*** |
| | (0.03) | (0.58) |
| ×MEXICO | 0.08*** | 1.53*** |
| | (0.02) | (0.39) |
| ×NETHERLANDS | 0.04* | 0.47 |
| | (0.02) | (0.47) |
| ×NEW ZEALAND | −0.02 | −0.18 |
| | (0.04) | (0.76) |
| ×NICARAGUA | 0.05 | −1.74 |
| | (0.07) | (2.11) |
| ×PERU | −0.01 | −0.85 |
| | (0.04) | (1.26) |
| ×PHILIPPINES | 0.06* | 1.35** |
| | (0.03) | (0.68) |
| ×POLAND | 0.17*** | 2.35*** |
| | (0.03) | (0.54) |
| ×PORTUGAL | 0.08** | 1.39* |
| | (0.03) | (0.81) |
| ×ROMANIA | 0.09** | 1.50 |
| | (0.04) | (0.95) |
| ×RUSSIA | 0.12** | 1.33 |
| | (0.05) | (1.27) |
| ×SINGAPORE | −0.01 | −0.88 |
| | (0.03) | (0.75) |
| ×SLOVAKIA | 0.04 | 0.44 |
| | (0.04) | (0.84) |
| ×SOUTH AFRICA | −0.02 | −1.36 |
| | (0.04) | (0.99) |
| ×SPAIN | 0.06*** | 0.75 |
| | (0.02) | (0.48) |
| ×SWEDEN | 0.09*** | 1.35*** |
| | (0.02) | (0.47) |
| ×SWITZERLAND | 0.00 | −0.67 |
| | (0.02) | (0.42) |
| ×TAIWAN | 0.02 | −0.02 |
| | (0.01) | (0.31) |
| ×THAILAND | 0.09*** | 1.84*** |
| | (0.02) | (0.46) |
| ×TURKEY | 0.16*** | 2.84*** |
| | (0.02) | (0.58) |
| ×UNITED KINGDOM | 0.04*** | 0.46 |
| | (0.01) | (0.30) |
| ×VIETNAM | 0.29*** | 5.32*** |
| | (0.03) | (0.53) |
| N | 1 948 501 | 1 948 501 |
| $R^2$ | 0.878 | 0.878 |

* p < 0.1, ** p < 0.05, *** p < 0.01

with $\widetilde{\ln p}_{io}(\sigma) \equiv \ln p_o - \overline{\ln p}_i(\sigma)$ and $\overline{\ln p}_i(\sigma) \equiv \sum_{o'} \ln p_{o'} s_{io'}(\sigma)$. Taking the first-order total derivative of equation (34) with respect to $\sigma$ gives

$$\frac{ds_o(\boldsymbol{\xi}(\sigma),\sigma)}{d\sigma} = \frac{\partial s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\sigma} + \sum_{o'\neq r} \frac{\partial s_{o'}(\boldsymbol{\xi}(\sigma),\sigma)}{\partial \xi_{o'}} \frac{\partial \xi_{o'}(\sigma)}{\partial\sigma} = 0. \tag{21}$$

Let $J+1$ be the number of varieties on the market. Equation (37) implies:

$$\frac{\partial \boldsymbol{\xi}(\sigma)}{\partial\sigma} = -B(\sigma)^{-1}A(\sigma), \tag{22}$$

with $A(\sigma)$ a $J\times 1$ vector and $B(\sigma)$ a $J\times J$ matrix such that $A_o(\sigma) \equiv \frac{1}{S_o}\frac{\partial s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\sigma}$ and $B_{o,o'}(\sigma) \equiv \frac{1}{S_o}\frac{\partial s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\xi_{o'}}$, respectively.

When $\sigma = 0$, consumers are identical:

$$\begin{cases} s_{io}(0) = & S_o \\ \widetilde{\ln p}_{io}(0) = & \widetilde{\ln p}_o \end{cases} \quad \forall i, \tag{23}$$

with $\widetilde{\ln p}_o \equiv \ln p_o - \overline{\ln p}$ and $\overline{\ln p} \equiv \sum_o S_o \ln p_o$. Plugging (36) and (39) into $A_o(0)$, we get

$$\begin{aligned} A_o(0) &= \frac{1}{S_o}\int s_{io}(0)\widetilde{\ln p}_{io}(0) G\epsilon_i\omega_{id}\,di \\ &= \widetilde{\ln p}_o G\int_i \epsilon_i\omega_{id}\,di \\ &= 0, \end{aligned} \tag{24}$$

where the third equality uses the fact that $\epsilon_i$ has a zero income-weighted mean. Assuming that $B(0)$ is invertible, from equation (38) and (40) we have

$$\frac{\partial \boldsymbol{\xi}(0)}{\partial\sigma} = 0.$$

**Second order derivative of $\boldsymbol{\xi}(\sigma)$** Next, we turn to the second order derivative. For each $j$, we have:

$$\begin{aligned} \frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\sigma^2} =& \frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\sigma^2} + 2\sum_{k\neq r}\frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\xi_k\partial\sigma}\frac{\partial\xi_k(\sigma)}{\partial\sigma} + \sum_{k\neq r}\frac{\partial s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\xi_k}\frac{\partial^2\xi_k(\sigma)}{\partial\sigma^2} \\ &+ \sum_{k\neq r}\sum_{l\neq k,r}\frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\xi_k\partial\xi_l}\frac{\partial\xi_k(\sigma)}{\partial\sigma}\frac{\partial\xi_l(\sigma)}{\partial\sigma} = 0 \end{aligned} \tag{25}$$

We want to evaluate this term for $\sigma = 0$. Since $\frac{\partial\xi_k}{\partial\sigma}(0) = 0$, the second and fourth terms from equation (41) disappear when $\sigma = 0$:

$$\frac{\partial^2 s_o(\boldsymbol{\xi}(0),0)}{\partial\sigma^2} + \sum_{k\neq r}\frac{\partial s_o(\boldsymbol{\xi}(0),0)}{\partial\xi_k}\frac{\partial^2\xi_k}{\partial\sigma^2}(0) = 0$$

We can rewrite this as

$$\frac{\partial^2 \xi_o(0)}{\partial \sigma^2} = -B(0)^{-1} D(0) \tag{26}$$

with $D(\sigma)$ a $J \times 1$ vector such that $D_o(\sigma) \equiv \frac{1}{S_o} \frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial \sigma^2}$:

$$\begin{aligned}
D_o(\sigma) &= \frac{1}{S_o} \int_i \frac{\partial s_{io}(\sigma)}{\partial \sigma} \breve{\ln p}_{io}(\sigma) G_d \epsilon_i \omega_{id} \, di - \frac{1}{S_o} \int_i s_{io}(\sigma) \left( \sum_k \ln p_k \frac{\partial s_{ik}(\sigma)}{\partial \sigma} \right) G_d \epsilon_i \omega_{id} \, di \\
&= \frac{1}{S_o} \int_i s_{io}(\sigma) \breve{\ln p}_{io}(\sigma) G_d \epsilon_i \breve{\ln p}_{io}(\sigma) G_d \epsilon_i \omega_{id} \, di - \frac{1}{S_o} \int_i s_{io}(\sigma) \left( \sum_k \ln p_k \left( s_{ik}(\sigma) \breve{\ln p}_{ik}(\sigma) G_d \epsilon_i \right) \right) G_d \epsilon_i \omega_{id} \, di \\
&= \frac{1}{S_o} \int_i s_{io}(\sigma) \breve{\ln p}^2_{io}(\sigma) G_d^2 \epsilon_i^2 \omega_{id} \, di - \frac{1}{S_o} \int_i s_{io}(\sigma) \left( \sum_k \ln p_k s_{ik}(\sigma) \breve{\ln p}_{ik}(\sigma) G_d^2 \epsilon_i^2 \omega_{id} \, di \right).
\end{aligned}$$

When $\sigma = 0$, this becomes

$$\begin{aligned}
&= \frac{1}{S_o} \left[ S_o \breve{\ln p}^2_o G_d^2 \left( \int_i \epsilon_i^2 \omega_{id} \, di \right) - S_o \sum_k \ln p_k S_k \breve{\ln p}_k G_d^2 \left( \int_i \epsilon_i^2 \omega_{id} \, di \right) \right] \\
&= \breve{\ln p}^2_o G_d^2 - \sum_k \ln p_k S_k \breve{\ln p}_k G_d^2 \, di \\
&= \frac{1}{\sigma^2} \left[ \breve{\ln p}^2_o \Sigma_d - \sum_k S_k \breve{\ln p}^2_k \Sigma_d \, di \right],
\end{aligned}$$

where the second equality uses the fact that $\int_i \epsilon_i^2 \omega_{id} \, di$ is to the (income-weighted) variance of $\epsilon_i$ (because $\epsilon_i$ has a zero mean), which is one. The second equality uses the definition $\Sigma_d \equiv \sigma^2 G_d^2$ as well as the fact that $\sum_k S_k \breve{\ln p}_k \Sigma_d \overline{\ln p} = 0$: deviations to the mean – $\breve{\ln p}_k$ – add up to zero.

Moving now to matrix $B(0)$ – the other term involved in $\frac{\partial^2 \xi_o(0)}{\partial \sigma^2}$ (equation 42):

$$\begin{aligned}
B_{o,o'}(\sigma) &\equiv \frac{1}{S_o} \frac{\partial s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial \xi'_o} \\
&= \begin{cases} \frac{1}{S_o} \int_i s_{io}(\sigma) s_{io'}(\sigma) \omega_{id} \, di & \text{if } o \neq o' \\ \frac{1}{S_o} \int_i s_{io}(\sigma)(1 - s_{io}(\sigma)) \omega_{id} \, di & \text{if } o = o' \end{cases}.
\end{aligned}$$

Therefore, we have

$$B_{o,o'}(0) = \begin{cases} S_{o'} & \text{if } o \neq o' \\ 1 - S_o & \text{if } o = o' \end{cases},$$

so that we can write $B(0) = I - M$ with $I$ the identity matrix and $M$ a $J \times J$ matrix of rank 1: $M_{o,o'} \equiv S_{o'}$. Using Miller (1981)'s lemma on the inverse of the sum of matrices, we have $B(0)^{-1} = I + \frac{1}{1 - trace(M)} M$, with $trace(M) = \sum_{k \neq r} S_k$. Plugging this expression for $B(0)^{-1}$ into (42):

$$\begin{aligned}
\frac{\partial^2 \xi(0)}{\partial \sigma^2} &= -M^{-1} D(0) = -D(0) - \frac{1}{1 - \sum_{k \neq r} S_k} M D(0) \\
&= -D(0) - \frac{1}{S_r} M D(0),
\end{aligned}$$

which implies

$$
\begin{aligned}
\frac{\partial^2 \xi_o(0)}{\partial \sigma^2} =& - D_o(0) - \frac{1}{S_r}\sum_{k\neq r} S_k D_k(0) \\
=& -\frac{1}{\sigma^2}\left\{\widetilde{\ln p_o}^2\Sigma_d - \sum_k S_k \widetilde{\ln p_k}^2\Sigma_d\right\} - \frac{1}{\sigma^2}\frac{1}{S_r}\left\{\sum_{k\neq r} S_k \widetilde{\ln p_k}^2 - \sum_{l\neq r} S_l \sum_k S_k \widetilde{\ln p_k}^2 \Sigma_d\right\} \\
=& -\frac{1}{\sigma^2}\left\{\widetilde{\ln p_o}^2\Sigma_d - \sum_k S_k \widetilde{\ln p_k}^2\Sigma_d\right\} \\
& - \frac{1}{\sigma^2}\frac{1}{S_r}\left\{\sum_k S_k \widetilde{\ln p_k}^2\Sigma_d - S_r \widetilde{\ln p_r}^2\Sigma_d - (1-S_r)\sum_k S_k \widetilde{\ln p_k}^2\Sigma_d\right\} \\
=& -\frac{1}{\sigma^2}\left(\widetilde{\ln p_o}^2\Sigma_d - \widetilde{\ln p_r}^2\Sigma_d\right) \\
=& -\frac{1}{\sigma^2}\widetilde{\ln p_o^2}\Sigma_d
\end{aligned}
$$

**Taylor expansion** To conclude, let us plug the expression for the first and second derivatives of structural error $\boldsymbol{\xi}(0)$ into Taylor expansion (35):

$$
\begin{aligned}
\xi_o(\sigma) &= \xi_o(0) + \frac{\partial \xi_o(0)}{\partial \sigma}\sigma + \frac{\partial^2 \xi_o(0)}{\partial \sigma^2}\frac{\sigma^2}{2} + O(\sigma^2) \\
&= \ln(S_o/S_r) - X_o'\beta - \alpha \ln p_o - \pi \ln p_o \mu_d + 0 - \frac{1}{2}\widetilde{\ln p_o^2}\Sigma_d + O(\sigma^2) \qquad (27)
\end{aligned}
$$

Equation (11) is obtained by swapping $\log(S_j/S_0)$ and $\xi_j(\sigma)$ in (45):

$$
\begin{aligned}
\ln S_o &= \ln S_r + X_o'\beta - \alpha \ln p_o - \pi \ln p_o \mu_d + 0 + \frac{1}{2}\widetilde{\ln p_o} + O(\sigma^2) \\
&= X_o'\beta - \alpha \ln p_o - \pi \ln p_o \mu_d + \Sigma_d \frac{1}{2}\widetilde{\ln p_o}^2 + \xi_o + (\alpha + \pi\mu_d)\ln P \qquad (28) \\
& \qquad (29)
\end{aligned}
$$

with
$$
(\alpha + \pi\mu_d)\ln P = -\sum_k \exp\left(X_k'\beta - \alpha \ln p_k - \pi \ln p_k \mu_d + \Sigma_d \frac{1}{2}\widetilde{\ln p_k}^2 + \xi_k + O(\sigma^2)\right)
$$

. Moreover, from (16) it is straightforward to show that $\Sigma_d$, the variance of $\alpha_i$, verifies $\Sigma_d = \lambda^2 V_d + \gamma^2$. Plugging this expression into (46) gives

$$
\ln S_o = X_o'\beta - \alpha \ln p_o - \pi \ln p_o \mu_d + \gamma^2 K_o + \lambda^2 V_d K_o + \xi_o + (\alpha + \pi\mu_d)\ln P \qquad (30)
$$

# D Model Extension

In this appendix, we generalize the model from section (3) by allowing for multiple random coefficients and multiple demographics. The regression equation that we derive from this general model nests equation (11) as a special case.

**The Extended Model** Let $X_{1,om}$ and $X_{2,om}$ be respectively the vector of linear and non-linear characteristics of variety $o$ in market $m$ ($X_{2,om}$ includes log prices). Preferences over linear characteristics are constant across consumers: $\beta_{1i} = \beta_1 \, \forall i$. By constrast, preferences for non-linear characteristics differ across consumers:

$$\beta_{2i} = \beta_2 + \Pi\mu_d + \Lambda V_d^{1/2}\nu_i + \Gamma\varepsilon_i. \tag{31}$$

with $\mu_d$ and $V_d$ the income-weighted mean and variance of the $n_D \times 1$ vector of demographic shocks $D_i$. $\nu_i$ is the standardized demographic shock: $D_i = \mu_d + V_d^{1/2}\nu_i$. $\varepsilon_i$ is a $n_\varepsilon \times 1$ vector of random shocks with zero mean and unit variance. $\nu_i$ and $\varepsilon_i$ are independent from each other and i.i.d. across consumers within destinations. $\Pi$ is a $n_{X_2} \times n_D$ matrix mapping destination-level mean demographics into destination-level mean preferences. $\Lambda$ is a $n_{X_2} \times n_D$ matrix mapping individual demographic deviations from the national mean into preference deviations. Finally, $\Gamma$ is a $n_{X_2} \times n_\varepsilon$ matrix mapping random shocks $\varepsilon_i$ into preferences.

Let $\Sigma_d$ be the the income-weighted variance of $\beta_{2i}$ in destination $d$. Since $\nu_i$ and $\varepsilon_i$ are uncorrelated, and both have zero mean and unit variance on each destination, we have

$$\begin{aligned} \Sigma_d &= \Lambda V_d \Lambda' + H, \\ \text{with} \quad H &\equiv \Gamma\Gamma'. \end{aligned}$$

Let $G_d$ be the $(n_{X_2} \times n_{X_2})$ matrix such that $\Sigma_d = \sigma^2 G_d G_d'$. We can re-write equation (31) as

$$\beta_{2i}(\sigma) = \beta_2 + \Pi\mu_d + \sigma G_d \epsilon_i. \tag{32}$$

with $\epsilon_i$ a $n_{X_2} \times 1$ vector of shocks whose income-weighted mean is zero and whose income-weighted variance is the identity matrix.[20] For any variable $x$, let $\underset{\cdot}{x}_o \equiv x_o - x_r$ be the deviation of $x_o$ relative to a reference variety $r$. The market share of variety $o$ in consumer $i$'s expenditure is

$$s_{io}(\sigma, \underset{\cdot}{\boldsymbol{\xi}}_m) = \frac{\exp\left(\underset{\cdot}{X}'_{1,om}\beta_1 + \underset{\cdot}{X}'_{2,om}\beta_{2i}(\sigma) + \underset{\cdot}{\xi}_{om}\right)}{1 + \sum_{o' \neq r} \exp\left(\underset{\cdot}{X}'_{1,o'm}\beta_1 + \underset{\cdot}{X}'_{2,o'm}\beta_{2i}(\sigma) + \underset{\cdot}{\xi}_{o'm}\right)},$$

with $\underset{\cdot}{\boldsymbol{\xi}}_m$ the vector of normalized demand quality in market $m$. The aggregate revenue market share of variety $o$ is

$$s_{om}(\sigma, \underset{\cdot}{\boldsymbol{\xi}}_m) = \int s_{io}\omega_{id} di, \tag{33}$$

with $\int \omega_{id} = 1$, $\omega_{id} \propto y_i$. Let $\underset{\cdot}{\boldsymbol{\xi}}_m(\sigma)$ be the vector of structural errors in market $m$, i.e. the vector of errors which, for a given value of the scaling parameter $\sigma$, equalize the market shares in the data and in the model:

$$\underset{\cdot}{\boldsymbol{\xi}}_m(\sigma): \quad \underset{\cdot}{\boldsymbol{\xi}} \text{ s.t. } s_{om}(\underset{\cdot}{\boldsymbol{\xi}}, \sigma) = S_{om} \quad \forall o \neq r, \tag{34}$$

We seek to obtain an analytical expression for the second-order expansion of $\underset{\cdot}{\boldsymbol{\xi}}_m(\sigma)$ around $\sigma = 0$

[20] Shock $\epsilon$ is related to $\nu$ and $\varepsilon_i$ through

$$\sigma G_d \epsilon_i = \Lambda V_d^{1/2}\nu_i + \Gamma\varepsilon_i$$

$$\underset{\sim}{\boldsymbol{\xi}}_m(\sigma) = \underset{\sim}{\xi}(0) + \left.\frac{\partial \xi(\sigma)}{\partial \sigma}\right|_{\sigma=0} \sigma + \left.\frac{\partial^2 \xi(\sigma)}{\partial \sigma^2}\right|_{\sigma=0} \frac{\sigma^2}{2} + O(\sigma^2). \tag{35}$$

**First-order derivative of $\underset{\sim}{\boldsymbol{\xi}}(\sigma)$** Hereafter, we drop subscript $m$ to save on notations. In this paragraph, we derive an expression for $\left.\frac{\partial \xi(\sigma)}{\partial \sigma}\right|_{\sigma=0}$. One important derivative that we will use throughout:

$$\frac{\partial s_{io}(\sigma, \underset{\sim}{\boldsymbol{\xi}})}{\partial \sigma} = s_{io}(\sigma, \underset{\sim}{\boldsymbol{\xi}}) \check{X}_{2,io}(\sigma)' G \epsilon_i \tag{36}$$

with $\check{X}_{2,io}(\sigma) \equiv X_{2,o} - \bar{X}_{2,i}(\sigma)$ and $\bar{X}_{2,i}(\sigma) \equiv \sum_{o'} X_{2,o'} s_{io'}(\sigma)$. Taking the first-order total derivative of equation (34) with respect to $\sigma$ gives

$$\frac{d s_o(\underset{\sim}{\boldsymbol{\xi}}(\sigma), \sigma)}{d\sigma} = \frac{\partial s_o(\underset{\sim}{\boldsymbol{\xi}}(\sigma), \sigma)}{\partial \sigma} + \sum_{o' \neq r} \frac{\partial s_{o'}(\underset{\sim}{\boldsymbol{\xi}}(\sigma), \sigma)}{\partial \underset{\sim}{\xi}_{o'}} \frac{\partial \xi_{o'}(\sigma)}{\partial \sigma} = 0. \tag{37}$$

Let $J + 1$ be the number of varieties on the market. Equation (37) implies:

$$\frac{\partial \underset{\sim}{\boldsymbol{\xi}}(\sigma)}{\partial \sigma} = -B(\sigma)^{-1} A(\sigma), \tag{38}$$

with $A(\sigma)$ a $J \times 1$ vector and $B(\sigma)$ a $J \times J$ matrix such that $A_o(\sigma) \equiv \frac{1}{S_o} \frac{\partial s_o(\underset{\sim}{\boldsymbol{\xi}}(\sigma), \sigma)}{\partial \sigma}$ and $B_{o,o'}(\sigma) \equiv \frac{1}{S_o} \frac{\partial s_o(\underset{\sim}{\boldsymbol{\xi}}(\sigma), \sigma)}{\partial \underset{\sim}{\xi}_{o'}}$, respectively.

When $\sigma = 0$, consumers are identical:

$$\begin{cases} s_{io}(0) = & S_o \\ \check{X}_{2,io}(0) = & \check{X}_{2,o} \end{cases} \quad \forall i, \tag{39}$$

with $\check{X}_{2,o} \equiv X_{2,o} - \bar{X}_2$ and $\bar{X}_2 \equiv \sum_o S_o X_{2,o}$. Plugging (36) and (39) into $A_o(0)$, we get

$$\begin{aligned} A_o(0) &= \frac{1}{S_o} \int s_{io}(0) \check{X}_{2,io}(0)' G \epsilon_i \omega_{id} \, di \\ &= \check{X}_{2,o}' G \int_i \epsilon_i \omega_{id} \, di \\ &= 0, \end{aligned} \tag{40}$$

where the third equality uses the fact that (i) $\epsilon_i$ has a zero mean and (ii) $\epsilon_i$ is independent from income $y_i$, which implies $\int \epsilon_i \omega_{id} \, di = \int \epsilon_i \, di = 0$. Assuming that $B(0)$ is invertible, from equation (38) and (40) we have

$$\frac{\partial \underset{\sim}{\boldsymbol{\xi}}(0)}{\partial \sigma} = 0.$$

**Second order derivative of $\boldsymbol{\xi}(\sigma)$** Next, we turn to the second order derivative. For each $j$, we have:

$$\begin{aligned}\frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\sigma^2} =& \frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\sigma^2} + 2\sum_{k\neq r}\frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\xi_k\partial\sigma}\frac{\partial\xi_k(\sigma)}{\partial\sigma} + \sum_{k\neq r}\frac{\partial s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\xi_k}\frac{\partial^2\xi_k(\sigma)}{\partial\sigma^2}\\ &+\sum_{k\neq r}\sum_{l\neq k,r}\frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\xi_k\partial\xi_l}\frac{\partial\xi_k(\sigma)}{\partial\sigma}\frac{\partial\xi_l(\sigma)}{\partial\sigma} = 0 \end{aligned}\tag{41}$$

We want to evaluate this term for $\sigma = 0$. Since $\frac{\partial\xi_k}{\partial\sigma}(0) = 0$, the second and fourth terms from equation (41) disappear when $\sigma = 0$:

$$\frac{\partial^2 s_o(\boldsymbol{\xi}(0),0)}{\partial\sigma^2} + \sum_{k\neq r}\frac{\partial s_o(\boldsymbol{\xi}(0),0)}{\partial\xi_k}\frac{\partial^2\xi_k}{\partial\sigma^2}(0) = 0$$

We can rewrite this as

$$\frac{\partial^2\xi_o(0)}{\partial\sigma^2} = -B(0)^{-1}D(0)\tag{42}$$

with $D(\sigma)$ a $J\times 1$ vector such that $D_o(\sigma)\equiv\frac{1}{S_o}\frac{\partial^2 s_o(\boldsymbol{\xi}(\sigma),\sigma)}{\partial\sigma^2}$:

$$\begin{aligned}D_o(\sigma) &= \frac{1}{S_o}\int_i\frac{\partial s_{io}(\sigma)}{\partial\sigma}\check{X}_{2,io}(\sigma)'G_d\epsilon_i\omega_{id}\,di - \frac{1}{S_o}\int_i s_{io}(\sigma)\left(\sum_k X'_{2,k}\frac{\partial s_{ik}(\sigma)}{\partial\sigma}\right)G_d\epsilon_i\omega_{id}\,di\\ &= \frac{1}{S_o}\int_i s_{io}(\sigma)\check{X}_{2,io}(\sigma)'G_d\epsilon_i\check{X}_{2,io}(\sigma)'G_d\epsilon_i\omega_{id}\,di - \frac{1}{S_o}\int_i s_{io}(\sigma)\left(\sum_k X'_{2,k}\left(s_{ik}\check{X}_{2,ik}(\sigma)'G_d\epsilon_i\right)\right)G_d\epsilon_i\omega_{id}\,di\\ &= \frac{1}{S_o}\int_i s_{io}(\sigma)\check{X}_{2,io}(\sigma)'G_d\epsilon_i\check{X}_{2,io}(\sigma)'G_d\epsilon_i\omega_{id}\,di - \frac{1}{S_o}\int_i s_{io}(\sigma)\left(\sum_k s_{ik}(\sigma)\check{X}_{2,ik}(\sigma)'G_d\epsilon_i X'_{2,k}\right)G_d\epsilon_i\omega_{id}\,di\\ &= \frac{1}{S_o}\int_i s_{io}(\sigma)\check{X}_{2,io}(\sigma)'G_d\epsilon_i\epsilon_i'G_d'\check{X}_{2,io}(\sigma)\omega_{id}\,di - \frac{1}{S_o}\int_i s_{io}(\sigma)\sum_k s_{ik}(\sigma)\check{X}_{2,io}(\sigma)'G_d\epsilon_i\epsilon_i'G_d'X_{2,k}\omega_{id}\,di.\end{aligned}$$

When $\sigma = 0$, this becomes

$$\begin{aligned}D_o(0) &= \frac{1}{S_o}\left[S_j\check{X}'_{2,o}G_d\left(\int_i\epsilon_i\epsilon_i'\omega_{id}\,di\right)G_d'\check{X}_{2,o} - S_j\sum_k S_k\check{X}'_{2,k}G_d\left(\int_i\epsilon_i\epsilon_i'\omega_{id}\,di\right)G_d'X_{2,k}\right]\\ &= \check{X}'_{2,k}G_dG_d'\check{X}_{2,o} - \sum_k S_k\check{X}'_{2,k}G_dG_d'X_{2,k}\\ &= \frac{1}{\sigma^2}\left[\check{X}'_{2,o}\Sigma_d\check{X}_{2,o} - \sum_k S_k\check{X}'_{2,k}\Sigma_d X_{2,k}\right] \qquad (43)\\ &= \frac{1}{\sigma^2}\left[\check{X}'_{2,o}\Sigma_d\check{X}_{2,o} - \sum_k S_k\check{X}'_{2,k}\Sigma_d\check{X}_{2,k}\right] \qquad (44)\end{aligned}$$

where the second equality uses the fact that $\int_i\epsilon_i\epsilon_i'\omega_{id}\,di = \int_i\epsilon_i\epsilon_i'\,di$ because $\epsilon_i$ is independent of $y_i$. Moreover, $\int_i\epsilon_i\epsilon_i'\,di$ is equal to the variance of $\epsilon_i$ (because $\epsilon_i$ has a zero mean), which is equal to the identity matrix. The third equality uses the definition $\Sigma_d\equiv\sigma^2G_dG_d'$. The fourth equality makes use of the fact that $\sum_k S_k\check{X}'_{2,k}\Sigma_d\bar{X} = 0$: deviations to the mean add up to zero.

Moving now to matrix $B(0)$ – the other term involved in $\frac{\partial^2 \xi_o(0)}{\partial \sigma^2}$ (equation 42):

$$
\begin{aligned}
B_{o,o'}(\sigma) &\equiv \frac{1}{S_o} \frac{\partial s_o(\boldsymbol{\xi}(\sigma), \sigma)}{\partial \xi'_o} \\
&= \begin{cases} \frac{1}{S_o} \int_i s_{io}(\sigma) s_{io'}(\sigma) \omega_{id}\, di & \text{if } o \neq o' \\ \frac{1}{S_o} \int_i s_{io}(\sigma)(1 - s_{io}(\sigma)) \omega_{id}\, di & \text{if } o = o' \end{cases}.
\end{aligned}
$$

Therefore, we have

$$
B_{o,o'}(0) = \begin{cases} S_{o'} & \text{if } o \neq o' \\ 1 - S_o & \text{if } o = o' \end{cases},
$$

so that we can write $B(0) = I - M$ with $I$ the identity matrix and $M$ a $J \times J$ matrix of rank 1: $M_{o,o'} \equiv S_{o'}$. Using Miller (1981)'s lemma on the inverse of the sum of matrices, we have $B(0)^{-1} = I + \frac{1}{1 - trace(M)} M$, with $trace(M) = \sum_{k \neq r} S_k$. Plugging this expression for $B(0)^{-1}$ into (42):

$$
\begin{aligned}
\frac{\partial^2 \xi(0)}{\partial \sigma^2} &= -M^{-1} D(0) = -D(0) - \frac{1}{1 - \sum_{k \neq r} S_k} M D(0) \\
&= -D(0) - \frac{1}{S_r} M D(0),
\end{aligned}
$$

which implies

$$
\begin{aligned}
\frac{\partial^2 \xi_o(0)}{\partial \sigma^2} =& - D_o(0) - \frac{1}{S_r} \sum_{k \neq r} S_k D_k(0) \\
&= -\frac{1}{\sigma^2} \left\{ \check{X}'_{2,o} \Sigma_d \check{X}_{2,o} - \sum_k S_k \check{X}'_{2,k} \Sigma_d \check{X}_{2,k} \right\} - \frac{1}{\sigma^2} \frac{1}{S_r} \left\{ \sum_{k \neq r} S_k \check{X}'_k \Sigma_d \check{X}_{2,k} - \sum_{l \neq r} S_l \sum_k S_k \check{X}'_{2,k} \Sigma_d \check{X}_{2,k} \right\} \\
&= -\frac{1}{\sigma^2} \left\{ \check{X}'_{2,o} \Sigma_d \check{X}_{2,o} - \sum_k S_k \check{X}'_{2,k} \Sigma_d \check{X}_{2,k} \right\} \\
&- \frac{1}{\sigma^2} \frac{1}{S_r} \left\{ \sum_k S_k \check{X}'_k \Sigma_d \check{X}_{2,k} - S_r \check{X}'_{2,r} \Sigma_d \check{X}_{2,r} - (1 - S_r) \sum_k S_k \check{X}'_{2,k} \Sigma_d \check{X}_{2,k} \right\} \\
&= -\frac{1}{\sigma^2} \left( \check{X}'_{2,o} \Sigma_d \check{X}_{2,o} - \check{X}'_{2,r} \Sigma_d \check{X}_{2,r} \right) \\
&= -\frac{1}{\sigma^2} \underset{\cdot}{\check{X}}'_{2,o} \Sigma_d \underset{\cdot}{\check{X}}_{2,o}
\end{aligned}
$$

**Taylor expansion** To conclude, let us plug the expression for the first and second derivatives of structural error $\underset{\cdot}{\boldsymbol{\xi}}(0)$ into Taylor expansion (35):

$$
\begin{aligned}
\underset{\cdot}{\xi}_o(\sigma) &= \underset{\cdot}{\xi}_o(0) + \frac{\partial \underset{\cdot}{\xi}_o(0)}{\partial \sigma} \sigma + \frac{\partial^2 \underset{\cdot}{\xi}_o(0)}{\partial \sigma^2} \frac{\sigma^2}{2} + O(\sigma^2) \\
&= \ln(S_o / S_r) - \underset{\cdot}{X}'_o \beta - \underset{\cdot}{X}'_{2,o} \Pi \mu_d + 0 - \frac{1}{2} \underset{\cdot}{\check{X}}'_{2,o} \Sigma_d \underset{\cdot}{\check{X}}_{2,o} + O(\sigma^2)
\end{aligned} \tag{45}
$$

Equation (11) is obtained by swapping $\ln(S_j / S_0)$ and $\xi_j(\sigma)$ in (45):

$$\ln S_o = \ln S_r + X_o'\beta + X_{2,o}'\Pi\mu_d + 0 + \frac{1}{2}\check{X}_{2,o}' + O(\sigma^2)$$

$$= X_o'\beta + X_{2,o}'\Pi\mu_d + \frac{1}{2}\check{X}_{2,o}'\Sigma_d\check{X}_{2,o} + \xi_o + (\alpha + \pi\mu_d)\ln P + O(\sigma^2) \tag{46}$$

$$\tag{47}$$

with

$$(\alpha + \pi\mu_d)\ln P = -\sum_k \exp\left(X_k'\beta + X_{2,k}'\Pi\mu_d + \frac{1}{2}\check{X}_{2,o}'\Sigma_d\check{X}_{2,k} + \xi_k + O(\sigma^2)\right)$$

. Moreover, from (32) it is straightforward to show that $\Sigma_d$, the variance of $\beta_{2,i}$, verifies $\Sigma_d = \Lambda V_d\Lambda' + H$, with $H \equiv \Gamma\Gamma'$. Plugging this expression into (46) gives

$$\ln S_o = X_o'\beta + X_{2,o}'\Pi\mu_d + \frac{1}{2}\check{X}_{2,o}'H\check{X}_{2,o} + \frac{1}{2}\check{X}_{2,o}'\Lambda V_d\Lambda'\check{X}_{2,o} + \xi_o + (\alpha + \pi\mu_d)\ln P + O(\sigma^2) \tag{48}$$

In the special case where log-income is the only demographics ($D_i = \ln y_i$) and log-prices the only non-linear characteristics ($X_{2,o} = \ln p_o$), then (48) collapses to equation (11).

**Estimation** (46) can be estimated by 2SLS. To get the list of regressors entering the estimation of (46), let us re-write its different terms using summations rather than matrices. We index non-linear product characteristics by $m$ or $n$ and demographics by $a$ or $b$. We drop subscript $o$. Let us start with $X_2'\Pi\mu_d$:

$$X_2'\Pi\mu_d = \sum_a^{n_D}\sum_{m=1}^{n_{X_2}} \Pi_{a,m}\mu_d[a]X_2'[m]$$

The parameter $\Pi$ involved in $X_2'\Pi\mu_d$ term can be estimated using a full set of interactions between product characteristics $X_2'[m]$ and the mean value for each demographics $\mu_d[a]$. Let us now study $\frac{1}{2}\check{X}_2'H\check{X}_2$:

$$\frac{1}{2}\check{X}_2'H\check{X}_2 = \sum_{m=1}^{n_{X_2}} \frac{1}{2}H_{m,m}\check{X}_2^2[m] + \sum_{m,n<m}^{n_{X_2}} H_{m,n}\check{X}_2[m]\check{X}_2[n]$$

The parameter $H$ involved in $\frac{1}{2}\check{X}_2'H\check{X}_2$ can be estimated through quadratic terms $\check{X}_2^2[m]$ for each product characteristics and a full set of interactions $\check{X}_2[m]\check{X}_2[n]$. Notice that the quadratic terms identify $\frac{1}{2}H_{m,m}$, not $H_{m,m}$. Let us now re-write the final term

$$\frac{1}{2}\check{X}_2'\Lambda V_d\Lambda'\check{X}_2 = \sum_a^{n_D}\sum_{m=1}^{n_{X_2}} \frac{1}{2}\Lambda_{a,m}^2 V_d[a,a]\check{X}_2[m]^2 + \sum_a^{n_D}\sum_{m,n<m}^{n_{X_2}} \Lambda_{a,m}\Lambda_{a,n}V_d[a,a]\check{X}_2[m]\check{X}_2[n]$$
$$+ \sum_{a,b<a}^{n_D}\sum_{m=1}^{n_{X_2}} \Lambda_{a,m}\Lambda_{b,m}V_d[a,b]\check{X}_2[m]^2 + \sum_{a,b<a}^{n_D}\sum_{m,n<m}^{n_{X_2}} 2\Lambda_{a,m}\Lambda_{b,n}V_d[a,b]\check{X}_2[m]\check{X}_2[n]$$

To be brought to the data, this term requires a full set of interactions between:

- the variance of demographics $V_d[a,a]$ and quadratic terms $\check{X}_2^2[m]$.
- the covariance of demographics $V_d[a,b]$ and quadratic terms $\check{X}_2^2[m]$.
- the variance of demographics $V_d[a,a]$ and interactions terms $\check{X}_2[m]\check{X}_2[n]$.

- the covariance of demographics $V_d[a,b]$ and interactions terms $\check{X}_2[m]\check{X}_2[n]$.

**Identification** Estimating $\Lambda$ from (48) by 2SLS may raise some over-identification issue. In equation (48), the $n_D \times n_{X_2}$ coefficients from matrix $\Lambda$ are identified from $\frac{1}{2}\check{X}'_{2,o}\Lambda V_d \Lambda' \check{X}_{2,o}$. The problem is that each coefficient shows up in multiple terms of $\frac{1}{2}\check{X}'_{2,o}\Lambda V_d \Lambda' \check{X}_{2,o}$. This overidentification means that in general (48) cannot be estimated through linear methods. There is one special case where this overidentification issue as absent: when demographics are uncorrelated ($V[a,b] = 0 \quad \forall a \neq b$) and when product characteristics do not depend on the same demographics ($\Lambda_{a,m}\Lambda_{a,n} = 0 \quad \forall a, \forall m \neq n$).

# E Simulation Results

In this subsection, we apply the FRAC estimator to simulated data. We want to assess the ability of FRAC to capture the realistic substitution patterns emerging from a mixed CES data generating process (DGP).

**Data Generating Process** We simulate the data for $n_m = 200$ markets, with $n_o = 25$ varieties on each market. Varieties have two observable characteristics: a scalar $X$ and log prices $\ln p$, both of which are drawn from a standard normal distribution:

$$X_{om}, \ln p_{om} \sim N(0,1).$$

Consumers have mixed CES preferences, such that $i$'s demand for variety $o$ is described by equation (4). The coefficient on $X$ is set to $\beta = 2$. We will consider three different distributions for $\alpha_i$, the random coefficient on log prices elasticity:

$$\alpha_i \sim \begin{cases} Normal(\mu_\alpha, V_\alpha) \\ LogNormal\left(\ln \mu_\alpha - \frac{1}{2}\ln\left(1 + \frac{V_\alpha}{\mu_\alpha^2}\right), \ln\left(1 + \frac{V_\alpha}{\mu_\alpha^2}\right)\right) \\ Uniform\left(-\sqrt{3}V_\alpha^{1/2} + \mu_\alpha, \sqrt{3}V_\alpha^{1/2} + \mu_\alpha\right) \end{cases}$$

These different distributions are parametrized to ensure that they have the same mean $\mu_\alpha = 3$ and variance $V_\alpha = 0.5$. The goal is to assess the influence of the random coefficient distribution on the estimation results. Figure 7 depicts the density of the alternative distributions.

The error of the model is $\xi \sim N(\mu_\xi, V_\xi)$ with $\mu_\xi = -5$ and $V_\xi = 0.5$. We approximate the aggregate demand through $n_c = 1000$ consumers. For each consumer, we draw a random coefficient $\alpha_i^{(f)}$, $f = \{\text{Normal, Log-Normal, Uniform}\}$ from each of the three distributions. Accordingly, we construct the consumer-level market shares and the aggregate market shares for each distribution

$$\begin{cases} s_{io}^{(f)} &= \frac{\exp\left(-\alpha_i^{(f)} \ln p_{om} + X'_{om}\beta + \xi_{om}\right)}{\sum_{k \in \mathcal{O}_m} \exp\left(-\alpha_i^{(f)} \ln p_{km} + X'_{km}\beta + \xi_{km}\right)} \\ s_{om}^{(f)} &= \frac{1}{n_c}\sum_{i=1}^{n_c} s_{io}^{(f)} \end{cases}, \quad f = \{\text{Normal, Log-Normal, Uniform}\}$$

**Estimation** From the simulated data, we estimate $\theta = \{\alpha, \gamma\}$ in three ways. First, we regress by OLS the log market shares on market fixed effects, $X$ and $\ln p$. This gravity equation is consistent with CES preferences. We refer to the resulting estimates as "CES" estimates. Second, we estimate $\theta$ from the linearization of the mixed CES preferences (equation 11). Specifically, we regress $\ln s^{(f)}$ on a set of market fixed effects, $X$, $\ln p$ and the artificial regressor

$$K^{(f)} \equiv \frac{1}{2}\left[\ln p_{om} - \sum_k s_{km}^{(f)} \ln p_{km}\right]^2.$$

FIGURE 7: Distribution of Random Price Coefficients

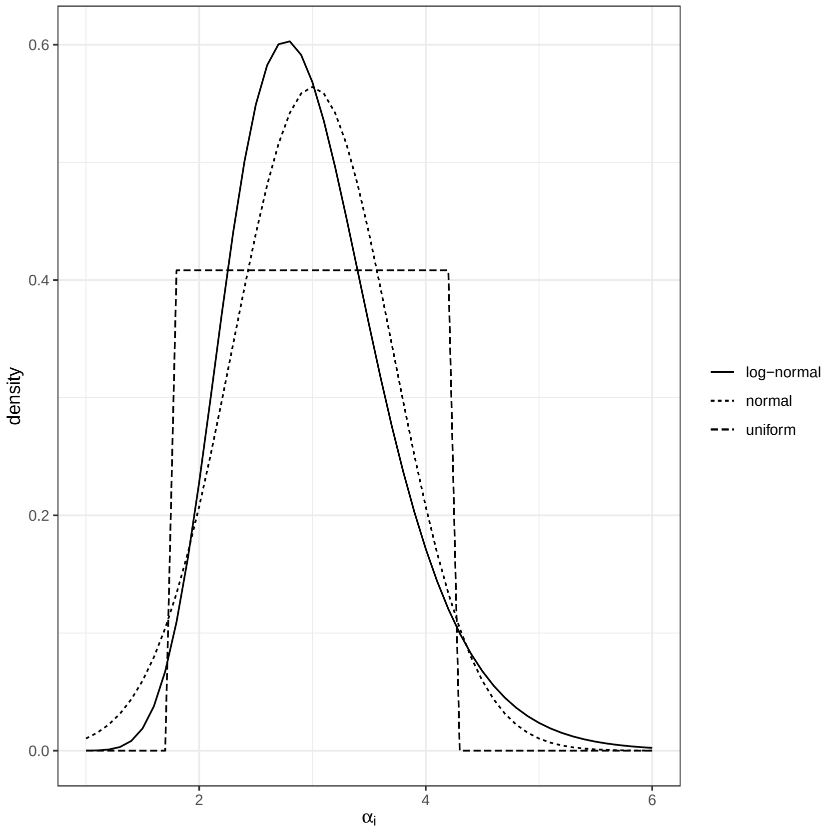


*Notes:* This figure depicts the density of the three alternative distributions (normal, log-normal and uniform) that we use to simulate $\alpha_i$, the random coefficient on prices. All these distributions have the same mean $\mu_\alpha = 3$ and the same variance $V_\alpha = 0.5$.

$K^{(f)}$ is endogenous because it depends on $s_{km}^{(f)}$. We therefore instrument it with

$$\hat{K}_{om}^{(f)} \equiv \frac{1}{2}\left[\ln p_{om} - \sum_k \exp\left(\widehat{\ln s_{km}^{(f)}}\right)\ln p_{km}\right]^2,$$

where $\widehat{\ln s_{om}^{(f)}}$ is the linear prediction of $\ln s_{om}^{(f)}$ based on $X$, $\ln p$ and market fixed effects. We refer to the resulting estimates as "FRAC" estimates.

Third, we estimate $\theta$ by non-linear GMM, following BLP. Let $Z_{om}$ be a set of instruments orthogonal to the $\xi_{om}$. Let $\boldsymbol{\xi}^{(f)}(\theta)$ be the vector of structural errors of the model, i.e. the vector of $\xi$'s equalizing the predicted market shares from the model to the actual market shares in the (simulated) data. $\hat{\theta}_{BLP}^{(f)}$ is obtained by minimizing the distance to zero of a set of moment conditions $E\left[Z\xi^{(f)}(\theta)\right]$. Importantly, recovering the structural error $\xi^{(f)}(\theta)$ requires to fully specify the distribution of random coefficient $\alpha_i$. We estimate $\hat{\theta}_{BLP}^{(f)}$ assuming that $\alpha_i$ is normally distributed, even if $f$ – the true distribution of the random coefficient – is not. We therefore refer to these estimates as "BLP-normal" estimates. Our set of BLP instruments is $Z = \{X, \ln p, \hat{K}^{(f)}\}$. As pointed out by Gandhi and Houde (2019), BLP instruments should include exogenous measures of local competition in the product space to identify $\gamma$, the dispersion in random coefficients. $\hat{K}_{om}^{(f)}$ is precisely a measure of local competition as it captures the price distance of variety $o$ to the average competitor.

Note that when the DGP for $\alpha_i$ is non-normal, both FRAC and BLP-normal are mis-specified, although for different reasons. On the one hand, FRAC does not require assumptions on the distribution

Table 6: Estimates from Simulated Data

| | CES | | | FRAC | | | BLP-normal | | |
|---|---|---|---|---|---|---|---|---|---|
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| RC distribution: | Normal | Log-normal | Uniform | Normal | Log-normal | Uniform | Normal | Log-normal | Uniform |
| $\beta$ $(X)$ | 1.974 | 1.978 | 1.977 | 1.988 | 1.987 | 1.987 | 1.990 | 1.987 | 1.987 |
| | (0.011) | (0.010) | (0.010) | (0.010) | (0.010) | (0.010) | (0.010) | (0.010) | (0.010) |
| $-\alpha$ $(\ln p)$ | -2.407 | -2.542 | -2.537 | -3.010 | -2.940 | -2.943 | -2.986 | -2.916 | -2.915 |
| | (0.017) | (0.013) | (0.013) | (0.027) | (0.026) | (0.027) | (0.025) | (0.025) | (0.025) |
| $\gamma^2$ $(K)$ | | | | 0.440 | 0.290 | 0.294 | 0.526 | 0.314 | 0.317 |
| | | | | (0.018) | (0.018) | (0.018) | (0.027) | (0.022) | (0.030) |
| N | 5,000 | 5,000 | 5,000 | 5,000 | 5,000 | 5,000 | 5,000 | 5,000 | 5,000 |

**Notes:** The true value of the parameters are $\beta = 2$, $-\alpha = 3$ and $\gamma^2 = 0.5$. Estimates from columns (1)-(3) are obtained through OLS regressions of log market shares on log prices and market fixed effects. Estimates from columns (4)-(6) are obtained through 2SLS regressions of log market shares on market fixed effects, log prices and artificial regressor $K$, using $\hat{K}$ as an instrument for $K$. Columns (7)-(9) are obtained from R command "BLPestimatoR". We assume that the random coefficient are normally distributed. We use $X$ as a linear product characteristic, $\ln p$ as non-linear product characteristic and $X$, $\ln p$ and $\hat{K}$ as instruments. All standard errors are clustered at the market level.

of $\alpha_i$. FRAC is therefore mis-specified as a second order approximation to the correctly specified model. On the other hand, BLP-normal estimates the exact model but under a mis-specified distribution of random coefficients. In the next paragraph, we compare the consequences of these different types of mis-specification for the estimated substitution patterns.

**Results: parameter estimates** Table 6 reports the different estimates of $\theta$ based on simulated data. All specifications deliver estimates of $\beta$ – the preference for product characteristic $X$ – which are comparable in magnitude and close to the true value $\beta = 2$. The same applies for $-\alpha$, except for the CES specification which tends to overerestimate it. Compared to the FRAC estimates, it suggests that $-\hat{\alpha}_{CES}$ picks up some of the effect of the omitted variable $K$.[21] Regarding $\gamma^2$ – the variance of the random coefficient – BLP-normal and FRAC deliver similar estimates. When the random coefficients are normal, both BLP-normal and FRAC estimates are close to the true value of 0.5. Specifically, BLP-normal is at 0.53 while FRAC is at 0.44. When the random coefficients are not normal but rather uniform or log-normal, both FRAC and BLP-normal are further away from the truth with an estimate for $\gamma^2$ close to 0.3. Therefore the mis-specification coming from linearizing the model (FRAC) and from assuming the wrong distribution (BLP-normal) deliver comparable biases in terms of structural parameters.

**Results: cross-elasticity estimates** The main strength of mixed preferences is to deliver rich substitution patterns. Therefore, beyond the estimates of the structural parameters, it is really the precision of the estimated substitution patterns that matters. It could be that an estimator does not deliver the true value of the structural parameter and yet delivers good estimates of cross-price elasticities. To study elasticity estimates, let us define $e_{o,c} \equiv \frac{1}{s_c}\frac{\partial \ln s_o}{\partial \ln p_c}$ as the cross-price elasticity of $o$ with respect to $c$, normalized by $s_c$. Under this normalization, the CES cross-price elasticity is constant across all varieties – $\hat{e}^{CES}_{o,c} = \hat{\alpha}^{CES}\,\forall o$ – which facilitates the comparison of predictions across estimators.

Figure 8 plots the relationship between the true and the estimated elasticity. Panel (a) and (b) depict this relationship in the case of a normally distributed RC. Unsurprisingly, BLP-normal performs well in this particular case, since the model is correctly specified. FRAC also does a good job at predicting the true elasticity, except for the upper tail of cross-price elasticities.

Next, we compare FRAC and BLP-normal when the true DGP is not normal, in which case both

[21] $\hat{\alpha}^{FRAC}$ measures the marginal effect of prices when $K = 0$, that is for a variety which is positioned in a very competitive region of the price distribution. By contrast, $\hat{\alpha}^{CES}$ measures the marginal effect of prices on average. Therefore it makes sense that $\hat{\alpha}^{CES} < \hat{\alpha}^{FRAC}$

estimators are mis-specified. It is therefore not clear ex-ante which should perform better. As depicted in panel (c), (d), (e) and (f), and similarly to the Normal case, both FRAC and BLP-normal are able to precisely predict the elasticities for most of the data when the RC is not normal. The impact of the DGP on the relative predictive preformance of FRAC and BLP-normal is concentrated on larger true elasticities. In the log-Normal case, BLP-normal is not as dominant over FRAC as in the Normal case. This is because in the log-normal case (i) FRAC predictions stand closer to the truth and (ii) BLP-normal tends to over-estimate large cross-price elasticities. In the case of a uniform DGP, FRAC predictions are even closer to actual values than in the log-Normal case, to the extent that FRAC clearly dominate BLP-Normal when it comes to predicting large values of cross-elasticities.

To sum up, both estimators deliver accurate and robust predictions for most of the data, except for the largest values. For extreme values, which estimator dominates depends on the (unknown) distribution of random coefficients. There is therefore no strict dominance of BLP-normal over FRAC in terms of prediction. Moreover, FRAC strictly dominates BLP-normal in terms of implementation and transparency of the identification.

# F Deriving Cross-Elasticities

From the definitions $\overline{\ln p} \equiv \sum_k s_k \ln p_k$ and $\widetilde{\ln p_o} = \ln p_o - \overline{\ln p}$, we get

$$\frac{d\overline{\ln p}}{d\ln p_c} = \sum_k \frac{d\ln s_k}{d\ln p_c} s_k \ln p_k + s_c \tag{49}$$

$$\frac{d\ln s_o}{d\ln p_c} = \frac{\Sigma_d}{2}\left[\frac{d\widetilde{\ln p_o}}{d\ln p_c} - \sum_k s_k \frac{d\widetilde{\ln p_k}}{d\ln p_c}\right] - \bar{\alpha}(\mathbb{1}_{o=c} - s_c) \tag{50}$$

$$\frac{d\widetilde{\ln p_o}}{d\ln p_c} = 2\widetilde{\ln p_o}\left(\mathbb{1}_{o=c} - \frac{d\overline{\ln p}}{d\ln p_c}\right) \tag{51}$$

Plugging (51) into (50) gives

$$\frac{\partial \ln s_{om}}{\partial \ln p_{cm}} = -\Sigma_d \frac{\partial \overline{\ln p}_m}{\partial \ln p_{cm}} \widetilde{\ln p}_{om} + \Sigma_d(\mathbb{1}_{o=c} - s_c)\widetilde{\ln p}_{cm} - \bar{\alpha}(\mathbb{1}_{o=c} - s_c)$$

which is nothing but main text equation (13). Plugging (13) into (49) gives the equilibrium expression for $\frac{\partial \overline{\ln p}_m}{\partial \ln p_{cm}}$ (equation 14).

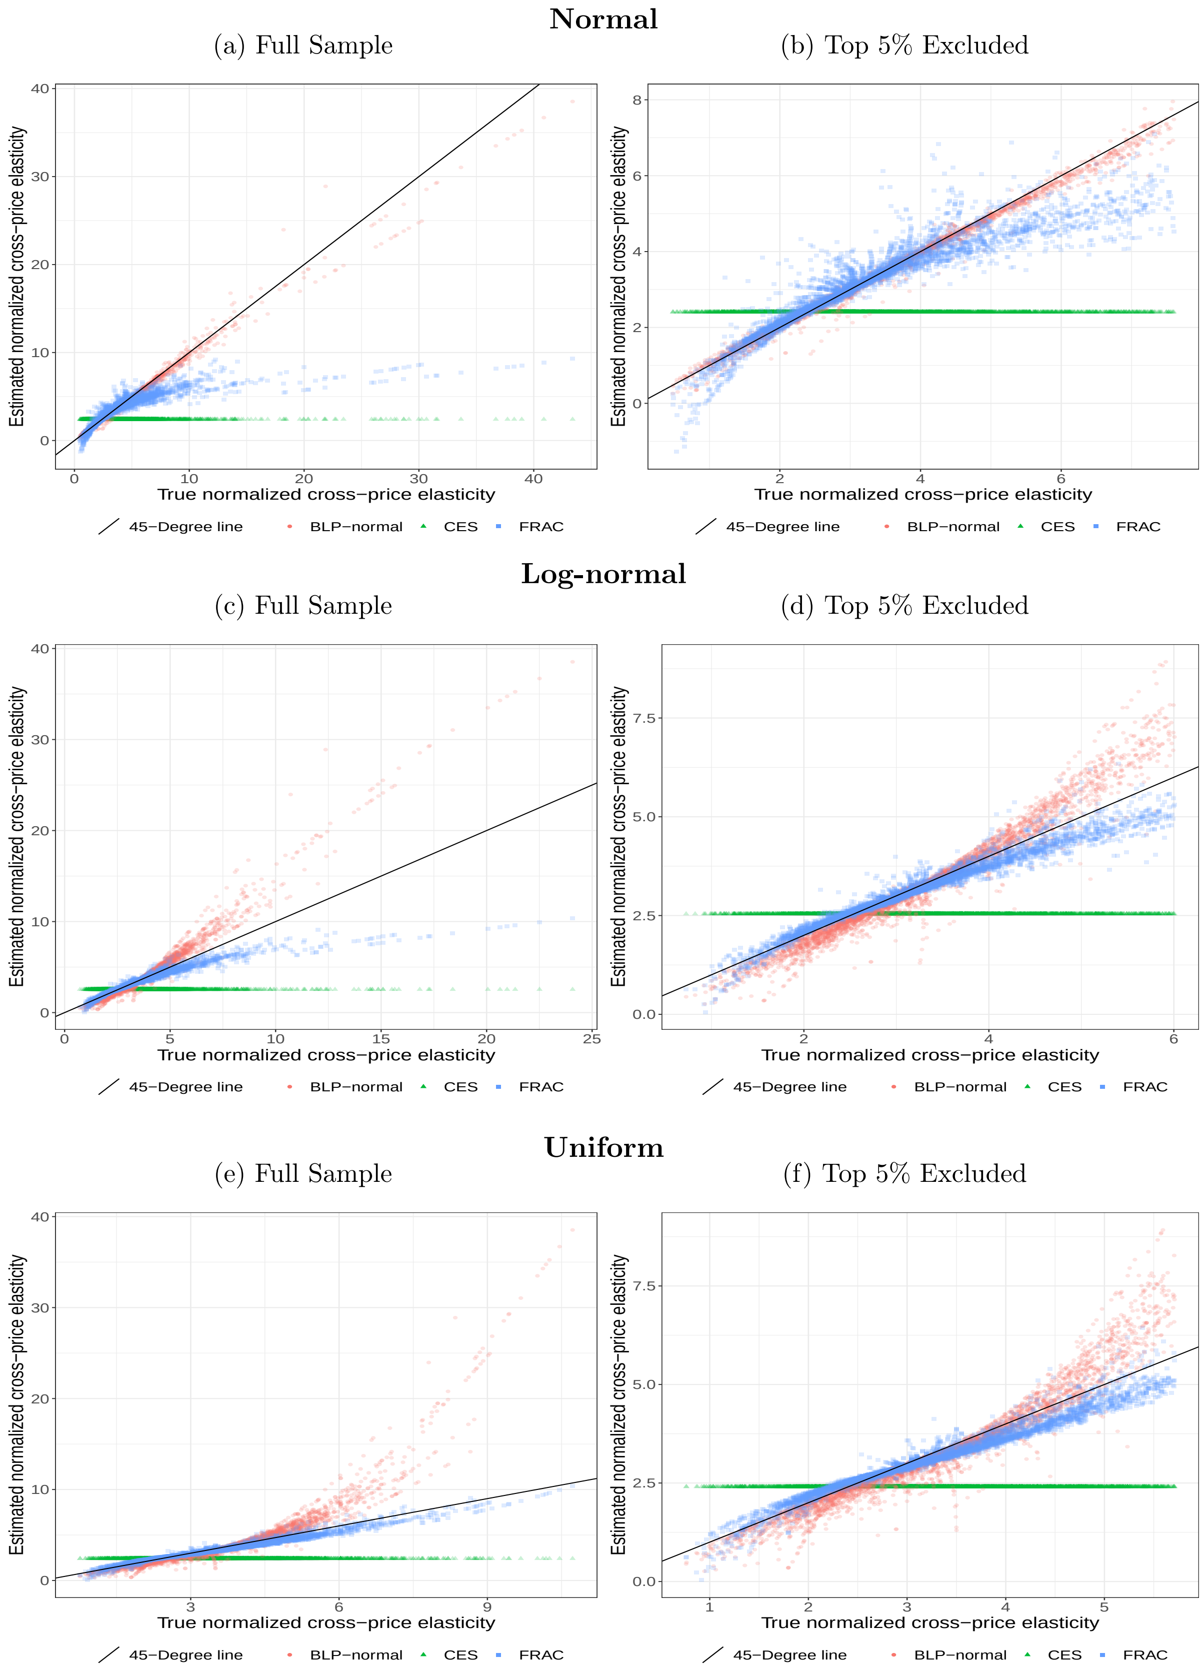


FIGURE 8: Normalized Cross-price Elasticity – Truth versus Estimates

This graph plots the normalized cross-price elasticity $e_{o,c} \equiv \frac{1}{s_c} \frac{\partial \ln s_o}{\partial \ln p_c}$ of variety $o$ with respect to a reference variety $c \neq o$. The x-axis is the true elasticity while the y-axis is the estimated elasticity, obtained either by CES, FRAC or BLP-normal. The left column of panels reports the full sample. In the right column, the top 5% largest true elasticities are excluded. In panels (a) and (b) the data is simulated assuming that random coefficients are normal. Panels (c), (d) and (e), (f) are respectively simulated with log-normal and uniform random coefficients.