Product Entry in the Global Automobile Industry

Alejandro Sabal
Yale University*

March 25, 2026

Abstract

This paper examines how national policies influence global market outcomes through firms' product portfolio decisions. I develop a method to estimate and solve a product entry game with multiple asymmetric firms making discrete choices, overcoming challenges arising from multiple equilibria, computational infeasibility, and selection on unobservables by using novel inequalities. Using data on global passenger vehicle sales, prices, and characteristics, I estimate large product portfolio fixed costs, which incentivize firms to offer similar products across markets to achieve scale. I analyze the effects of production and consumption subsidies favoring U.S. brands, and find that product entry exacerbates profit-shifting towards protected brands and that heterogeneity in consumer preferences plays a key role in shaping the global repercussions of policies.

Keywords: discrete choice methods, moment inequalities, oligopoly, entry games, global firms, automobile industry

*I am grateful to Gene M. Grossman, Kate Ho, and Eduardo Morales for their invaluable guidance and support. I am grateful to discussants Felix Montag, Dennis Novy, and Charlie Porcher. I would like to thank Elena T. Aguilar, Nicholas Buchholz, Juanma Castro-Vincenzi, Francesco Fabbri, Allan Hsiao, Michael Jenuwine, Adam Kapor, Jakub Kastl, Hugo Lhuillier, Alessandro Lizzeri, Simon Margolin, Paola Moscariello, George Nikolakoudis, Ezra Oberfield, Stephen Redding, Richard Rogerson, Carol Shou, and Martin Vaeth for detailed comments. This research was supported by the International Economics Section (IES) at Princeton University, Fundación Ramón Areces, and Cowles Foundation at Yale University. Additional materials for this paper can be found at: https://github.com/sabalalejandro/sabal/blob/main/papers/product_entry_auto/additional_materials.pdf. All errors are my own. This version of the paper: March 25, 2026. Email: Alejandro Sabal (alejandro.sabal@yale.edu).

1 Introduction

Product mix decisions account for important changes in the market structure of the automobile industry. In the United States, the number of firms selling at least one product barely changed from 1980 to 2018, while product offerings doubled (Grieco et al. 2023). Policy debates now focus on how major subsidies, such as those in the Inflation Reduction Act or European Green Deal, will reshape the industry and affect market outcomes globally. Yet, despite product entry's importance in an industry subject to significant government intervention, its role in shaping national policy effects on global markets remains little understood.

This paper examines how national policies shape firms' global product offerings and market outcomes. Firms choose portfolios to maximize profits given heterogeneous global preferences and policies, making interdependent choices: developing costly differentiated products encourages multi-country sales to achieve scale, while additional product offerings can steal business from own and rival products. To deal with a model incorporating these features tractably, I derive novel inequalities that bound entry probabilities, enabling tractable estimation and solution of the model under various policy scenarios.

The framework I propose allows for heterogeneity in preferences across product characteristics as in Berry et al. (1995) and incorporates both product portfolio and market entry decisions. The model features a finite set of firms, each endowed with a set of potential products. Firms make portfolio and market entry decisions in a sequential game with three stages. In the first stage, firms choose which subset of their potential products to include in their global product portfolio, subject to a fixed cost per product line. In the second stage, firms choose which subset of their portfolio to offer in each market, subject to market entry fixed costs. In the third and final stage, firms set prices for each product in each market.

Methodologically, the key contribution of this paper is to show how to estimate fixed costs and compute the impact of counterfactual policies in entry games with multiple asymmetric firms, each making multiple discrete choices. Three main challenges arise in such settings. First, the existence and uniqueness of Nash equilibria are not guaranteed. Second, solving for the equilibria of the model may be computationally infeasible, particularly in settings with a large number of heterogeneous players, each with a large set of potential choices, as in the automobile industry. Finally, the third complication is to allow for selection on the unobserved (to the researcher) component of firms' fixed cost shocks.

I jointly overcome such challenges, both for estimation and for computing counterfactual exercises, using novel inequalities that bound the probabilities of firms' portfolio and market entry decisions. The inequalities are derived under two assumptions: unobserved rival fixed

cost shocks and submodularity of variable profits with respect to product offerings. A benefit of the first property is that it guarantees the existence of a pure strategy Nash equilibrium, which is not ensured with complete information. The second property implies that the change in variable profits from offering a product in a market declines with the set of offered products.[1] I combine these two familiar assumptions to develop a method to estimate and solve oligopolistic multi-product entry games.

I apply submodularity to derive bounds on the gains from offering a product in a market, which I use to derive necessary conditions for entry that hold across all equilibria and overcome the selection problem. Fan and Yang (2025) also employ submodularity to derive necessary conditions for entry, though in a complete rather than incomplete-information framework. Such conditions permit integrating the unobserved component of the fixed costs, under any assumed distribution, to obtain bounds on firms' choice probabilities that depend on the fixed cost parameters and expectations over rivals' actions.[2]

de Paula (2013) reviews the literature on the econometric analysis of games with multiple equilbria. As in Ciliberto and Tamer (2009), I derive a novel set of moment inequlities that are valid under equilibrium multiplicity. Compared to their approach, my method does not rely on computing the equilibria of the model under any set of parameters, making it computationally feasible even in high-dimensional discrete settings.[3] I derive my moment inequalities by further bounding firms' choice probabilities in two steps. First, I use convex upper and concave lower bounds of the cumulative distribution function (CDF) of fixed costs. This relates to Dickstein and Morales (2018), Dickstein et al. (2024), and Porcher et al. (2024), who use convex odds functions or linear approximations to derive moment inequalities in a single-agent setting, though I directly bound choice probabilities in a game. Second, I apply Jensen's inequality to construct moments that depend on the observed realization of rivals' offerings. This extends the insights from Pakes (2010) and Dickstein and Morales (2018) to an incomplete-information game, where firms' expectations are over rivals' endogenous entry decisions rather than over exogenous ex-post realized variables.

To evaluate the effects of policies given estimated parameters, I develop a tractable solution algorithm for multi-product entry games, based on the derived bounds on firms' choice probabilities. My solution algorithm has two main advantages relative to existing approaches that assume complete information on firms' fixed cost shocks. First, it provides

[1] The change in variable profits from offering a product in a given market when the set of products offered in the market is Ω is smaller than when the set of products offered in that market is $\Omega' \subseteq \Omega$.

[2] Submodularity is stronger than required for estimation, but plays a more critical role in the solution method.

[3] With only 10 firms each making 10 discrete choices, there are $2^{100} \approx 10^{30}$ possible entry configurations. It is infeasible to compute all of them.

bounds on the equilibrium distribution of entry decisions even in oligopolistic settings with more than two firms.[4] Second, my solution approach does not rely on approximation methods or equilibrium selection assumptions and bounds any equilibrium of the entry game. I find informative bounds on counterfactual outcomes in practice.

The solution method starts by evaluating changes in profits for each product at the most competitive conditions (all potential products are simultaneously offered) and least competitive conditions (no other potential product is offered). This yields weak upper and lower bounds on the probability that any product is offered in a market. I then show how to simulate tighter upper bounds by evaluating the change in expected profits from offering a product using the initialized *lower* bounds on rival product offering probabilities. Similarly, I simulate tighter lower bounds by integrating rivals' entry decisions using the initialized *upper* bound rival offering probabilities. I prove that iterating on this procedure yields monotonically tighter bounds on the joint equilibrium distribution of offerings in each country, which I use to compute bounds on market outcomes such as consumer surplus.[5]

I estimate the model with 2019 *IHS Markit* data on the universe of new passenger vehicle registrations in a representative set of countries. I complement this main dataset with gravity variables from CEPII; PPP and Gini data from the World Bank; and the MRI Simmons 2019 U.S. Crosstab Report. The latter provides information relating vehicle characteristics to buyers' demographics. With these cross-sectional data, I estimate a heterogeneous agent demand model similar to Berry et al. (1995) for passenger vehicles using micro-moments as in Petrin (2002).

Demand and marginal cost estimates yield average own-price elasticities and markups consistent with previous estimates in the automotive industry (i.e., Berry et al. 1995, Grieco et al. 2023). I find substantial heterogeneity in price sensitivity across the income distribution and a considerable home market bias, as in Coşar et al. (2017); consumers are willing to pay on average over $1500 to purchase from a local brand, all else equal. My marginal cost estimates reveal a positive relationship between cost and distance to the brand's headquarters country, as well as larger costs of producing larger or more powerful vehicles.

Using my method, I obtain fixed cost estimates that reveal important economies of scale at the product level. The estimates show that, with 95% confidence, the median product-line fixed cost is $138–548 million, versus $8–15 million for market entered per product.[6] The

[4] Jia (2008) also uses a procedure to bound the Nash equilibria in a duopolistic entry game, but the procedure in that paper is not valid when there are more than two players.

[5] The bounds of the equilibrium distribution of product offerings in each market are bounds in the sense of first-order stochastic dominance.

[6] My estimates align with industry (IHS Global) estimates of product development costs in the car industry.

large gap between product and market entry costs reflects scale economies, incentivizing firms to offer similar product bundles across markets. I also estimate variances in fixed cost distributions that reveal substantial unobserved heterogeneity, underscoring the need to account for factors known to firms but unobserved by researchers.

I use my solution algorithm to assess national policy effects on global consumer and firm outcomes through two experiments. First, echoing large-scale packages like the Inflation Reduction Act, I study a 20% marginal cost reduction for U.S. brands.[7] Second, reflecting rising protectionist consumption subsidies, I analyze a 50% U.S. consumption subsidy for U.S. brands. Several countries have implemented consumer-side subsidies of similar magnitudes to support specific vehicle types.[8] These experiments disentangle how product entry shapes global effects of various types of policies, particularly when implemented in large markets like the United States.

The marginal cost subsidy affects global outcomes along both the intensive margin (prices and quantities) and the extensive margin (product offerings). It increases American brand dominance worldwide, raising their market shares by at least 2.6 percentage points in Japan and over 13.7 in the United Kingdom. American brands' variable profits also rise across all markets, nearly tripling in several. In contrast, non-American brands experience declines in both profits and market shares due to intensified competition.

A key finding is that product portfolio adjustments amplify these effects. American brands expand their portfolios anticipating higher profits, while non-U.S. brands reduce theirs, expecting lower competitiveness. Ignoring this adjustment would significantly understate the rise in U.S. brand shares and profits. Across most markets, entry accounts for over 25% of the increase in both the lower and upper bound of U.S.-brand market shares. Thus, accounting for portfolio responses is crucial for understanding firm-level outcomes.

Consumers worldwide benefit from access to cheaper American products. However, portfolio responses do not account for much of the rise in consumer surplus, as benefits from U.S. entry are offset by losses from non-U.S. exits. Surplus gains vary by country: consumers in Brazil and Mexico, who are more price-sensitive, see increases of over 10%. American consumers value favored brands more due to home bias, and thus benefit more than their counterparts in other rich nations.

Unlike the production subsidy, the U.S. consumption subsidy does not alter cost nor preference fundamentals abroad. It raises U.S. consumer surplus by over 43% and shifts

[7] I do not model the production location decisions of firms, so I reduce the marginal cost of all products produced by brands headquartered in the United States.

[8] For instance, in China, consumer subsidies on electric vehicles peaked at around 40–60% in 2014, according to the Environmental Energy Study Institute (Lu 2018). I show quantitatively that even a consumer subsidy this large leads to small cross-country effects.

U.S. brand shares and profits upward in the United States, with limited effects elsewhere. While the policy induces entry of U.S.-branded products both domestically and abroad, the new products are unpopular in unsubsidized markets overseas and therefore have little effect on foreign market shares.

Three main takeaways emerge. First, both policies raise consumer surplus overall, with beneficial product entry by favored brands outweighing the losses from disadvantaged brands' product exit. Second, entry magnifies profit shifting toward advantaged brands. Third, despite affecting global product offerings, U.S. protectionist consumer subsidies mainly have local effects on market structure, while production subsidies drive global shifts in firm and consumer outcomes.

This paper relates to previous work studying the effects of government policies in the automobile industry. My framework contributes to Berry et al. (1995), Goldberg (1995), Petrin (2002), Coşar et al. (2017), Grieco et al. (2023), and Allcott et al. (2024) by providing a model of product entry in the international automobile industry to study the cross-market effects of national policies on global consumer and firm-level outcomes. While such papers focus on pricing behavior and market power in the passenger vehicle industry, I study international product entry. My methodological contribution to this literature is to show how to estimate and solve a product entry game while accounting for horizontal and vertical product differentiation, heterogeneous consumer preferences, and strategic pricing.

I develop a multi-product entry model related to Eizenberg (2014), Wollmann (2018), Bontemps et al. (2023), Fan and Yang (2025), and Montag (2024), which also define potential entry opportunities and use moment inequalities to estimate fixed costs. My paper differs in two key dimensons. First, I model entry in a global setting, allowing firms to develop products and exploit scale economies by selling them in multiple markets. Second, I assume unobserved rival fixed cost shocks rather than complete information, enabling tractable counterfactuals without low-dimensional restrictions, approximations, or equilibrium selection rules. Like Fan and Yang (2025) but unlike Eizenberg (2014) and Wollmann (2018), my moment inequalities impose no support restrictions and identify both the mean and the variance of fixed costs, while requiring a distributional assumption.

This paper contributes to the international trade literature on interdependent global firm decisions (e.g., Tintelnot 2016; Antràs et al. 2017; Morales et al. 2019; Head and Mayer 2019; Alfaro-Urena et al. 2023; Castro-Vincenzi 2024; Castro-Vincenzi et al. 2024; Head et al. 2024) with two main differences. First, I study global product portfolio choices rather than production, sourcing, or dynamic entry. Second, in my model, firms are strategic rather than atomistic. The model incorporates mechanisms from prior work in trade—scale economies (Krugman 1980; Venables 1987; Thomas 2011; Costinot et al. 2019), oligopolies

(Atkeson and Burstein 2008), and multi-product firms (Bernard et al. 2011; Mayer et al. 2021)—but provides a quantitative framework for policy analysis with multi-product firms, strategic behavior, and product entry. Unlike much of this literature, I allow for heterogeneous consumers and products, flexible substitution patterns, and strategic behavior, and hold general equilibrium variables fixed to focus on industry equilibria.

Methodologically, this paper contributes to the literature on solution methods in settings with multiple interdependent discrete choices. In single-agent settings, Arkolakis et al. (2023), Alfaro-Urena et al. (2023), and Castro-Vincenzi et al. (2024) provide solution algorithms that exploit knowledge of complementarity or substitutability across discrete choices. In a strategic setting, Seim (2006) solves an incomplete-information game with ex-ante symmetric firms. Jia (2008) provides an algorithm to solve a complete-information entry game with two asymmetric firms making multiple entry decisions. My contribution is to provide a method to bound any equilibrium of an entry game featuring more than two asymmetric firms, each making multiple discrete choices.

The remainder of the paper is organized as follows. Section 2 overviews the data and the industry setting. Section 3 develops a model of multi-product and multi-market entry. Section 4 provides bounds on firms' choice probabilities and explains how to use them to bound the fixed cost parameters. Section 5 reports estimation results. Section 6 develops a solution algorithm to bound the equilibrium distribution of product offerings in each market. Section 7 evaluates the effects of U.S. policies favoring domestic brands on global market outcomes. Section 8 concludes.

2 Data and Industry Setting

My primary source of data is information on the universe of new passenger vehicle registrations in the year 2019 in 12 countries: Australia, Brazil, China, France, Germany, India, Italy, Japan, Mexico, Spain, the United Kingdom, and the United States. I obtain these data from *IHS Markit*.[9] The data contain manufacturer-suggested retail prices (MSRP) in USD, units sold, and other characteristics such as fuel type, body type, horsepower, length, height, width, and weight. The data are at the quarterly-trim-country level and include 177 different brands and 73 different parent companies.[10] I merge this dataset with the brand's headquarters country and with information from CEPII on the distance between

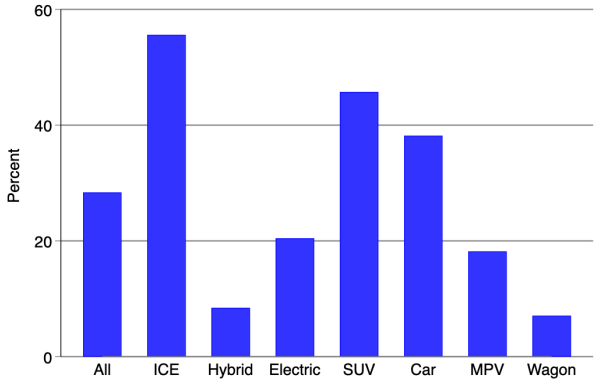
[9] Previous work has used versions of these data, e.g., Coşar et al. (2017) and Head and Mayer (2019).

[10] A trim is a definition used by manufacturers to identify a vehicle's special features and level of equipment at a finer level than a nameplate. An example of a nameplate is a Toyota Corolla. Within the nameplate, there is typically differentiation across trims, for instance, the Toyota Corolla ZR or the Toyota Corolla Ascent, which may have different characteristics like horsepower.

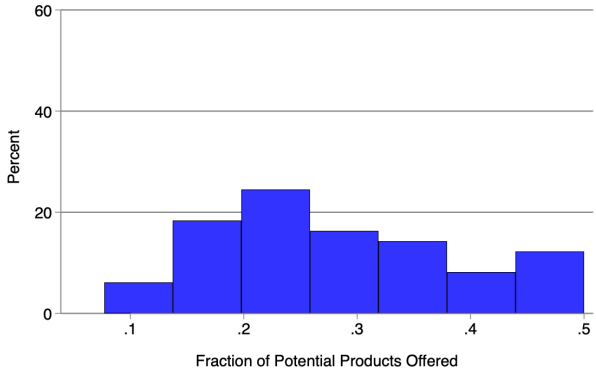
the destination country and the brand's headquarters country.

Figure 1: Fraction of Potential Products Offered Globally

Panel A: Across Product Categories



Panel B: Across Firms



Note: In Panel A, the "All" category reports the percentage of potential products (potential brand-body type-fuel type combinations) that are sold in at least one market. All other categories report the fraction of potential products of that category that are sold in at least one market. "ICE" stands for internal combustion engine and "MPV" stands for multi-purpose vehicle. Panel B reports the distribution of the fraction of potential products offered in at least one market across parent companies (e.g., Ford).

To obtain product market shares at the year-country level, I divide units sold in 2019 by market size. I follow Grieco et al. (2023) and define market size as the product of the number of households and the average number of vehicles per household in each country in 2019, divided by the average tenure of car ownership, which is assumed to be 5 years.

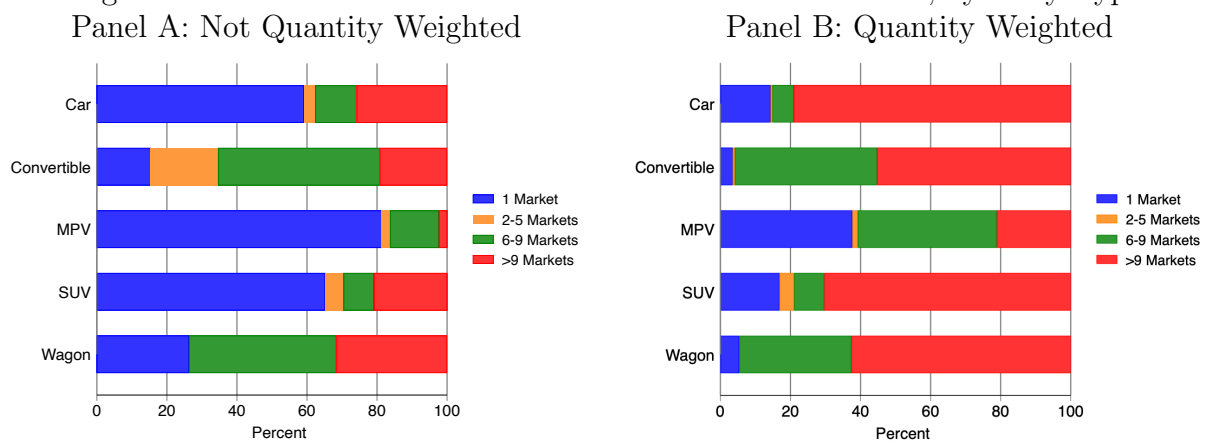
I also obtain micro-moments relating the income of vehicle buyers to the prices of vehicles they purchase from the MRI-Simmons 2019 Crosstab Report for the United States. I obtain market-level data on PPP income per capita and Gini coefficients from the World Bank. Assuming income follows a log-normal distribution in each country, income per capita and the Gini coefficient yield the country-specific location and scale parameters.

I refer to parent companies (e.g., Ford) as firms. Each firm has a set of brands (e.g., Ford, Lincoln). I define a product as a brand-body type-fuel type combination. The set of possible body types includes Cars, SUVs, wagons, multi-purpose vehicles (MPVs), and convertibles.[11] Possible fuel types are internal combustion engine (ICE), (plug-in) electric, or hybrid. For instance, a potential product is a Lincoln SUV with an ICE.[12]

[11] As stated in Online Appendix A, I define a Car to be a sedan, hatchback, or coupe.

[12] This is a different level of aggregation than in most of the literature on automobiles, which usually aggregates across trims at the brand-nameplate level (e.g., Head and Mayer 2019, Grieco et al. 2023), e.g., Toyota Corolla. My product definition permits the extrapolation of the full set of product characteristics to all potential products, whereas it is not possible to take a stance on a brand's potential nameplates.

Figure 2: Number of Markets Offered Conditional on Portfolio, by Body Type



Note: Both Panel A and Panel B condition on the products that are observed to be offered in at least one market in the data. For each body type, Panel A reports the percentage of products within that body type that are offered in different number-of-market categories. Panel B reports the quantity share of each number-of-market category across body type categories.

I aggregate the remaining characteristics of these products (e.g., size or horsepower) by taking a quantity-weighted average across all trims within this category. In Online Appendix A, I discuss the procedure used to impute product characteristics for potential products that I do not observe in my sample. After dropping brands specializing in luxury products, there are 49 parent companies, 130 brands, and 1530 potential products.[13]

Figure 1 reports the fraction of potential products that firms sell in at least one market and how it varies across product categories (Panel A) and across firms (Panel B). Firms offer fewer than 30% of their potential products across all markets. Across the firm distribution, no firm offers more than 50% of its potential products.

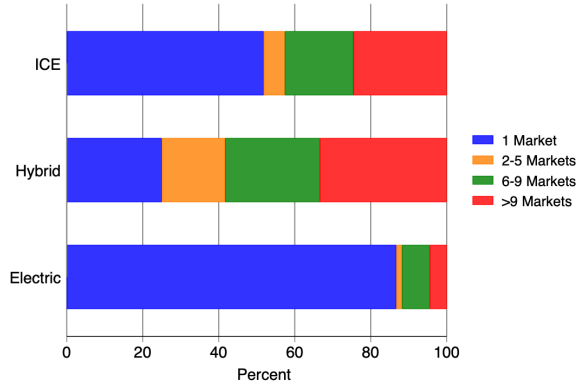
Figure 1 also shows that global product offerings vary substantially across product categories, which suggests a response to demand conditions that vary across product types. While more than 40% of potential SUV products are offered in at least one market, fewer than 10% of potential hybrid or wagon products are offered.

Figures 2 and 3 condition on the set of products sold in at least one market and show that across all body types and fuel types, a substantial fraction of products are sold in a single market. For instance, around 60% of SUVs or Cars are sold in a single market. Still, a considerable fraction of such products (for these categories, around 20%) is sold in over 9 markets. There is also significant heterogeneity in entry patterns across product categories,

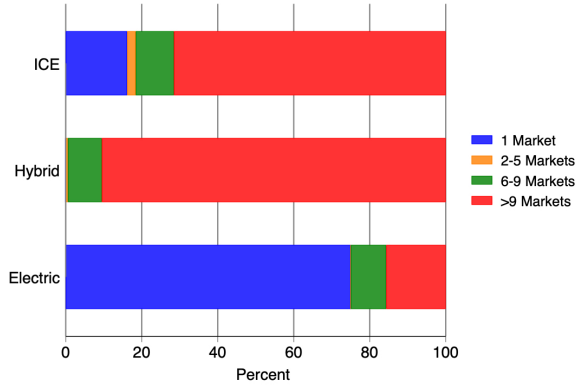
[13]Brands like Ferrari, Maserati, Lamborghini, or Rolls-Royce have small quantity shares and substantially higher prices across markets. Thus, I excluded them from the estimation sample to avoid issues with small market shares and because they belong to a substantially different segment of the market.

Figure 3: Number of Markets Offered Conditional on Portfolio, by Fuel Type

Panel A: Not Quantity Weighted



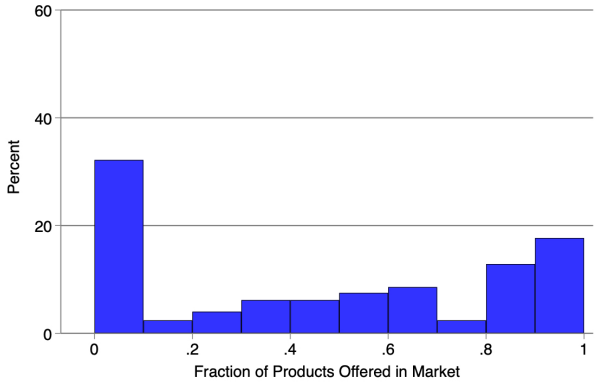
Panel B: Quantity Weighted



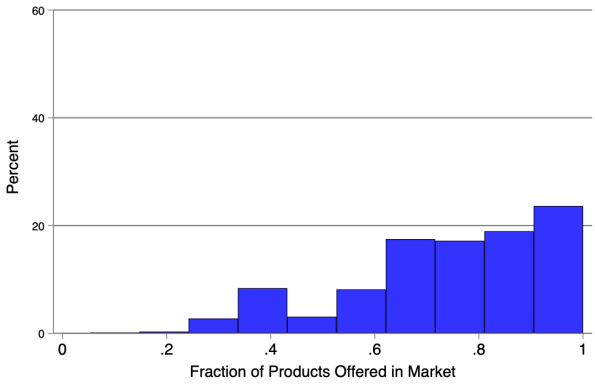
Note: Both Panel A and Panel B condition on the products that are observed to be offered in at least one market in the data. For each fuel type, Panel A reports the percentage of products within that fuel type that are offered in different number-of-market categories. Panel B reports the quantity share of each number-of-market category across fuel type categories.

Figure 4: Fraction of Products Offered

Panel A: Not Quantity Weighted



Panel B: Quantity Weighted



Note: Panel A plots the distribution of the fraction of products offered across firm-market pairs. Panel B weights by the total units sold by the firm in the market.

with the most extreme example being electric vehicles, the majority of which are sold in a single market. The right panel in Figure 2 shows that weighting by quantity changes these statistics. Products offered in over 9 markets account for the vast majority of units sold (70-80% for Cars and SUVs). This shows that product market shares differ substantially, with some capturing a large share of world demand.

Finally, automotive firms not only serve multiple markets but also offer a portfolio of multiple products. Both Panel A and Panel B in Figure 4 show that, in most cases, firms

sell a strict subset of their products in each market. Panel A shows that only in around 30% of firm-markets there is near-zero entry. Moreover, only in 17.7% (23.6% if quantity-weighted) of the cases do firms sell all their products in a market.

3 Model of Strategic Global Multi-Product Offerings

In this section, I develop a model of strategic product portfolio choice and market entry. A set $\mathcal{F} = \{1, ..., f, ..., F\}$ of firms competes in a set $\mathcal{M} = \{1, ..., m, ..., M\}$ of markets.[14] A given firm f is endowed with a set of potential products. I denote a firm's potential set of products as $\mathcal{A}^f$. I denote the union of all potential products across firms by $\mathcal{A}$.

I model firms' portfolio and market entry choices in three stages. In Stage 1, each firm f realizes a set of product portfolio fixed cost shocks $\{\nu_j^g\}_{j \in \mathcal{A}^f}$ for each of their potential products, and upon observing these, chooses which products to introduce in its global product portfolio, $\mathcal{G}^f$. In Stage 2, each firm f realizes a set of market entry fixed cost shocks $\{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}$ and chooses which subset of products $\Omega_m^f \subseteq \mathcal{G}^f$ in its portfolio to offer in each potential market m.[15] Finally, in Stage 3, each firm realizes a set of demand and marginal cost shocks for each product offered in each market and chooses which prices to charge. Firms solve the model by backward induction.

A key assumption is that throughout Stages 1 and 2, firms do not observe their rivals' fixed cost shocks when making their entry decisions. Thus, in Stages 1 and 2, firms play a Bayesian Nash equilibrium in product entry decisions.

Assumption 1 (Unobservability of Rival Fixed Cost Shocks) *Firms' fixed cost shocks* $\{\nu_j^g\}_{j \in \mathcal{A}^f}$ *and* $\{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}$ *are private information at the time of making product portfolio and market entry decisions throughout Stages 1 and 2 of the game.*

Assumption 1 is a departure from previous papers such as Jia (2008), Ciliberto and Tamer (2009), Eizenberg (2014), Wollmann (2018), or Fan and Yang (2025), which assume that firms perfectly observe other firms' fixed cost shocks when making their entry decisions. This assumption guarantees the existence of a pure strategy equilibrium in both Stage 1 and Stage 2 of the game.[16] This is an advantage relative to complete information, where

[14] I hold the set of firms fixed due to evidence that product entry has been more important than firm entry in leading to changes in market structure in this industry from 1980-2018 (see Grieco et al. 2023). However, the model can be extended to allow for firm/brand entry provided their potential products' characteristics are defined.

[15] I lighten notation by referring to collections of variables across markets $\{Y_m\}_{m \in \mathcal{M}}$ as $\{Y_m\}_m$.

[16] The proof follows using Theorem 4 from Milgrom and Weber (1985) and Theorem 3.1 from Balder (1988).

the existence of a pure strategy equilibrium is not guaranteed in games featuring strategic substitutes with more than two players.[17]

I now describe the three stages of the game in reverse order.

3.1 Stage 3

3.1.1 Demand and Marginal Costs

The demand side of the model follows a mixed logit specification, as in Berry et al. (1995). Buyer i in market m chooses to purchase a new passenger vehicle j from the set of offerings Ω_m or the outside option so as to maximize utility. Within a market, buyers have idiosyncratic demographic characteristics, which determine how much they value the different attributes of each good, as well as their distaste for prices. The indirect utility that buyer i in market m derives from product j is given by,

$$U_{ijm} = \beta_m + \beta_{b(j)} - \alpha_i p_{jm} + \beta^x \boldsymbol{X}_{jm} + \xi_{jm} + \varepsilon_{ijm} \tag{1}$$

$$= \tilde{\delta}_{jm} - \alpha_i p_{jm} + \xi_{jm} + \varepsilon_{ijm}. \tag{2}$$

In equation (1), the vector $\boldsymbol{X}_{jm}$ denotes a set of non-price attributes of product $j \in \Omega_m$ or brand-market characteristics. I allow buyers to have different preferences across markets over (a subset of) these characteristics. The vector $\boldsymbol{X}_{jm}$ includes market identifiers interacted with (i) a dummy denoting whether the product is electric or hybrid, (ii) a set of dummies for different body type categories, and (iii) size. It also includes horsepower/weight, horsepower, and a home market dummy denoting whether the brand's headquarters are located in market m.[18] $\beta_{b(j)}$ is a brand fixed effect, and ξ_{jm} is a product-market demand shock that is realized at this stage of the game and, subsequently, perfectly observed by all buyers and firms. The variable $\tilde{\delta}_{jm}$ denotes the mean non-price utility of product j in market m net of the demand shock ξ_{jm}. Buyer i has a marginal utility of income $\alpha_i = \exp(\alpha_0 + \alpha_1 \log(income_i) + \alpha_2 China_i + \sigma^y u_i)$, where u_i are i.i.d. normal shocks. I allow distaste for prices to be different in China, conditional on income, in light of the unique policy environment that characterizes China during my sample period, which makes consumers seemingly less price sensitive. Finally, ε_{ijm} denotes an i.i.d. Type 1 Extreme Value distributed preference shock

[17] In Additional Materials B.1, I provide an example of a complete information game with strategic substitutes and 3 players where no pure strategy Nash equilibrium exists. Magnolfi and Roncoroni (2022) provide an identification approach for discrete games under a weaker informational assumption but do not deal with within-firm interdependencies and their approach is not feasible to implement in high-dimensional settings. Additional materials can be found at the following link: https://github.com/sabalalejandro/sabal/blob/main/papers/product_entry_auto/additional_materials.pdf.

[18] Horsepower/weight proxies inversely for fuel efficiency, which is an omitted variable in my dataset.

that is buyer-product-market specific. The mean utility of the outside good (good 0) is normalized to zero, such that $U_{i0m} = \varepsilon_{i0m}$.

The above specification of indirect utility yields buyer-specific logit probabilities that can then be integrated over the joint distribution of market-specific buyer demographics and taste shocks to obtain market shares s_{jm}. Then, the quantity of product j sold in market m can be obtained by multiplying the market share s_{jm} by the market size M_m.

The marginal cost of supplying a unit of product j in market m is given by,

$$\begin{aligned}\log(c_{jm}) &= \gamma_m + \gamma_{b(j)} + \gamma_{fueltype(j)} + \gamma_1 \log(hp_j) + \gamma_2 \log(hpw_j) \\ &\quad + \gamma_3 \log(size_j) + \gamma_4 \log(dist_{jm}) + \omega_{jm} \\ &= \tilde{c}_{jm} + \omega_{jm} \end{aligned} \tag{3}$$

where hp and hpw denote horsepower and horsepower/weight, respectively, and $dist_{jm}$ denotes the distance between market m and the headquarters country of brand b. The term γ_m is a market fixed effect and $\gamma_{b(j)}$ is a brand fixed effect. The variable $\tilde{c}_{jm}$ denotes the mean (log) marginal cost of product j in market m net of the product-market marginal cost shock ω_{jm}. Marginal costs, therefore, depend on observed characteristics and a marginal cost shock ω_{jm} that is realized at this stage of the game.

Demand and marginal cost shocks $\{\xi_{jm}, \omega_{jm}\}_{j \in \Omega_m, m}$ are i.i.d. and I assume they follow a bivariate normal distribution.[19]

3.1.2 Pricing

Firms set prices according to a complete-information, Nash-Bertrand pricing game in each market m, as in Berry et al. (1995).

When firms choose prices, demand and marginal cost shocks $\{\xi_{jm}, \omega_{jm}\}_{j \in \Omega_m}$ are known by all firms for all products Ω_m offered in market m. Each firm chooses its prices to maximize its variable profits, given by,

$$\begin{aligned}\Pi_m^{f,3} &= \sum_{j \in \Omega_m^f} (p_{jm} - c_{jm}) M_m s_{jm}(\Omega_m^f, \Omega_m^{-f}, \boldsymbol{p}_m^f, \boldsymbol{p}_m^{-f}, \boldsymbol{\xi}_m) \\ &= \sum_{j \in \Omega_m^f} \pi_{jm}(\Omega_m^f, \Omega_m^{-f}, \boldsymbol{p}_m^f, \boldsymbol{p}_m^{-f}, \boldsymbol{\xi}_m, \boldsymbol{\omega}_m).\end{aligned}$$

I use boldface notation to denote vectors of variables. The variable $\boldsymbol{p}_m^f$ is the vector of

[19] Other papers have made alternative assumptions on the distribution of these shocks. For instance, Wollmann (2018) draws from the empirical distribution of these shocks when computing counterfactual experiments.

prices in market m across products Ω_m^f offered by firm f, $\boldsymbol{p}_m^{-f}$ denotes the prices charged for products Ω_m^{-f} offered by firms other than firm f, and $(\boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$ stand for the demand and marginal cost shocks for all products Ω_m offered in market m.

I denote equilibrium variable profits given demand and marginal cost shocks, and product offerings Ω_m, by,

$$\pi_{jm}^*(\Omega_m^f, \Omega_m^{-f}, \boldsymbol{\xi}_m, \boldsymbol{\omega}_m) := \pi_{jm}(\Omega_m^f, \Omega_m^{-f}, \boldsymbol{p}_m^{f,*}, \boldsymbol{p}_m^{-f,*}, \boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$$

where $\boldsymbol{p}_m^{f,*}$ denotes the equilibrium price vector that emerges in market m at $(\Omega_m^f, \Omega_m^{-f})$ given demand and marginal cost shocks.[20]

3.2 Stage 2: Market Entry Decisions

At this stage, firms have chosen their global product portfolio $\mathcal{G}^f$ and realize a set of market entry fixed cost shocks $\{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}$ for each product in their portfolio in each market. Firms make market entry decisions conditional on what they know; the information set of firm f in Stage 2 is, $\mathcal{I}^{f,2} := (\mathcal{G}^f, \{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}, \mathcal{I})$, where set $\mathcal{I}$ denotes the component of the information set that is common knowledge to all firms and is given by, $\mathcal{I} := (\{\tilde{\delta}_{jm}\}_{j \in \mathcal{A}, m}, \{\tilde{c}_{jm}\}_{j \in \mathcal{A}, m}, \mathcal{A}^f, \mathcal{A}^{-f})$. Set $\mathcal{I}$ also includes knowledge of the distributions of unobserved demand, marginal cost, and fixed cost shocks, policies in the counterfactual, and the equilibrium strategies that firms play.[21]

Given this information, each firm f chooses offerings in each market Ω_m^f so as to maximize expected profits. That is, it solves,

$$\Pi_m^{f,2}(\mathcal{I}^{f,2}) = \max_{\Omega_m^f \subseteq \mathcal{G}^f} \sum_{j \in \mathcal{G}^f} O_{jm}[\mathbb{E}[\pi_{jm}^*(\Omega_m^f, \Omega_m^{-f}, \boldsymbol{\xi}_m, \boldsymbol{\omega}_m)|\mathcal{I}] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e)] \tag{4}$$

where $O_{jm} = 1$ if and only if $j \in \Omega_m^f$. The variable $F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e)$ denotes the market entry fixed cost of offering product j in market m. Assumption 2 makes a distributional assumption for these fixed costs.

Assumption 2 (Market Entry Fixed Costs) *The fixed cost of offering product $j \in \mathcal{G}^f$ in market m is given by, $F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) = \exp(Z_{jm}'\theta_e + \sigma_e \nu_{jm}^e)$, where Z_{jm} is $\mathcal{I}$-measurable and $\nu_{jm}^e | \mathcal{I} \sim_{i.i.d.} Normal(0,1)$.*

[20] Without random coefficients, Nocke and Schutz (2018) show that the pricing equilibrium is unique. To the best of my knowledge, there is no proof for the case with random coefficients on prices. I apply the Morrow and Skerlos (2011) approach to reliably find an equilibrium, following the analysis in Conlon and Gortmaker (2020). The focus of this paper is on dealing with equilibrium multiplicity in the entry, rather than the pricing game.

[21] In summary, $\mathcal{I}$ denotes knowledge of the data-generating process.

In the empirical implementation, I assume that Z_{jm} is a constant.[22]

Firms choose which subset of their portfolio to offer in each market while best responding to other firms' product entry strategies and taking into account cannibalization across their products. The expectation in equation (4) is with respect to the distribution of $(\boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$ that are realized in Stage 3 and the distributions of $\boldsymbol{\nu}^{-f,e}$ and $\boldsymbol{\nu}^{-f,g}$ determining rival firms' offerings. Due to Assumption 1, only the information in the common part of the information set $\mathcal{I}$ is useful to predict variable profits, so I remove conditioning on the private information component in the conditional expectation in equation (4). When a firm chooses which products to offer in a market, it takes expectations over the market structure against which it will be competing and the markups it will be able to charge for each product it potentially sells in such a market. Such expectations are conditional on firms' knowledge of the policy environment subsumed in $\mathcal{I}$, which firms anticipate affects rival firms' offerings in equilibrium.

3.3 Stage 1: Global Portfolio Decision

In the first stage of the game, firms realize a set of product portfolio fixed cost shocks $\{\nu_j^g\}_{j\in\mathcal{A}^f}$ for each of their potential products. The information set of firm f is therefore, $\mathcal{I}^{f,1} := (\{\nu_j^g\}_{j\in\mathcal{A}^f}, \mathcal{I})$.

Each firm f chooses its product portfolio by maximizing expected profits,

$$\Pi^{f,1}(\mathcal{I}^{f,1}) = \max_{\mathcal{G}^f\subseteq\mathcal{A}^f} \mathbb{E}\Big[\sum_m \Pi_m^{f,2}(\mathcal{G}^f, \boldsymbol{\nu}_m^{f,e}, \mathcal{I})\big|\mathcal{I}\Big] - \sum_{j\in\mathcal{A}^f} G_j F_j^g(\nu_j^g; \theta_g, \sigma_g), \tag{5}$$

where $G_j = 1$ if and only if $j \in \mathcal{G}^f$.

The variable $F_j^g(\nu_j^g; \theta_g, \sigma_g)$ denotes the fixed cost of adding product j into the firm's product portfolio. Assumption 3 makes a distributional assumption for these fixed costs.

Assumption 3 (Portfolio Fixed Costs) *The fixed cost of introducing product $j \in \mathcal{A}^f$ into firm f's product portfolio is given by, $F_j^g(\nu_j^g; \theta_g, \sigma_g) = \exp(Z_j'\theta_g + \sigma_g\nu_j^g)$, where Z_j is $\mathcal{I}$-measurable and $\nu_j^g|\mathcal{I} \sim_{i.i.d.} Normal(0,1)$.*

In the empirical implementation, I assume that Z_j is a constant.

[22]Note that this specification rules out economies and diseconomies of scope, as the fixed costs are independent of the firm's decisions to offer other products. In Online Appendix F, I explain how the model and method could be extended to allow for such interdependencies through fixed costs and why estimates of such parameters are noisy given that there are relatively few instances in which a firm introduces zero products in a market.

Firms choose their portfolio so as to best respond to other firms' portfolio and market entry strategies. The portfolio decision affects the expected profits of firm f in all potential markets since including a product in its portfolio gives the firm the option to sell it in any country in the second stage. As in Stage 2, firm f's expectation is over other firms' portfolio and market entry shocks $\boldsymbol{\nu}^{-f,e}$ and $\boldsymbol{\nu}^{-f,g}$ determining rival firms' offerings in each market, and ex-post demand and marginal cost shocks. At this stage of the game, firm f also takes expectations over its own market entry fixed cost shocks $\boldsymbol{\nu}^{f,e}$, which it does not realize until Stage 2. Assumption 1 implies that only information contained in $\mathcal{I}$ is useful to predict market-level profits $\Pi_m^{f,2}$, so I remove conditioning on the private information component in the conditional expectation in equation (5). The product-specific fixed cost F_j^g induces a complementarity across markets. Greater expected profitability for product j in market m can lead to the introduction of product j into the firm's global product portfolio and its offering in other markets in Stage 2.

3.4 Marginal Value and Submodularity

I introduce the concept of the marginal value of offering a product in a given market.

Definition 1 *The marginal value of offering product j in market m at market structure Ω_m under demand and marginal cost shocks $(\boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$ is given by,*

$$\Delta_{jm}(\Omega_m^f, \Omega_m^{-f}; \boldsymbol{\xi}_m, \boldsymbol{\omega}_m) := \underbrace{\pi_{jm}^*(\Omega_m^f, \Omega_m^{-f}, \boldsymbol{\xi}_m, \boldsymbol{\omega}_m)}_{\textit{variable profits from } j \textit{ in } m} \tag{6}$$
$$+ \underbrace{\sum_{j' \neq j, j' \in \Omega_m^f} \left[\pi_{j'm}^*(\Omega_m^f, \Omega_m^{-f}, \boldsymbol{\xi}_m, \boldsymbol{\omega}_m) - \pi_{j'm}^*(\Omega_m^f \setminus \{j\}, \Omega_m^{-f}, \boldsymbol{\xi}_m, \boldsymbol{\omega}_m)\right]}_{\textit{cannibalization: change in variable profits for other products when } j \textit{ is offered}}.$$

It is the change in variable profits from offering product j in market m for firm f when the initial market structure is given by $\Omega_m^f \setminus \{j\}$ and under demand and marginal cost shocks in market m given by $(\boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$.

Throughout the remainder of the paper, I will leverage the following property.

Assumption 4 (Submodularity) *The marginal value $\Delta_{jm}(\Omega_m^f, \Omega_m^{-f}; \boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$ is monotone decreasing in Ω_m^f and Ω_m^{-f} with respect to the set-inclusion order at any demand and marginal cost shocks, $(\boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$.*

Assumption 4 is more likely to hold when (i) marginal costs are constant and when (ii) the unobserved shocks $(\boldsymbol{\xi}, \boldsymbol{\omega})$ are independent of firms' portfolio and market entry decisions. The

latter assumption is guaranteed due to the timing assumption on the demand and marginal cost shocks, which are realized in Stage 3. This timing assumption permits using standard techniques (i.e., Berry et al. 1995) to estimate demand and marginal costs.

As in Fan and Yang (2025), I find that at estimated parameters and under thousands of $(\boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$ draws there is no violation of Assumption 4, which is consistent with the economic forces present in the model.[23]

4 Bounds on Choice Probabilities and Moment Inequalities

I propose a set of moment inequalities to bound fixed-cost parameters $(\theta_e, \sigma_e, \theta_g, \sigma_g)$ that jointly overcomes the challenges of equilibrium multiplicity, computational infeasibility, and selection on unobserved (to the researcher) fixed cost shocks, under Assumptions 1-4.

First, I show how to derive bounds on firms' choice probabilities. Then, I show how to use these bounds to derive moment inequalities that partially identify the fixed cost parameters. Throughout, I assume that all demand and marginal cost parameters have been estimated. Together with the Nash-Bertrand pricing assumption in Stage 3, such parameters are sufficient to obtain variable profits for each product-market pair given any demand and marginal cost shock realizations $(\boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$ and market structures Ω_m. I then compute the marginal value of product j in market m, $\Delta_{jm}(\Omega_m^f, \Omega_m^{-f}; \boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$, under any offerings and demand and marginal cost shocks.

4.1 Bounds on Choice Probabilities

In this section, I derive the upper bound on the probability of market entry. In Online Appendix B, I provide the full derivations of the upper and lower bounds on the probabilities of market entry and product development following a similar recipe. Bounding firms' choice probabilities is the first step toward both deriving moment inequalities for estimation and deriving bounds on the entry equilibria.

Consider a firm in Stage 2 of the game that has chosen a product portfolio $\mathcal{G}^f$ at Stage 1. The derivation of the bounds on market entry probabilities proceeds in a sequence of steps. I begin by writing down a revealed-preference condition that must hold at any equilibrium. Let $\mathbb{1}_{\Omega_m^f}$ be an indicator denoting that firm f chooses to offer bundle Ω_m^f in market m. Then,

[23] Assumption 4 is a property of many of the payoff functions in the entry literature, including Ciliberto and Tamer (2009).

at any equilibrium,

$$(\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}})\, \mathbb{1}\big\{\mathbb{E}[\Delta_{jm}(\Omega_m^f, \Omega_m^{-f}) \mid \{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}, \mathcal{I}, \mathcal{G}^f] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0\big\}$$
$$= (\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}})\, \mathbb{1}_{\Omega_m^f}. \tag{7}$$

Equation (7) says that conditional on firm f choosing bundle Ω_m^f containing product j, or bundle $\Omega_m^f \setminus \{j\}$, then it chooses bundle Ω_m^f if and only if its expected profits conditional on the information set yield weakly greater profits.

Next, I deal with the selection problem caused by lack of knowledge of the distribution of ν_{jm}^e conditional on bundle Ω_m^f or $\Omega_f \setminus \{j\}$ being chosen by firm f. To deal with the selection problem, I first leverage Assumption 4 to obtain an upper bound for the marginal value of offering product j in market m, and write,

$$(\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}})\, \mathbb{1}\big\{\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f}) \mid \{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}, \mathcal{I}, \mathcal{G}^f] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0\big\}$$
$$\geq (\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}})\, \mathbb{1}_{\Omega_m^f}. \tag{8}$$

Importantly, in equation (8), the expected marginal value term inside the indicator function no longer depends on Ω_m^f. Summing all inequalities (8) across all bundles Ω_m^f containing product j and noting that the event that product j is offered in market m ($O_{jm} = 1$) is equivalent to the event that some bundle containing product j is offered, I obtain,

$$\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f}) | \mathcal{I}, \mathcal{G}^f, \{\nu_{km}^e\}_{k \in \mathcal{G}^f, m}] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0\} \geq O_{jm}. \tag{9}$$

In Online Appendix B, I similarly derive a lower bound inequality,

$$\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\mathcal{G}^f, \Omega_m^{-f}) | \mathcal{I}, \mathcal{G}^f, \{\nu_{km}^e\}_{k \in \mathcal{G}^f, m}] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) < 0\} \geq 1 - O_{jm}. \tag{10}$$

Inequalities (9) and (10) hold for *any* entry opportunity for product $j \in \mathcal{G}^f$ in market m (*independently* of which bundle Ω_m^f is actually chosen), thus overcoming the selection problem.[24] Inequality (9) says that if product j is offered in market m ($O_{jm} = 1$), necessarily an upper bound on the expected profit gain from doing so must be weakly positive. Under Assumption 4, I obtain an upper bound on the expected profit gain by evaluating the marginal value as if product j were the only product to be offered by the firm in the market, i.e., at

[24] Fan and Yang (2025) deal with the selection problem in a multiple discrete choice setting by bounding the probability of entry below by the probability that entry is dominant and above by the probability that entry is not dominated. They similarly implement their inequalities using the submodularity of variable profits, though in a setting of complete rather than incomplete information.

bundle $\{j\}$.[25] At this extreme bundle, there can be no cannibalization. Inequality (10) says that if j is not offered in market m ($O_{jm} = 0$), necessarily a lower bound on the expected gain from this choice must be negative. Assumption 4 says that the extreme bundle that minimizes such possible gains is $\mathcal{G}^f$, at which cannibalization is maximized. These bundles are not chosen optimally by the firm given its private information $\{\nu^e_{jm}\}_{j\in\mathcal{G}^f,m}$, which eliminates the selection problem generated by not knowing the distribution of ν^e_{jm} conditional on the observed and optimal bundle Ω^f_m.

The next step is to deal with the strategic nature of the problem. Under Assumption 1, knowledge of $\{\nu^e_{km}\}_{k\in\mathcal{G}^f,m}$ is useless to predict rival firms' offerings decisions. It follows that $\mathbb{E}[\Delta_{jm}(\{j\},\Omega^{-f}_m)|\mathcal{I},\mathcal{G}^f,\{\nu^e_{km}\}_{k\in\mathcal{G}^f,m}] = \mathbb{E}[\Delta_{jm}(\{j\},\Omega^{-f}_m)|\mathcal{I},\mathcal{G}^f]$, and similarly for the lower bound inequality (10). The key insight now is that in inequalities (9) and (10), the unobserved fixed cost shock ν^e_{jm} (which enters additively) is independent of the conditional expectation operators inside the indicator functions; that is, independent of $\mathcal{I}$ and $\mathcal{G}^f$. For firm f, knowing ν^e_{jm} is not useful for predicting the bounds on its expected gain in variable profits. Thus, I take expectations conditional on $\mathcal{I}$ on both sides of inequalities (9)-(10) and use the known distribution of $\nu^e_{jm}|\mathcal{I}$ to obtain bounds on the probability that any product j in the firm's portfolio is offered in market m. Letting Γ_{jm} denote the CDF of the log-normal market entry fixed cost,

$$\begin{aligned}\Gamma_{jm}(\mathbb{E}[\Delta_{jm}(\mathcal{G}^f,\Omega^{-f}_m)|\mathcal{I},\mathcal{G}^f];\theta_e,\sigma_e) &\leq \mathbb{P}(O_{jm}=1|\mathcal{I},\mathcal{G}^f)\\ &\leq \Gamma_{jm}(\mathbb{E}[\Delta_{jm}(\{j\},\Omega^{-f}_m)|\mathcal{I},\mathcal{G}^f];\theta_e,\sigma_e). \qquad (11)\end{aligned}$$

The inequalities in (11) provide bounds on the ex-ante probability that product $j \in \mathcal{G}^f$ is offered in market m. The lower bound is the probability that the market entry fixed cost is smaller than the change in variable profits from offering product j in market m obtained under maximal cannibalization; the upper bound is the probability that the market entry fixed cost is smaller than the upper bound on the variable profit gains.

In Online Appendix B.2, I derive bounds on the probability, $\mathbb{P}(G_j = 1|\mathcal{I})$, that product j is introduced in the firm's global product portfolio using similar methodological insights as in the derivations of inequalities (9) and (10). Instead of bounding the marginal value of offering a product in a market, the derivations bound the change in value from introducing a product into the firm's portfolio. An additional complication arises, which is that I have to deal with subgame perfection: firms make portfolio decisions in Stage 1, anticipating that they will make optimal market entry decisions in Stage 2. In Online Appendix B.2 I show

[25] Note that the derivation of equations (9)-(10) only require bundles that maximize and minimize the marginal value of each product j in each market. Submodularity is stronger than required for the derivation of the moment inequalities, but plays a more important role in the solution algorithm derived in Section 6.

how to deal with this issue and that,

$$\tilde{\Lambda}_j\Big(\big\{\mathbb{E}[\Delta_{jm}(\mathcal{A}^f,\Omega_m^{-f})|\mathcal{I}]\big\}_m\Big) \leq \mathbb{P}(G_j=1|\mathcal{I}) \leq \tilde{\Lambda}_j\Big(\big\{\mathbb{E}[\Delta_{jm}(\{j\},\Omega_m^{-f})|\mathcal{I}]\big\}_m\Big), \tag{12}$$

where $\tilde{\Lambda}_j$ is an increasing function in each of its M (number of markets) arguments and depends both on the distribution of portfolio fixed costs Λ_j and the distributions of market entry fixed costs across markets $\{\Gamma_{jm}\}_m$. Intuitively, bounds on the probability of portfolio introduction depend on (i) bounds on the additional variable profits that the product can earn in each market, (ii) the distribution of portfolio fixed costs determining how costly it is to include the product in the firm's portfolio, and (iii) the distributions of market entry fixed costs which determine how costly it is to offer the product in each market conditional on being included in the firm's portfolio.

In summary, I showed how to derive bounds on firms' choice probabilities in a strategic setting with multiple firms making multiple discrete choices under Assumptions 1-4.

4.2 Moment Inequalities using Convex/Concave Bounds of the Fixed Cost CDF

In this subsection, I show how to use the bounds on firms' market entry probabilities to derive moment inequalities that are useful for the estimation of (θ_e, σ_e). In Online Appendix B.2, I similarly show how to derive moment inequalities that partially identify (θ_g, σ_g) based on the bounds on the probability of portfolio introduction.

The inequalities in (11) provide bounds on the ex-ante probability of product offerings, but these depend on firms' expectations over rivals' product offerings Ω_m^{-f} in the equilibrium that generates the data. I develop moment inequalities that do not require solving the model and only depend on observed data and parameters.

To obtain an upper bound moment inequality for estimation of fixed cost parameters (θ_e, σ_e), I bound the term on the right-hand side of the inequalities in (11) with a convex upper bound of Γ_{jm}, which I denote by $\overline{\Gamma}_{jm}$.[26] This step is related to how Dickstein and Morales (2018), Dickstein et al. (2024), and Porcher et al. (2024) use a convex odds or linear approximation function to derive moment inequalities in a single-agent setting. Bounding choice probabilities directly yields,

$$\overline{\Gamma}_{jm}(\mathbb{E}[\Delta_{jm}(\{j\},\Omega_m^{-f})|\mathcal{I},\mathcal{G}^f];\theta_e,\sigma_e) \geq \mathbb{E}[O_{jm}|\mathcal{I},\mathcal{G}^f]. \tag{13}$$

[26] There are many potential convex upper bounds of any CDF function. In Online Appendix B, I discuss the implementation in detail. Figure B.1 plots convex upper bounds of a log-normal CDF.

Since $\overline{\Gamma}_{jm}$ is convex, I apply Jensen's inequality,

$$\mathbb{E}[\overline{\Gamma}_{jm}(\Delta_{jm}(\{j\}, \Omega_m^{-f}); \theta_e, \sigma_e) - O_{jm} | \mathcal{I}, \mathcal{G}^f] \geq 0. \tag{14}$$

Inequality (14) is a conditional moment inequality that depends on the realization of rivals' offerings, Ω_m^{-f}. This relates to previous work that derives moment inequalities that depend on the ex-post realization of payoff-relevant variables (i.e., Pakes 2010, Pakes et al. 2015, Dickstein and Morales 2018); I extend the tricks developed in such papers to an incomplete-information game where expectations are over rivals' endogenous actions. Due to Assumption 1, entry decisions are conditionally independent across firms, which permits using the realized set of entry decisions for estimation, while simultaneously allowing for selection on unobserved (own) fixed cost shocks. Under complete information, inequality (14) would be invalid due to violation of this independence result. The unobservability of rival fixed cost shocks reduces the issue of selection on unobservables, yielding model-consistent bounds that are tighter than the analogous moment inequalities obtained under the assumption of complete information.[27] Using similar arguments and a concave lower bound for the log-normal CDF, which I denote by $\underline{\Gamma}_{jm}$, I derive,

$$\mathbb{E}[\underline{\Gamma}_{jm}(\Delta_{jm}(\mathcal{G}^f, \Omega_m^{-f}); \theta_e, \sigma_e) - O_{jm} | \mathcal{I}, \mathcal{G}^f] \leq 0. \tag{15}$$

Conditional moment inequalities (14) and (15) partially identify (θ_e, σ_e). I also prove the corresponding conditional moment inequalities that partially identify (θ_g, σ_g), which are based on bounds on the probability of product portfolio introduction, given by the inequalities in (12). The derivations in Online Appendix B.2 show how to use the bounds on the probability of portfolio introduction to obtain moment inequalities (40) and (48) using similar insights as used for inequalities (14)-(15). Theorem 1 summarizes the result.

Theorem 1 *The set of parameter vectors consistent with conditional moment inequalities (14), (15), (40), and (48) contains the true parameter vector* $(\theta_e, \sigma_e, \theta_g, \sigma_g)$.

Proof. See Online Appendix B. ■

[27] Under complete information, one could obtain the analogous model-consistent inequality,

$$\mathbb{E}[\overline{\Gamma}_{jm}(\Delta_{jm}(\{j\}, \emptyset); \theta_e, \sigma_e) - O_{jm} | \mathcal{I}] \geq 0.$$

Under Assumption 4, the largest gain from offering product j in market m obtains when product j is the only product (other than the outside option) offered in market m. This extreme bundle is required under complete information because it is the only way to obtain a known bound that is independent of the (unobserved) information set of any given firm. The bounds are less informative when applied to the same data because $\Delta_{jm}(\{j\}, \emptyset) \geq \Delta_{jm}(\{j\}, \Omega_m^{-f})$ for all $j \in \mathcal{A}$ and all $m \in \mathcal{M}$.

My moment inequality approach is based on moments that bound firms' choice probabilities, which makes them suitable for computing counterfactual exercises. In my setting, Assumption 1 provides a model of firms' expectations over rivals' actions, which I exploit in Section 6 to provide a solution method. In other settings, one could use the moment inequalities directly to compute counterfactuals. For instance, the convex bounds on choice probabilities can be applied to multiple (or single) discrete-choice single-agent settings where expectations are over an exogenous, ex-post realized, payoff-relevant variable.

Though I assume that the distributions of both product portfolio and market entry fixed costs are log-normal, the arguments discussed in Section 4 apply under different assumptions on the fixed cost distributions.[28] I assume a log-normal specification because it is a natural benchmark in a setting with significant product-market heterogeneity.

In Online Appendix G.2, I simulate a fully solvable version of the model to show that the inequalities are more informative in settings with low substitutability across products within firms and a larger number of firms. Substiutability across firms does not affect the informativeness of the moment inequalities in simulations. Intuitively, using "extreme" bundles to deal with the selection problem is more costly when the returns from offering a product in a market depend greatly on which other products the firm is selling there. The moment inequalities are highly informative in settings with substantial product differentiation or single-product firms.

4.2.1 Identification

My approach yields moment inequalities based on the derived bounds on entry probabilities. Excessive product portfolio costs can be rejected if they imply that the upper bound on the probability of portfolio introduction (see inequality (40) in the Online Appendix) is smaller than the empirical probability of portfolio introduction. Insufficient portfolio costs can be rejected if they imply that the lower bound on the probability of portfolio introduction (see inequality (48) in the Online Appendix) is larger than the empirical probability. Similarly, market entry fixed cost parameters can be rejected whenever the empirical probability of market entry (conditional on the observed portfolios) does not lie within the bounds given by inequalities (14) and (15). The product portfolio and market entry fixed costs rationalize the cross-sectional entry patterns discussed in Section 2: firms do not develop all of their potential products and do not sell all developed products in all markets.

[28] More precisely, the log-normal assumption only matters for the construction of the convex/concave bound of the fixed cost CDFs. Under a different distributional assumption and, therefore, a different CDF, the convex upper bound and concave lower bound functions would differ, but all remaining arguments would still be valid.

The inequalities identify not only the location parameters of the log-normal distributions but also the scale parameters, determining the variance of (log) fixed costs. Too large a scale parameter (holding fixed the location parameter) can be rejected because it would imply too low a probability of product introduction for highly profitable products. This is because a large scale parameter makes the tails of the fixed cost distributions fatter, therefore making the implied bounds on the probability of product introduction tend towards 0.5 irrespective of the magnitude of Δ_{jm}. Similarly, scale parameters that are too small can be rejected because they would imply that the rate of product introduction tends to 1 for highly profitable deviations, which is not what is observed in the data. Intuitively, the variance of fixed costs must be large enough to rationalize market entry and portfolio deviations that are apparently highly profitable and yet are not undertaken by firms.

Identifying the second moment of the fixed cost distribution is a key advantage relative to previous approaches in the multi-product entry literature, such as Eizenberg (2014), Wollmann (2018), or Bontemps et al. (2023), which only identify the mean. By estimating all the parameters describing the fixed cost distributions, one can accurately integrate firms' fixed cost shocks, which is necessary to solve for model-consistent counterfactual outcomes.

To summarize, I showed how to derive moment inequalities in a setting with multiple discrete choices in a computationally feasible manner robust to equilibrium non-uniqueness. The inequalities solely depend on changes in variable profits, which can be computed using the Stage 3 model, observable entry decisions O_{jm} and Ω_m^{-f}, and fixed cost parameters.

4.3 Moments and Inference

To use the conditional moment inequalities from Theorem 1 for estimation, I construct instrument functions that are positive-valued and exogenous to obtain unconditional moment inequalities.[29] My approach is similar to two-stage least squares: I project endogenous variables onto exogenous variables and then use the first-stage predicted values to compute instruments. Online Appendix D.2 provides additional details about the implementation.

First, I simulate S draws of demand and marginal cost shocks $(\boldsymbol{\xi}, \boldsymbol{\omega})$ using the fitted bivariate normal distribution. Second, I use the model for Stage 3 of the game to compute $\widehat{\Delta}_{jm}(\{j\}, \Omega_m^{-f})$, $\widehat{\Delta}_{jm}(\mathcal{G}^f, \Omega_m^{-f})$, and $\widehat{\Delta}_{jm}(\mathcal{A}^f, \Omega_m^{-f})$, where the $\widehat{\Delta}$ notation denotes that such marginal values are averaged across the S simulation draws for $(\boldsymbol{\xi}, \boldsymbol{\omega})$. Third, I project each of the endogenous $\widehat{\Delta}_{jm}(\{j\}, \Omega_m^{-f})$, $\widehat{\Delta}_{jm}(\mathcal{G}^f, \Omega_m^{-f})$, and $\widehat{\Delta}_{jm}(\mathcal{A}^f, \Omega_m^{-f})$ on exogenous ($\mathcal{I}-$measurable) market size M_m, and interactions of market identifiers with $\tilde{\delta}_{jm}$, $\tilde{c}_{jm}$ and $\tilde{\delta}_{jm} \times \tilde{c}_{jm}$ using PPML, where $\tilde{\delta}_{jm}$ and $\tilde{c}_{jm}$ are defined in equations (2) and (3). I then

[29] More precisely, one requires instruments that are positive and $\mathcal{I}$-measurable.

use the predicted values from the PPML model to construct instruments. Introducing notation, let $\hat{\bar{x}}_{jm}$, $\hat{\underline{x}}^h_{jm}$, and $\hat{\underline{x}}^l_{jm}$ be the predicted values of $\widehat{\Delta}_{jm}(\{j\}, \Omega_m^{-f})$, $\widehat{\Delta}_{jm}(\mathcal{G}^f, \Omega_m^{-f})$, and $\widehat{\Delta}_{jm}(\mathcal{A}^f, \Omega_m^{-f})$, respectively. For inequality (14), I build instruments based on percentile category bins of $\hat{\bar{x}}_{jm}$ i.e., $\mathbb{1}\{\hat{\bar{x}}_{jm} \in Q_\tau(\hat{\bar{x}}_{jm})\}$, where Q_τ denotes a percentile category bin τ.[30] Similarly, I construct instruments for inequality (15) according to $\mathbb{1}\{\hat{\underline{x}}^h_{jm} \in Q_\tau(\hat{\underline{x}}^h_{jm})\}$. For Stage 1, upper and lower bound inequalities, I construct percentile category bin indicators, respectively, of the form,

$$\mathbb{1}\Big\{ \sum_{m\in\mathcal{M}} \hat{\bar{x}}_{jm} \in Q_\tau\Big(\sum_{m\in\mathcal{M}} \hat{\bar{x}}_{jm}\Big)\Big\} \text{ and } \mathbb{1}\Big\{ \sum_{m\in\mathcal{M}} \hat{\underline{x}}^l_{jm} \in Q_\tau\Big(\sum_{m\in\mathcal{M}} \hat{\underline{x}}^l_{jm}\Big)\Big\}.$$

The instruments are positive and exogenous ($\mathcal{I}$-measurable), depending only on factors known to firms at entry. [31] Interacting moments with indicators for exogenous profitability levels identifies fixed cost parameters: unprofitable products give upper bounds (costs must be low enough for some to enter), while profitable products give lower bounds (costs must be high enough for some not to enter). I implement the moment inequalities by computing the empirical counterparts of the moments in Theorem 1, reported in Online Appendix D.3.

To conduct inference, I proceed sequentially by constructing confidence sets performing a grid search and using Andrews and Soares (2010). First, I construct a confidence set for (θ_e, σ_e) using the empirical analogs of inequalities (14)-(15), interacted with the instruments. Then, I construct a confidence set for (θ_g, σ_g) using the empirical analogs of the portfolio introduction moment inequalities (see (40) and (48) in the Online Appendix) interacted with the instruments, and the confidence set obtained for (θ_e, σ_e) in the first step.[32] For the (θ_e, σ_e) confidence set, I rely on asymptotics as $JM \to \infty$. The (θ_g, σ_g) confidence set relies on asymptotics as $J \to \infty$.[33]

In Online Appendix G.3, I simulate a solvable version of the model and show that the Andrews and Soares (2010) confidence sets perform well in terms of coverage even if a single cross-section is used to estimate the fixed costs. I replicate my estimation procedure with simulated data and show that undercoverage is very limited even when using a robust

[30] For instance, if we have data on X_i, and we construct $Q_1(X_i)$, $Q_2(X_i)$, and $Q_3(X_i)$, then $x_i \in Q_1(X_i)$ if and only if x_i lies below the 33th percentile of the empirical distribution of X_i.

[31] See Online Appendix D.2 for the construction of the instruments. In the main specifications, I use 3 bins for the Stage 2 inequalities, 5 bins for the Stage 1 inequalities, and polynomials of the PPML-predicted values.

[32] I apply a Bonferroni correction to account for the multiple-testing issue.

[33] An issue with inference in this setting is that, due to the strategic nature of the model, expectational errors are correlated, leading to the violation of statistical independence across observations. Menzel (2016) develops an asymptotic theory for discrete action games with many players. As the number of players goes to infinity, firms can more precisely forecast all payoff-relevant aspects of the market structure due to the law of large numbers.

variance-covariance matrix, particularly when the number of firms is large.

5 Estimation Results

In this section, I report the estimates of model parameters. The timing of the model implies that estimation can be done in 3 steps. First, I estimate demand and marginal cost parameters. Then, I implement the two-step procedure described in Section 4 to sequentially estimate $(\theta_e, \sigma_e, \theta_g, \sigma_g)$.

5.1 Demand and Marginal Costs

Demand estimation follows Petrin (2002). I use Gandhi and Houde (2019) differentiation instruments and micro-moments that match the probability that buyers of different incomes purchase vehicles in different price ranges. The key identifying assumption leverages the assumption that demand and marginal cost shocks are realized after firms make their portfolio and market entry choices. This guarantees that $\mathbb{E}[(\xi_{jm}, \omega_{jm})|\mathcal{I}] = 0$, which makes standard Berry et al. (1995) or Gandhi and Houde (2019) instruments valid. Additional details about the implementation and the moments used for estimation are included in Online Appendix D.1.

Table 1 reports the demand and marginal cost estimates. The signs and magnitudes of all estimates are consistent with previous estimates in the literature. Moreover, all of the key parameter estimates are significant at the 95% significance level. The mean share-weighted implied own-price elasticity is -8.41 across all countries, higher (in absolute value) than that obtained for the United States in Grieco et al. (2023) (they find it to be -6.37 in 2018). The coefficient of -0.790 for the interaction of each buyer's income with its price sensitivity is highly significant and in line with previous estimates of around -1 in the literature (e.g., Coşar et al. 2017 obtain a point estimate of -0.997). The coefficient on the home market dummy shows that consumers are willing to pay a substantial premium to purchase from a local brand. The coefficient of 1.01 implies that, on average, across the world income distribution, consumers' additional utility from a home brand is valued, all else equal, equivalently to a $1,518 U.S. dollars decrease in price. Consumers significantly value vehicles with greater horsepower and vehicles with less horsepower/weight, which is negatively correlated with fuel efficiency (an omitted variable in my dataset). Finally, interacting the price coefficient with a China dummy proves to be important. The coefficient of -1.51 is highly significant and shows that all else equal, Chinese consumers appear to be less price-sensitive. I hypothesize that this is due to aggressive and unobserved demand-side policies that render the MSRP

Table 1: Stage 3 Parameter Estimates

	Parameter	Std. error
Demand		
Mean parameters		
price (α_0)	2.88	0.690
home market	1.01	0.153
horsepower (log)	5.00	2.46
horsepower/weight (log)	-2.19	1.44
Non-linear parameters		
Income (α_1)	-0.790	0.117
China	-1.51	0.297
Shock Std (σ^y)	0.809	0.131
Marginal Costs (log)		
electric	0.340	0.051
hybrid	0.272	0.030
horsepower/weight (log)	-0.426	0.111
horsepower (log)	1.00	0.112
size (log)	0.251	0.209
distance to brand HQ (log)	0.062	0.007
Observations	1,414	
Mean Share-Weighted Implied Own Price Elasticity	-8.41	
Percent Implied Negative Marginal Costs	0	

Note: The demand specification includes body-type and electric-hybrid dummies interacted with market dummies. It also includes size interacted with market dummies. Both specifications include brand and market fixed effects. Standard errors are clustered at the brand level.

an overestimate of consumer prices in China. I interpret this dummy as controlling for such policies, which I will hold fixed in counterfactual exercises.

Marginal cost estimates align with prior institutional knowledge of the industry. In my sample period, supplying an electric vehicle is substantially more costly than a hybrid or internal combustion engine (ICE) vehicle. The marginal cost of ICE vehicles is around 31% (40%) lower than that of a hybrid (electric) vehicle. It is also more costly to supply a vehicle with higher horsepower and size, though the latter coefficient is not statistically significant. It is less costly to supply vehicles with higher horsepower/weight. Finally, it is also substantially more expensive to supply markets that are distant from the brand's headquarters country. A 1% increase in distance from the headquarters country raises the marginal cost of supplying a given market by 0.06%.

Marginal costs are recovered from prices given the Nash-Bertrand conduct assumption, so they contain information on optimal sourcing decisions from 2019 production locations.

Table 2: Stages 1 and 2 Parameter Confidence Set Limits

	95% Confidence Set Limits
Stage 2: Market Entry Fixed Cost	
θ_e (Location)	[-4.8, -4.2]
σ_e (Scale)	[2.9, 4.4]
Stage 1: Product Fixed Cost	
θ_g (Location)	[-1.8, -0.6]
σ_g (Scale)	[1.6, 3.3]
Observations - Stage 2	3240
Observations - Stage 1	739

Note: Confidence sets are computed using Andrews and Soares (2010). First, I implement a grid search to compute a 97.5% confidence set for parameters (θ_e, σ_e) using the Stage 2 moment inequalities. Then, I use the Stage 1 moment inequalities to compute a 97.5% confidence set for (θ_g, σ_g), evaluating the moments at the accepted values of (θ_e, σ_e). The Bonferroni correction yields a 95% confidence set for all 4 parameters. Marginal values are in billions of U.S. dollars. Online Appendix Figure E.1 plots the full grid of accepted values.

I treat these as fixed in counterfactual experiments and interpret my marginal cost parameters and fixed effects as capturing firms' sourcing possibilities. Given demand estimates and the Nash-Bertrand pricing assumption, I compute $\Delta_{jm}(\mathcal{G}^f, \Omega_m^{-f})$, $\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})$, and $\Delta_{jm}(\{j\}, \Omega_m^{-f})$, which are the "data" required for my moment inequality procedure.[34]

5.2 Fixed Costs

Using the moment inequality procedure, I separately identify the distributions of global portfolio and market entry fixed costs. As shown in Section 2, both are apparently important for explaining cross-sectional product offerings. Estimating their magnitudes is the first step toward quantifying policy effects. High portfolio and low entry costs imply similar varieties offered across countries, making globalization or market-access restrictions highly consequential. Conversely, low portfolio and high entry costs imply independent markets with limited global integration effects. Table 2 reports the estimates.[35]

The estimates show that the distribution of product fixed costs has a higher median than the distribution of market entry fixed costs. The estimates imply that the median fixed cost of adding a product to a firm's portfolio is \$138-549 million.[36] Meanwhile, the estimates

[34] As in previous work in this literature (e.g., Wollmann 2018, Fan and Yang 2025), I will treat these marginal values as data when estimating the fixed costs.

[35] These results are highly robust to the choice of instruments. In Online Appendix E, I show that even with more instrument bins, it is not possible to reject the model by obtaining an empty confidence set.

[36] A downside of this product definition is that it does not differentiate across products within this categories (e.g., Toyota Corollas v.s. Honda Civics). An advantage of aggregating at a coarser rather than at a finer level is that it reduces the concern that product development costs are shared across product categories.

imply that the median fixed cost of offering a product in a market is \$8-15 million per product. This suggests that scale economies at the product level are important, implying interdependence in global market outcomes in the automobile industry.

Moreover, the estimates reveal sizeable scale parameters for both the global portfolio and market entry fixed cost (log-normal) distributions, which determine the variance of (log) fixed costs from the perspective of the researcher. Such scale parameters rationalize why some firms make some apparently unprofitable choices and why they do not make some seemingly profitable choices. Thus, these parameters reflect complexities that firms observe, and researchers do not.

The magnitudes of the estimates accord well with other evidence in the literature. Wollmann (2018) estimates that, in the truck industry, it costs \$5-25 million on average to introduce a product into the United States. My median market entry costs of \$8-15 million per product lie within this range. I follow the approach in Wollmann (2018) to convert my cross-sectional estimates into "dynamic" estimates. Under a discount factor of 0.9 (a hurdle rate of 0.1), my estimates imply a median of \$1.4-5.5 billion per product.[37] IHS Global reports that the cost of developing and maintaining product lines in the automotive industry is \$1-6 billion per product, which aligns with my estimates (Autoblog 2010).

A limitation of identifying fixed costs from cross-sectional product offerings is that it makes it impossible to fully distinguish fixed from sunk costs. While incorporating dynamics is definitely an important direction for future research, my framework captures many important features of the automobile industry and is capable of separately identifying product portfolio and market entry fixed costs.

6 Solution Method for Product Entry Games

I propose a solution method to bound any Bayesian Nash equilibrium of the global multi-product entry game based on the bounds on entry probabilities derived in Section 4.1.

Given any information set $\mathcal{I}$, and implied equilibrium distribution of product offerings in market m, which I denote by $\boldsymbol{\mu}_m^*$, the algorithm yields bounds $\underline{\boldsymbol{\mu}}_m$ and $\overline{\boldsymbol{\mu}}_m$ such that $\underline{\boldsymbol{\mu}}_m \leq_{FOSD} \boldsymbol{\mu}_m^* \leq_{FOSD} \overline{\boldsymbol{\mu}}_m$.[38] This solution approach does not rely on approximation methods nor equilibrium selection rules (e.g., order of entry assumptions) and is computationally

[37] This approach converts fixed costs estimated from a cross-section into sunk costs in a model that assumes that firms choose product offerings myopically without anticipating future changes in market structure.

[38] Note that the bounds are not necessarily equilibria, as is the case with supermodular games. The solution approach can also be used to build moment inequalities for estimation. In my setting, even though it is feasible to compute counterfactuals given parameters, computing solution bounds under many parameter vectors (as required for estimation) is still computationally expensive.

feasible to implement even in settings with many firms and discrete choices.[39]

Throughout this section, I treat the set of fixed cost parameters $(\theta_e, \sigma_e, \theta_g, \sigma_g)$ determining the distributions of global portfolio and market entry fixed costs as known for ease of exposition. Recall that the CDF of the fixed cost of offering product j in market m is denoted by Γ_{jm}, while the CDF of the global portfolio fixed cost is denoted by Λ_j.

The solution method relies on the baseline choice-probability bounds discussed in Section 4.1, given by inequalities (11) and (12). Importantly, in these inequalities, each firm's expectation is with respect to rival firms' equilibrium distribution of entry decisions in each market m, $\boldsymbol{\mu}_m^*$. In Online Appendix C, I show that provided fixed cost draws are independent conditional on $\mathcal{I}$, the multiple of the lower and the upper bounds in inequalities (11) and (12) bound the ex-ante probability that product j is offered in market m, $\mathbb{P}(O_{jm} = 1|\mathcal{I})$.

Clearly, because (11) and (12) depend on $\boldsymbol{\mu}_m^*$, which is the object of interest in this section, these inequalities are not immediately useful to solve the model. Instead, I exploit Assumption 4 to devise Algorithm 1, described below.

Algorithm 1 *Let $\mathbb{E}_{\boldsymbol{\mu}}[\Delta_{jm}(\Omega_m^f, \Omega_m^{-f})|\mathcal{I}]$ denote expectations conditional on $\mathcal{I}$ according to probability measure $\boldsymbol{\mu}$ over rival firms' offerings. The algorithm operates in the following steps, given fixed cost parameters $(\theta_e, \sigma_e, \theta_g, \sigma_g)$:*

1. *(Initialization) Compute, for each potential product $j \in \mathcal{A}$ and each market $m \in \mathcal{M}$, $\overline{\mu}_{jm}^{1,cond} = \Gamma_{jm}\big(\mathbb{E}\big[\Delta_{jm}(\{j\}, \emptyset)|\mathcal{I}\big]\big)$, $\underline{\mu}_{jm}^{1,cond} = \Gamma_{jm}\big(\mathbb{E}\big[\Delta_{jm}(\mathcal{A}^f, \mathcal{A}^{-f})|\mathcal{I}\big]\big)$, $\overline{\mu}_{jm}^{1,port} = \tilde{\Lambda}_j\big(\big\{\mathbb{E}[\Delta_{jm}(\{j\}, \emptyset)|\mathcal{I}]\big\}_m\big)$, and $\underline{\mu}_{jm}^{1,port} = \tilde{\Lambda}_j\big(\big\{\mathbb{E}[\Delta_{jm}(\mathcal{A}^f, \mathcal{A}^{-f})|\mathcal{I}]\big\}_m\big)$. At initialization, expectations are solely with respect to demand and marginal cost shocks that are realized at Stage 3 of the game.*

2. *(Ex-ante Rival Entry Probability Bounds) For each $j \in \mathcal{A}$ and each $m \in \mathcal{M}$, compute $\overline{\mu}_{jm}^1 = \overline{\mu}_{jm}^{1,port}\overline{\mu}_{jm}^{1,cond}$ and $\underline{\mu}_{jm}^1 = \underline{\mu}_{jm}^{1,port}\underline{\mu}_{jm}^{1,cond}$. Denote by $\underline{\boldsymbol{\mu}}_m^1$ and $\overline{\boldsymbol{\mu}}_m^1$ the vectors of probability bounds in market m.*

3. *(Iteration) Compute new probability bounds $\overline{\mu}_{jm}^{2,cond} = \Gamma_{jm}\big(\mathbb{E}_{\underline{\boldsymbol{\mu}}_m^1}\big[\Delta_{jm}(\{j\}, \Omega_m^{-f})|\mathcal{I}\big]\big)$, $\underline{\mu}_{jm}^{2,cond} = \Gamma_{jm}\big(\mathbb{E}_{\overline{\boldsymbol{\mu}}_m^1}\big[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})|\mathcal{I}\big]\big)$, $\overline{\mu}_{jm}^{2,port} = \tilde{\Lambda}_j\big(\big\{\mathbb{E}_{\underline{\boldsymbol{\mu}}_m^1}[\Delta_{jm}(\{j\}, \Omega_m^{-f})|\mathcal{I}]\big\}_m\big)$, and $\underline{\mu}_{jm}^{2,port} = \tilde{\Lambda}_j\big(\big\{\mathbb{E}_{\overline{\boldsymbol{\mu}}_m^1}[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})|\mathcal{I}]\big\}_m\big)$ for each $j \in \mathcal{A}$ and each $m \in \mathcal{M}$. Expectations are over rival firms' entry decisions according to the probability bounds, as well as demand and marginal cost shocks.*

[39] Common heuristic approaches in the literature approximate Nash equilibria by iterating over deviations from a given action vector in payoff-improving directions, such as the "greedy" algorithm or simulated annealing.

4. *(Repetition) At iteration k of the algorithm, compute unconditional probability bounds* $\overline{\mu}_{jm}^{k} = \overline{\mu}_{jm}^{k,port}\overline{\mu}_{jm}^{k,cond}$ *and* $\underline{\mu}_{jm}^{k} = \underline{\mu}_{jm}^{k,port}\underline{\mu}_{jm}^{k,cond}$*, and use* $\underline{\boldsymbol{\mu}}_{m}^{k}$ *and* $\overline{\boldsymbol{\mu}}_{m}^{k}$ *to compute updated probability bounds* $\overline{\mu}_{jm}^{k+1,cond} = \Gamma_{jm}\big(\mathbb{E}_{\underline{\boldsymbol{\mu}}_{m}^{k}}\big[\Delta_{jm}(\{j\},\Omega_{m}^{-f})|\mathcal{I}\big]\big)$, $\underline{\mu}_{jm}^{k+1,cond} = \Gamma_{jm}\big(\mathbb{E}_{\overline{\boldsymbol{\mu}}_{m}^{k}}\big[\Delta_{jm}(\mathcal{A}^{f},\Omega_{m}^{-f})|\mathcal{I}\big]\big)$, $\overline{\mu}_{jm}^{k+1,port} = \tilde{\Lambda}_{j}\big(\big\{\mathbb{E}_{\underline{\boldsymbol{\mu}}_{m}^{1}}[\Delta_{jm}(\{j\},\Omega_{m}^{-f})|\mathcal{I}]\big\}_{m}\big)$*, and* $\underline{\mu}_{jm}^{k+1,port} = \tilde{\Lambda}_{j}\big(\big\{\mathbb{E}_{\overline{\boldsymbol{\mu}}_{m}^{k}}[\Delta_{jm}(\mathcal{A}^{f},\Omega_{m}^{-f})|\mathcal{I}]\big\}_{m}\big)$.

5. *(Convergence) Iterate until* $||\boldsymbol{\mu}_{m}^{k+1} - \boldsymbol{\mu}_{m}^{k}||$ *is small enough in each market* $m \in \mathcal{M}$.

Theorem 2 states the properties of Algorithm 1.

Theorem 2 *Under Assumptions 1-4, Algorithm 1 converges monotonically to bounds, in the sense of first-order stochastic dominance, of any equilibrium distribution of product offerings in each market m given any information set $\mathcal{I}$. That is, for any iteration $k > 0$ and any $m \in \mathcal{M}$,*

$$\underline{\boldsymbol{\mu}}_{m}^{k-1} \leq_{FOSD} \underline{\boldsymbol{\mu}}_{m}^{k} \leq_{FOSD} \boldsymbol{\mu}_{m}^{*} \leq_{FOSD} \overline{\boldsymbol{\mu}}_{m}^{k} \leq_{FOSD} \overline{\boldsymbol{\mu}}_{m}^{k-1}.$$

Proof. See Online Appendix C. ■

At step 1 (Initialization), submodularity permits the computation of the weakest possible bounds on the marginal value of any product j in each market m. The highest marginal value of product j in each market is obtained by evaluating it at the least competitive conditions possible: only product j (and the outside option) is offered in market m. The lowest possible marginal value of product j in market m is obtained under the hypothesis that all potential products are simultaneously being offered, the most competitive conditions possible. Submodularity also plays a key role guaranteeing the monotonicity property of the algorithm, which implies that $\underline{\boldsymbol{\mu}}_{m}^{k} \leq_{FOSD} \boldsymbol{\mu}_{m}^{*} \leq_{FOSD} \overline{\boldsymbol{\mu}}_{m}^{k}$ at any iteration step $k > 0$. To prove monotonicity, I use the property that a probability distribution X first-order stochastically dominates Y if and only if for every monotonically decreasing function f, $\mathbb{E}[f(X)] \leq \mathbb{E}[f(Y)]$.[40] Intuitively, an increase in lower-bound rival entry probabilities can only decrease the entry-probability upper bounds; a decrease in upper-bound rival entry probabilities can only increase in the entry-probability lower bounds.

The private information assumption is also important.[41] First, it guarantees that at least one PSNE exists. Second, it ensures that product-market entry decisions are independent (conditional on $\mathcal{I}$) across firms. Together with Assumptions 2-3, I show in Online Appendix C.1 that Algorithm 1 generates two J-dimensional random vectors in each market, $\underline{\boldsymbol{\mu}}_{m}^{k}$ and

[40] The proof in Online Appendix C requires that entry decisions across firms be either complements or substitutes. Therefore, the method can be extended to cases where within-firm interdependencies are flexible but known to the researcher, provided that firms exhibit increasing or decreasing differences with respect to rival entry decisions.

[41] For the solution method, one could allow for a complete information component that is unobserved to the researcher but observed to firms. What matters is that there is a component that is private information.

$\overline{\boldsymbol{\mu}}_m^k$, that consist of independent Bernoulli random variables. Setting within-firm bundles to extreme (low or high) cannibalization eliminates within-firm dependence; independent private shocks eliminate across-firm dependence. While $\boldsymbol{\mu}_m^*$ has a complicated joint distribution due to interdependencies across entry decisions stemming from the combinatorial problem, the bounding random vectors are composed of mutually independent random variables. This property facilitates the simulation of the expectations in Steps 3-4 of Algorithm 1: one simply simulates T sets of $J \times M$ uniform draws and uses $\underline{\boldsymbol{\mu}}_m^k$ and $\overline{\boldsymbol{\mu}}_m^k$ as thresholds determining product-market entry.[42]

Algorithm 1 iteratively "eliminates" dominated strategies by exploiting Assumption 4. If a product's offer probability is high under the most competitive conditions, it must be at least as high in any equilibrium; if low under the least competitive conditions, it must be no higher in equilibrium. Since all firms can compute these bounds given $\mathcal{I}$, iterating yields the equilibrium bounds in Theorem 2.

Intuitively, the informativeness of the inequalities will be affected by substitutability of products within firms (more differentiation makes the method more informative) and by the severity of the equilibrium multiplicity problem. In Additional Materials D, I show that in a simple case with two firms each making a binary choice, the solution method converges to $[\underline{p}_1, \overline{p}_1]$ and $[\underline{p}_2, \overline{p}_2]$, where $\underline{p}_i$ ($\overline{p}_i$) is the smallest (largest) entry probability across equilibria.

To summarize, the method provides a mapping from parameters and exogenous product characteristics and the policy environment to bounds on the joint distribution of product offerings implied by the Bayesian Nash conduct assumption. This is related to how, in Berry et al. (1995), exogenous variables and parameters, together with the Nash-Bertrand assumption, yield equilibrium markups.

7 The Impact of National U.S. Policies on Global Market Outcomes

Aggressive industrial and trade policies are globally on the rise. In the automotive industry, policies promoting clean vehicles have been implemented across jurisdictions. Such policies include highly generous production and consumption tax credits or subsidies as part of policy packages such as the Inflation Reduction Act in the United States or the European Green Deal. Ambitious policies have long been prevalent in China.

To assess the impact of subsidies implemented in a large market on the global structure of the industry, I study the effects of marginal cost and consumer subsidies favoring

[42] Online Appendix C.2 provides the practical implementation details.

U.S.-headquartered brands.[43] The counterfactual experiments disentangle the effects of such policies on market outcomes in other countries, both through the intensive and the extensive margin in a world of interdependent product offerings. They also demonstrate the quantitative importance of product entry in determining profit-shifting and consumer surplus effects in other markets.

To assess the quantitative impact of national policies on global market outcomes while allowing for endogenous product portfolio responses, I use the algorithm from Section 6. For instance, because consumer surplus is increasing in the set of products offered, I bound expected consumer surplus given any policy environment $\mathcal{I}$ using, $\mathbb{E}_{\underline{\boldsymbol{\mu}}_m}[CS_m|\mathcal{I}] \leq \mathbb{E}_{\boldsymbol{\mu}^*_m}[CS_m|\mathcal{I}] \leq \mathbb{E}_{\overline{\boldsymbol{\mu}}_m}[CS_m|\mathcal{I}]$. I follow a similar approach to bound other outcomes.[44]

Computing counterfactual experiments in a manner robust to equilibrium multiplicity is crucial, particularly if policymakers are interested in profit-shifting motives, where the identity of entrants is of interest. Under different equilibria, the subsidy can potentially have different effects on specific firms' profits, which is important for policymakers.

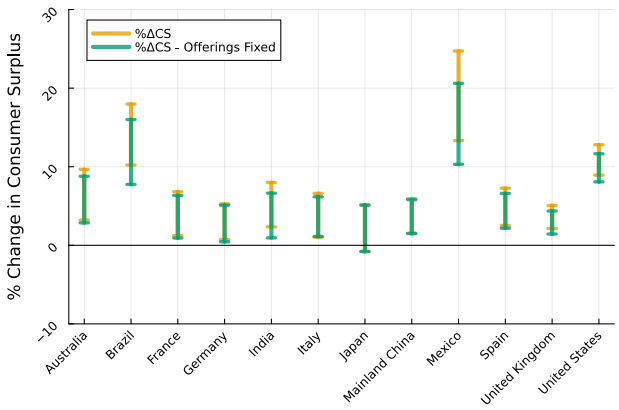
7.1 Effect of a 20% Marginal Cost Subsidy on American Brands

I study the equilibrium effects of a 20% reduction in the marginal cost c_{jm} of each product potentially offered by a brand with headquarters in the United States in any market. The U.S. brands that receive the production subsidy include Buick, Cadillac, Chevrolet, Chrysler, Dodge, Ford, GMC, Jeep, Lincoln, and Tesla. The Inflation Reduction Act targets clean and energy-efficient vehicles with 10% production subsidies for stages like critical mineral or electrode active material production, in addition to a subsidy of $35 per kilowatt-hour for battery cell manufacturing and $10 per kilowatt-hour for battery module assembly, according to the 45X production tax credit (Banks 2023). A wide range of additional state-level production incentives have also been implemented (Good Jobs First 2024).

[43] I consider a policy that favors only brands headquartered in the United States, which prevents me from making arbitrary assumptions regarding the production location of potential products that are not sold in my sample. Publicly available data shows that brands with U.S. headquarters are the primary recipients of U.S. federal and state subsidies (Good Jobs First 2024). It is also a good approximation of policies that have been implemented in China, which have overwhelmingly benefited brands with local headquarters.

[44] To simulate a lower (upper) bound on brand-level shares and profits, I integrate other brands' entry choices using $\overline{\boldsymbol{\mu}}_m^{-b}$ ($\underline{\boldsymbol{\mu}}_m^{-b}$) and brand b's entry choices using $\underline{\boldsymbol{\mu}}_m^{b}$ ($\overline{\boldsymbol{\mu}}_m^{b}$). I report bounds on counterfactual outcomes under a point in the confidence set: $(\theta_e, \sigma_e, \theta_g, \sigma_g) = (-4.5, 3.6, -1.2, 2.5)$. This point lies in the midpoint of the 95% confidence set limits. In counterfactual experiments, I compute expectations with respect to the solved bounds on the equilibrium distribution of offerings in each market and over demand and marginal cost shocks. In Additional Materials A, I show that the conclusions are robust to different points in the confidence set.

Figure 5: Change in Consumer Surplus



Note: This figure plots, for each country, bounds on the expected change in consumer surplus (relative to the outside option) following a 20% marginal cost reduction for U.S. brands. The expectation is with respect to bounds on the probability distribution of firms' offerings and demand/marginal cost shocks. The intervals in green show the change in expected consumer surplus using the bounds on the distribution of product offerings before the policy is implemented.

7.1.1 Consumer Surplus

Figure 5 reports the policy's effects on consumer surplus across countries, with orange intervals capturing total effects (price, quantity, and product offerings) and green intervals holding the distribution of product offerings fixed. Many countries see large minimum gains—over 8.92% in the United States, Mexico, and Brazil, and over 2.34% in India. The 20% drop in U.S. production costs benefits poorer nations like Brazil more than richer ones like Germany, reflecting stronger price sensitivity. The United States also gains more than other rich countries due to home bias, as American consumers prefer domestic brands.

The effects of the policy on consumer surplus do not seem to reflect changes on the extensive margin. While the upper and lower bounds of the green intervals lie slightly below those of the orange intervals, such differences are small. This suggests either that (i) product entry does not change in response to the policy or (ii) there are offsetting changes in market structure from the point of view of consumers. To shed light on this, I now report the effects on brand-level outcomes.

7.1.2 Products Offered

Figure 6 shows the effect of the policy on the composition of products offered across countries. Panel A shows that the policy leads to an expansion of the set of products offered by U.S. brands throughout the globe. Lower marginal costs generate greater profitability in all markets, which leads to a greater number of product offerings. Panel B shows that because of competition from U.S. brands, non-U.S. brands tend to downsize their product

offerings across markets, with both the upper bound and lower bound number of non-U.S. varieties declining across the globe.

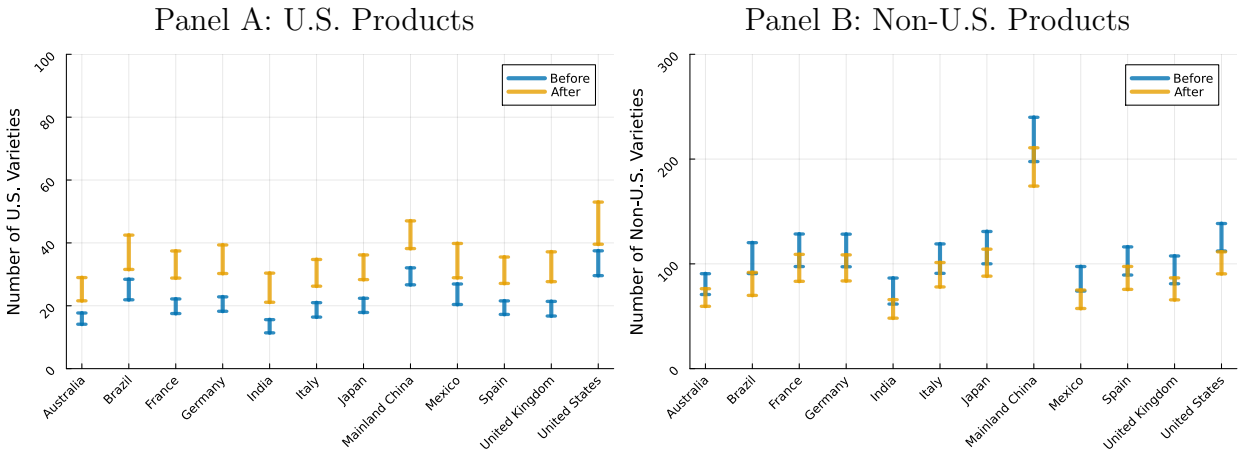
7.1.3 Market Shares and Variable Profits

Figure 7 shows the overall effect of the policy on the market shares of U.S. brands. Comparing the blue intervals to the orange intervals in Panel A, there are significant and heterogeneous effects of the policy worldwide. In the United Kingdom, the increase in the share of U.S. brands is greater than 13.6 percentage points, whereas in Japan, only 2.5-4.1 percentage points. Stateside, U.S.-brand market shares increase at least 9 percentage points.

Importantly, product entry amplifies the increase in the share of U.S. brands, as seen by comparing the green to the orange intervals in Panel A. On average across jurisdictions, product entry accounts for 29.6% and 26.4% of the rise in the upper and the lower bound of the U.S.-brand market shares, respectively. Knowledge of the subsidy causes U.S. brands to expect higher profitability in all markets, which enables them to strengthen their position by offering more products. Meanwhile, non-U.S. brands anticipate reduced profitability due to their relatively greater cost of production. Their expected market share declines by more after accounting for the endogenous exit of non-U.S. brands.

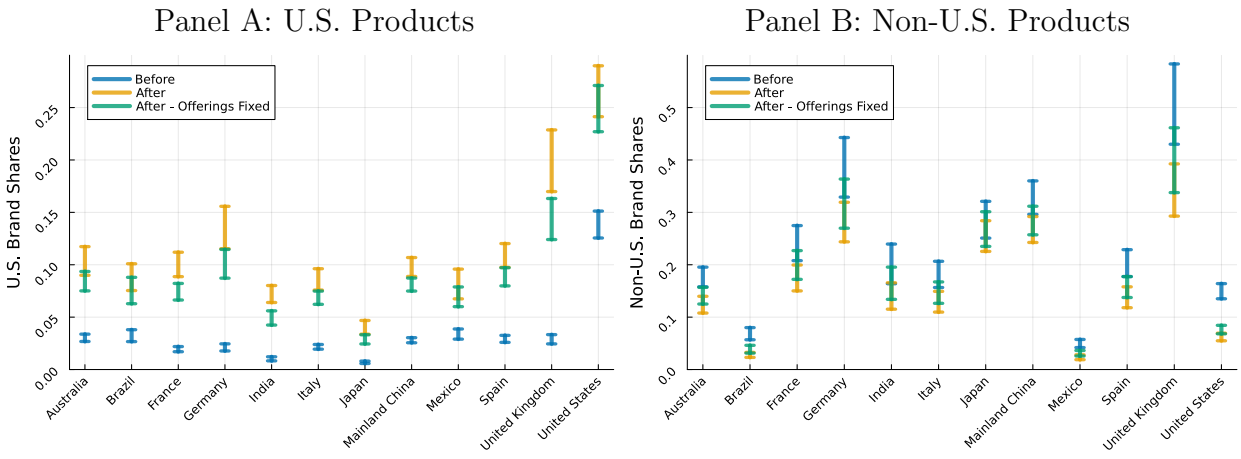
In Figure H.1 of Online Appendix H, I plot bounds on expected (log) total variable profits of U.S. brands before and after the subsidy is implemented. The plot shows very substantial

Figure 6: Number of Products Offered



Note: Panel A displays bounds on the expected number of U.S.-branded products offered across countries before (blue) and after (orange) a 20% reduction in U.S. brands' marginal costs. The expectation is with respect to bounds on the probability distribution of firms' offerings and demand/marginal cost shocks. Panel B displays the corresponding bounds on the expected number of non-U.S.-branded products.

Figure 7: Market Shares



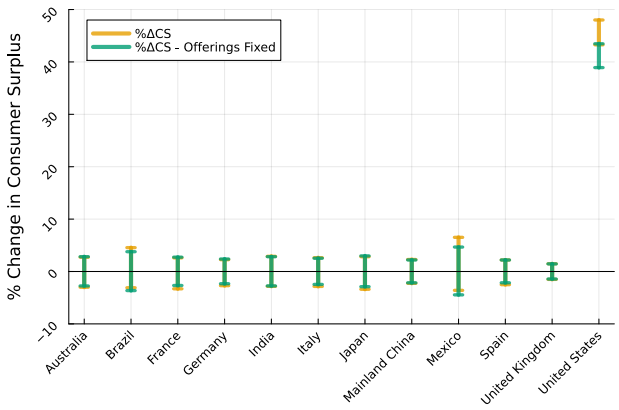
Note: Panel A displays bounds on the expected total market share of U.S. brands across countries before (blue) and after (orange) a 20% reduction in U.S. brands' marginal costs. The expectation is with respect to bounds on the probability distribution of firms' offerings and demand/marginal cost shocks. The intervals in green use the bounds on the distribution of product offerings in each market before the policy is implemented. Panel B displays the corresponding bounds on the expected total market share of non-U.S. brands before and after the policy.

increases in profits of more than 1 log point (172%) in countries like the United Kingdom, India, or Germany and very sizeable increases in variable profits across most jurisdictions, with the smallest minimum increase in the United States. As with market shares, the rise in profits (comparing the green to the orange intervals) would be underestimated if one ignored endogenous product offering changes. A similar and reverse story can be told for non-U.S. brands, who observe a decline in their expected variable profits across most markets, though the decline is not as large as the rise in U.S. brands' profits.

7.2 Effect of a 50% Consumption Subsidy on American Brands

In this section, I study the effect of a 50% U.S. consumption subsidy on products offered by U.S. brands. The motivation for studying the effect of a consumer-side subsidy is the large and generous consumer-side policies implemented throughout the world providing incentives to purchase clean vehicles. As with production-side subsidies, many of these policies have implicitly or explicitly favored local brands over foreign brands. In China, a highly protected automotive market since its inception, subsidies favoring Chinese-made electric vehicles covered around 40-60% of their price in 2014 (Lu 2018). In the United States, clean vehicles are subsidized with tax credits of up to $7,500, with Buy American incentives contingent on local final assembly.

Figure 8: Change in Consumer Surplus



Note: This figure plots, for each country, bounds on the expected change in consumer surplus (relative to the outside option) following a 50% consumer subsidy for U.S.-branded products in the United States. The expectation is with respect to bounds on the probability distribution of firms' offerings and demand/marginal cost shocks. The intervals in green show the change in expected consumer surplus using the bounds on the distribution of product offerings before the policy is implemented.

7.2.1 Consumer Surplus and Products Offered

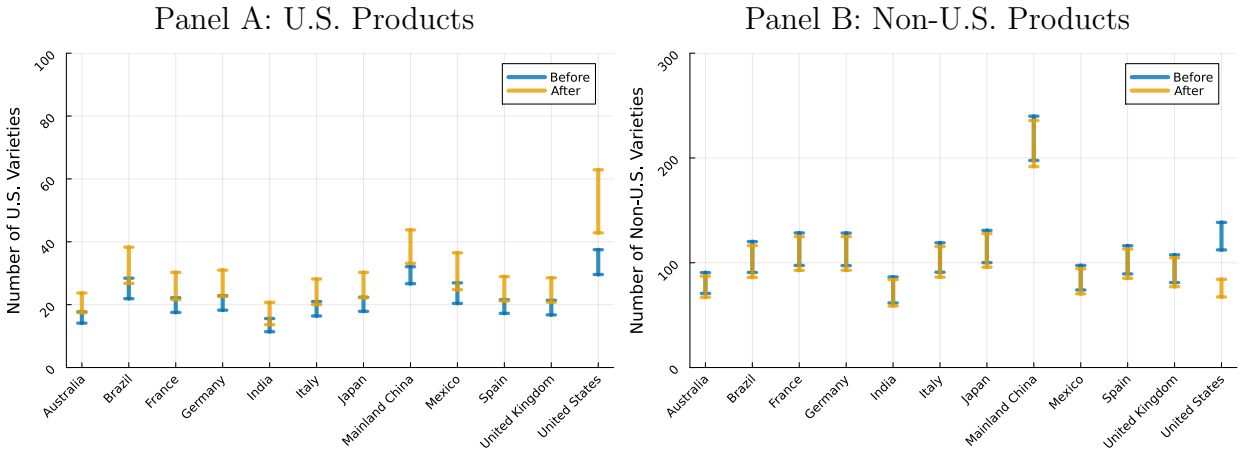
First, I report the effects of the policy on consumer surplus worldwide. Figure 8 shows a large minimum increase in consumer surplus in the United States following the policy of over 43.2%. Moreover, ignoring product entry would lead to underestimating the rise in consumer surplus. This large U.S. subsidy has small effects on non-U.S. consumer surplus.

Figure 9 shows the effect of the policy on the composition of products offered across countries. Panel A shows that the policy leads to an expansion of the sets of products offered by U.S. brands globally. The large consumer subsidy on U.S. brands induces the entry of U.S.-branded products into firms' portfolios, which firms can choose to offer overseas. In the United States, the policy leads to a contraction of varieties offered by foreign brands in light of their reduced profitability. Interestingly, non-U.S. brands do not significantly contract their product offerings in other markets. This suggests that the decline in profits resulting from the discriminatory policy in the United States is not large enough to induce the global exit of many non-U.S. products.

7.2.2 Market Shares and Variable Profits

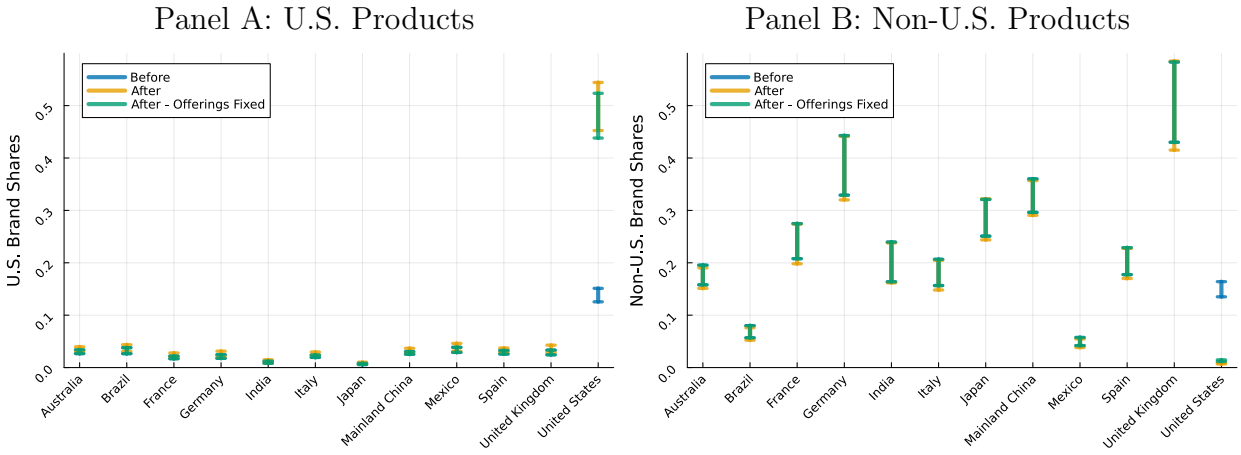
Figure 10 shows that the policy boosts U.S. brands' market share in the United States by over 30 percentage points, while non-U.S. brands' share declines by at least 12.5 percentage points—driven mainly by the intensive margin rather than new product introductions. New U.S. products tend to be unpopular, limiting their market impact. Consequently, U.S.

Figure 9: Number of Products Offered



Note: Panel A displays bounds on the expected number of U.S.-branded products offered across countries before (blue) and after (orange) a 50% consumer subsidy on U.S.-branded products in the United States. The expectation is with respect to bounds on the probability distribution of firms' offerings and demand and marginal cost shocks. Panel B displays the corresponding bounds on the expected number of non-U.S.-branded products before and after the policy.

Figure 10: Market Shares



Note: Panel A displays bounds on the expected total market share of U.S. brands across countries before (blue) and after (orange) a 50% consumer subsidy on U.S.-branded products in the United States. The expectation is with respect to bounds on the probability distribution of firms' offerings and demand and marginal cost shocks. The intervals in green use the bounds on the distribution of product offerings in each market before the policy is implemented. Panel B displays similar bounds on the expected total market share of non-U.S. brands before and after the policy.

brand dominance changes little abroad: despite expanded offerings overseas (Figure 9), weak intensive-margin effects mean consumers show limited interest, unlike in the earlier case where cheaper U.S. products eroded foreign brands' positions.

Figure H.1 in Online Appendix H shows the effects on variable profits. The effects reported in Panel A and Panel B mirror the previously discussed effects on market shares. While the national consumer subsidy favoring U.S. brands has large effects on U.S. brand dominance in the United States, there is little profit shifting in other markets. Unlike the production subsidy, the U.S. consumer subsidy influences foreign markets only through product entry, making its global impact smaller.

8 Conclusion

National policies shape global markets through firms' portfolio decisions. Studying strategic entry in high-dimensional discrete settings is challenging due to complex interdependent decisions, equilibrium multiplicity, and selection on unobservables. This paper develops a method to estimate and solve an incomplete-information entry game with multiple asymmetric firms while jointly overcoming such hurdles. I derive novel inequalities that bound ex-ante entry probabilities before fixed costs are realized and form the basis of both the estimation and solution method.

The methods developed in this paper permit quantifying the effects of various policies (e.g., international mergers, product bans, product-specific subsidies/tariffs) on equilibrium market structure. Many sectors, like automobiles, feature differentiated products, heterogeneous consumers, and multinational (or multi-market) oligopolies. Limitations remain: incorporating dynamics and richer interdependencies, especially via fixed or sunk costs in platform-sharing settings, are key directions for future work.

References

Alfaro-Urena, Alonso, Juanma Castro-Vincenzi, Sebastian Fanelli, and Eduardo Morales (2023): "Firm Export Dynamics in Interdependent Markets," *Mimeo*.

Allcott, Hunt, Reigner Kane, Max Maydanchik, and Joseph Shapiro (2024): "The Effects of "Buy American": Electric Vehicles and the Inflation Reduction Act," *Mimeo*.

Andrews, Donald W. K. and Gustavo Soares (2010): "Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection," *Econometrica*, 78 (1), 119–157.

Antràs, Pol, Teresa C. Fort, and Felix Tintelnot (2017): "The Margins of Global

Sourcing: Theory and Evidence from US Firms," *American Economic Review*, 107 (9), 2514–64.

Arkolakis, Costas, Fabian Eckert, and Rowan Shi (2023): "Combinatorial Discrete Choice: Theory and Application to Multinational Production," *Mimeo*.

Atkeson, Andrew and Ariel Burstein (2008): "Pricing-to-Market, Trade Costs, and International Relative Prices," *American Economic Review*, 98 (5), 1998–2031.

Autoblog (2010): "Why Does It Cost So Much For Automakers To Develop New Models?" https://shorturl.at/bOLX6.

Balder, Erik J. (1988): "Generalized Equilibrium Results for Games with Incomplete Information," *Mathematics of Operations Research*, 13 (2), 265–276.

Banks, Leo (2023): "How Inflation Reduction Act Electric Vehicle Incentives Are Driving a U.S. Manufacturing Renaissance," https://shorturl.at/ZP71Y.

Bernard, Andrew B, Stephen J Redding, and Peter K Schott (2011): "Multiproduct Firms and Trade Liberalization," *The Quarterly Journal of Economics*, 126 (3), 1271–1318.

Berry, Steven, James Levinsohn, and Ariel Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63 (4), 841–890.

Bontemps, Christian, Cristina Gualdani, and Kevin Remmy (2023): "Price Competition and Endogenous Product choice in Networks: Evidence from the US Airline Industry," *Mimeo*.

Castro-Vincenzi, Juan Manuel (2024): "Climate Hazards and Resilience in the Global Car Industry," *Mimeo*.

Castro-Vincenzi, Juan Manuel, Eugenia Menaguale, Eduardo Morales, and Alejandro Sabal (2024): "Market Entry and Plant Location in Multiproduct Firms," *Mimeo*.

Ciliberto, Federico and Elie Tamer (2009): "Market Structure and Multiple Equilibria in Airline Markets," *Econometrica*, 77 (6), 1791–1828.

Conlon, Christopher and Jeff Gortmaker (2020): "Best Practices for Differentiated Products Demand Estimation with PyBLP," *The RAND Journal of Economics*, 51 (4), 1108–1161.

Costinot, Arnaud, Dave Donaldson, Margaret Kyle, and Heidi Williams (2019): "The More We Die, The More We Sell? A Simple Test of the Home-Market Effect*," *The Quarterly Journal of Economics*, 134 (2), 843–894.

Coşar, Kerem, Paul Grieco, Shengyu Li, and Felix Tintelnot (2017): "What Drives Home Market Advantage?" *Journal of International Economics*, 110.

de Paula, Áureo (2013): "Econometric Analysis of Games with Multiple Equilibria," *Annual Review of Economics*, 5 (Volume 5, 2013), 107–131.

Dickstein, Michael, Jihye Jeon, and Eduardo Morales (2024): "Patient Costs and Physicians' Information," *Mimeo*.

Dickstein, Michael J and Eduardo Morales (2018): "What Do Exporters Know?*," *The Quarterly Journal of Economics*, 133 (4), 1753–1801.

Eizenberg, Alon (2014): "Upstream Innovation and Product Variety in the U.S. Home PC Market," *The Review of Economic Studies*, 81 (3 (288)), 1003–1045.

Fan, Ying and Chenyu Yang (2025): "Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety," *Journal of Political Economy*, 133 (6), 1886–1931.

Gandhi, Amit and Jean-François Houde (2019): "Measuring Substitution Patterns in Differentiated-Products Industries," *Mimeo*.

Goldberg, Pinelopi Koujianou (1995): "Product Differentiation and Oligopoly in International Markets: The Case of the U.S. Automobile Industry," *Econometrica*, 63 (4), 891–951.

Good Jobs First (2024): "Subsidy Tracker Database," https://www.goodjobsfirst.org/subsidy-tracker.

Grieco, Paul L E, Charles Murry, and Ali Yurukoglu (2023): "The Evolution of Market Power in the U.S. Automobile Industry," *The Quarterly Journal of Economics*, qjad047.

Head, Keith and Thierry Mayer (2019): "Brands in Motion: How Frictions Shape Multinational Production," *American Economic Review*, 109 (9), 3073–3124.

Head, Keith, Thierry Mayer, Marc Melitz, and Chenying Yang (2024): "Industrial Policies for Multi-Stage Production: The Battle for Battery-Powered Vehicles," *Mimeo*.

Jia, Panle (2008): "What Happens When Wal-Mart Comes to Town: An Empirical Analysis of the Discount Retailing Industry," *Econometrica*, 76 (6), 1263–1316.

Krugman, Paul (1980): "Scale Economies, Product Differentiation, and the Pattern of Trade," *The American Economic Review*, 70 (5), 950–959.

Lu, Jieyi (2018): "Comparing U.S. and Chinese Electric Vehicle Policies," https://shorturl.at/mvLmj.

Magnolfi, Lorenzo and Camilla Roncoroni (2022): "Estimation of Discrete Games with Weak Assumptions on Information," *The Review of Economic Studies*, 90 (4), 2006–2041.

Mayer, Thierry, Marc J Melitz, and Gianmarco I.P. Ottaviano (2021): "Product Mix and Firm Productivity Responses to Trade Competition," *Review of Economics and Statistics*, 103 (5), 874–891.

Menzel, Konrad (2016): "Inference for Games with Many Players," *The Review of Economic Studies*, 83 (1 (294)), 306–337.

Milgrom, Paul R and Robert J Weber (1985): "Distributional Strategies for Games with Incomplete Information," *Mathematics of Operations Research*, 10 (4), 619–632.

Montag, Felix (2024): "Mergers, Foreign Competition, and Jobs: Evidence from the U.S. Appliance Industry," *Mimeo*.

Morales, Eduardo, Gloria Sheu, and Andrés Zahler (2019): "Extended Gravity," *The Review of Economic Studies*, 86 (6), 2668–2712.

Morrow, W. Ross and Steven J. Skerlos (2011): "Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand," *Operations Research*, 59 (2), 328–345.

Nocke, Volker and Nicolas Schutz (2018): "Multiproduct-Firm Oligopoly: An Aggregative Games Approach," *Econometrica*, 86 (2), 523–557.

Pakes, A. (2010): "Alternative Models for Moment Inequalities," *Econometrica*, 78 (6), 1783–1822.

Pakes, A., J. Porter, Kate Ho, and Joy Ishii (2015): "Moment Inequalities and Their Application," *Econometrica*, 83 (1), 315–334.

Petrin, Amil (2002): "Quantifying the Benefits of New Products: The Case of the Minivan," *Journal of Political Economy*, 110 (4), 705–729.

Porcher, Charly, Eduardo Morales, and Thomas Fujiwara (2024): "Measuring Information Frictions in Migration Decisions: A Revealed-Preference Approach," *Mimeo*.

Seim, Katja (2006): "An Empirical Model of Firm Entry with Endogenous Product-Type Choices," *The RAND Journal of Economics*, 37 (3), 619–640.

Thomas, Catherine (2011): "Too Many Products: Decentralized Decision Making in Multinational Firms," *American Economic Journal: Microeconomics*, 3 (1), 280–306.

Tintelnot, Felix (2016): "Global Production with Export Platforms*," *The Quarterly Journal of Economics*, 132 (1), 157–209.

Venables, Anthony (1987): "Trade and Trade Policy with Differentiated Products: A Chamberlinian-Ricardian Model," *Economic Journal*, 97 (387), 700–717.

Wollmann, Thomas G. (2018): "Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles," *American Economic Review*, 108 (6), 1364–1406.

ONLINE APPENDIX

Product Entry in the Global Automobile Industry

March 25, 2026

Alejandro Sabal

A Data Cleaning and Imputation of Product Characteristics

In this section, I explain the procedure to obtain the estimation sample from the raw data. I first describe preliminary data cleaning procedures. Then, I describe in detail how I aggregate and impute product characteristics to all potential products $\mathcal{A}$.

A.1 Preliminary Cleaning

The *IHS Markit* 2019 sample contains information on the universe of passenger vehicle registrations for 12 countries: Australia, Brazil, France, Germany, India, Italy, Japan, Mainland China, Mexico, Spain, the United Kingdom, and the United States. The data are at the quarterly-trim-country level.

I define 3 fuel type categories using information on fuel type in the dataset. First, I define "Electric" vehicles as any vehicle that fits in a plug-in category. For instance, I observe Plug-In Electric or Plug-in (Petrol or Diesel) vehicles. Any such fuel type will be considered part of the"Electric" category. Second, I define a "Hybrid" vehicle as any non-plug-in vehicle with a fuel type in the hybrid category. Finally, I define an internal combustion engine vehicle ("ICE") as any vehicle that is neither hybrid nor plug-in, be it Petrol or Diesel.

The dataset does not contain the vast majority of large pickups or vans, so I drop the small number of instances of such items from my sample. I use the SUV and Wagon categories to define the respective body types. I define the Convertible body type as including convertibles, retractable hardtops, and roadsters. Finally, I define the Car category as including either sedans, hatchbacks, or coupes. This gives 5 different possible body types.

In the sample, I observe some products in extreme price ranges. Since the goal of this paper is, in part, to estimate product development fixed costs, I drop the most luxurious products, which are likely designed and produced with very different technologies than the typical passenger vehicle. I drop all products with observed prices higher than $150,000 or

pertaining to the following brands: Lamborghini, McLaren, Bentley, Ferrari, Aston Martin, Maserati, Bugatti, and Rolls-Royce.

A.2 Demand Estimation Sample

I aggregate the data at the product-year level, where the product definition is a brand-body type-fuel type combination. I obtain total units sold by summing across all quarters and trims corresponding to my product definition. I aggregate the remaining characteristics (including horsepower, weight, length, width, and height) by taking a quantity-weighted average of such characteristics across all quarter-trim observations within my product definition. I merge this dataset with information from CEPII containing the distance between any two countries (population-weighted), which gives the distance between the headquarters country of the brand and the destination country. Market size is the product of the number of households and the average number of vehicles per household in market m, divided by the average tenure of car ownership, which I assume is 5 years, as in Grieco et al. (2023). The number of households in each country in 2019 and the average number of vehicles per household are obtained from the World Population Review and publicly available government sources (e.g., Department of Transportation).

I obtain market shares for each product in each market by dividing the total units sold by market size. As typical in the automobile literature that estimates mixed logit demand models, I drop very small market shares. More precisely, I drop all products with shares smaller than 0.00001 within a market. I only do this for demand and marginal cost estimation. In the estimation of fixed costs, I treat these products as part of the set of products in firms' global portfolios as well as in the set of product offerings in such markets.

A.3 Fixed Cost Estimation Sample and Product Characteristic Imputation

For the fixed cost estimation sample, I require a dataset at the (potential) product-market level that contains the observed characteristics of each product-market pair. This is what is required to use the demand and marginal cost model to predict variable profits under any arbitrary market structure for each possible entry opportunity. Given that I defined the set of potential products as the set of all brand-body type-fuel type combinations, I need to impute all observed characteristics that enter the demand and marginal cost specification for all potential products. These characteristics include horsepower, horsepower/weight, and size (length $\times$ width). To do this, I use the characteristics of the observed products.

I impute size for all products not observed in the *IHS Markit* sample according to the following sequential procedure: (1) using the mean size of observed products of the same body type sold by the same brand; (2) if there are no such products, I use the mean size of observed products of the same body type sold by the same parent company; (3) if there are no such products, I use the mean size across observed products of said body type.

I impute horsepower (and horsepower/weight) for all products not observed in the *IHS Markit* sample according to the following procedure: (1) using the mean horsepower (horsepower/weight) of observed products of the same fuel type and body type sold by the same parent company; (2) if there are no such products, I use the mean horsepower (horsepower/weight) of observed products with the same body type and fuel type offered in that country; (3) if there are no such products, I use the mean horsepower (horsepower/weight) of all products with the same fuel type and body type.

Finally, I treat a product as offered in a market ($O_{jm} = 1$) if and only if at least one unit is sold in that market. This is standard in previous papers studying product entry. I define a product as being in a firm's product portfolio ($G_j = 1$) if and only if it is observed to be sold in at least one market in my sample. This definition assumes that there are no products that are only offered in markets that are not included in my sample. This is a decent assumption, given that the markets in my sample account for most of the global demand. A caveat is that, in my model, a firm could add a product to its portfolio but face high market entry costs everywhere, leading it not to sell the product. Misclassifying such cases as $G_j = 0$ when $G_j = 1$ would understate the share of products introduced and likely overstate portfolio fixed costs. This bias should be minimal, since with large portfolio costs and independent entry-cost shocks, it is unlikely a firm would develop a product and then sell it nowhere.

B Proof of Theorem 1 (Moment Inequalities)

In this section, I derive the key results that I use for the estimation of $(\theta_e, \sigma_e, \theta_g, \sigma_g)$. In Section B.1, I derive moment inequalities that bound (θ_e, σ_e). In Section B.2, I derive moment inequalities that bound $(\theta_e, \sigma_e, \theta_g, \sigma_g)$. Recall from the main text that Γ_{jm} denotes the CDF of F^e_{jm}. I let Λ_j denote the CDF of F^g_j. I introduce notation and write, $\mathbb{1}_{\mathcal{C}^f} = \mathbb{1}\{\mathcal{C}^f \text{ is chosen by firm } f\}$, where $\mathcal{C}^f$ is a bundle of products chosen by firm f.

B.1 Stage 2: Market Entry Inequalities

I start by proving formally the upper bound inequality (9). At this stage, firm f has already chosen portfolio $\mathcal{G}^f$. By revealed preference and best response, for any $j \in \mathcal{G}^f$,

$$\begin{aligned}(\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}})\, \mathbb{1}\big\{\mathbb{E}[\Delta_{jm}(\Omega_m^f, \Omega_m^{-f}) \mid \{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}, \mathcal{I}, \mathcal{G}^f] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0\big\} \\ = (\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}})\, \mathbb{1}_{\Omega_m^f}. \end{aligned} \tag{16}$$

Equation (16) says that conditional on firm f choosing Ω_m^f or $\Omega_m^f \setminus \{j\}$, the firm chooses Ω_m^f if and only if it is weakly preferred to $\Omega_m^f \setminus \{j\}$.

Under Assumption 4, the largest possible change in expected profits from introducing product j in market m can be obtained by offering no product other than j in market m. Thus, I obtain inequality,

$$\begin{aligned}(\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}}) \mathbb{1}\{\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f}) | \{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}, \mathcal{I}, \mathcal{G}^f] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0\} \\ \geq (\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}}) \mathbb{1}_{\Omega_m^f}. \end{aligned} \tag{17}$$

In inequality (17), inside the indicator function, the expected marginal value is evaluated at a set of entry decisions that is not optimal and is therefore independent of $\{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}$ conditional on $\mathcal{I}$. Also, due to Assumption 1, rivals' product offerings Ω_m^{-f} are independent of $\{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}$ conditional on $\mathcal{I}$. Thus, removing conditioning on the shocks,

$$\begin{aligned}(\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}}) \mathbb{1}\{\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f}) | \mathcal{I}, \mathcal{G}^f] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0\} \\ \geq (\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}}) \mathbb{1}_{\Omega_m^f}. \end{aligned} \tag{18}$$

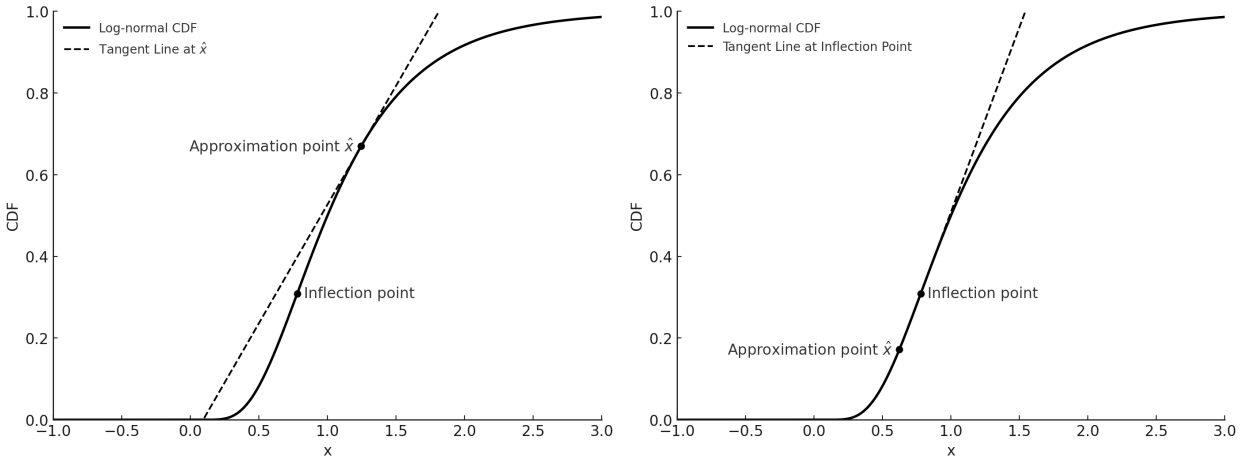
Note that inequality (18) holds for all bundles Ω_m^f containing product j, and that the term $\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f}) | \mathcal{I}, \mathcal{G}^f] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0\}$ is now independent of both Ω_m^f and $\Omega_m^f \setminus \{j\}$. Therefore, realizing that $O_{jm} = \sum_{\Omega_m^f : j \in \Omega_m^f} \mathbb{1}_{\Omega_m^f}$ and $1 - O_{jm} = \sum_{\Omega_m^f : j \in \Omega_m^f} \mathbb{1}_{\Omega_m^f \setminus \{j\}}$, I sum inequality (18) across all such mutually exclusive bundles to obtain,

$$\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f}) | \mathcal{I}, \mathcal{G}^f] - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0\} \geq O_{jm}. \tag{19}$$

That is, if product j is offered in market m, the largest possible change in profits from offering product j in market m must be weakly positive. The term $\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f}) | \mathcal{I}, \mathcal{G}^f]$ is independent of ν_{jm}^e, so I take expectations conditional on $\mathcal{I}$ and $\mathcal{G}^f$ to derive,

$$\Gamma_{jm}\big(\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f}) | \mathcal{I}, \mathcal{G}^f]; \theta_e, \sigma_e\big) \geq \mathbb{E}[O_{jm} | \mathcal{I}, \mathcal{G}^f], \tag{20}$$

Figure B.1: Convex Upper Bounds of a Log-Normal CDF



where recall that Γ_{jm} denotes the cumulative distribution function (CDF) of F^e_{jm} given θ_e and σ_e. Inequality (20) gives an upper bound on the probability that product j is offered in market m, conditional on $\mathcal{I}$ and $\mathcal{G}^f$. The key insight for estimation is to now use a convex upper bound of CDF Γ_{jm}. If F^e_{jm} is log-normal, it has inflection point $\check{x}_{jm}(\theta_e, \sigma_e) = \exp(Z'_{jm}\theta_e - \sigma_e^2)$, so one can define,

$$\begin{aligned}\overline{\Gamma}^1_{jm}(x; \theta_e, \sigma_e) &:= \Gamma_{jm}(x; \theta_e, \sigma_e)\mathbb{1}\{x < \check{x}_{jm}(\theta_e, \sigma_e)\} \\ &+ [\Gamma_{jm}(\check{x}; \theta_e, \sigma_e) + \gamma_{jm}(\check{x}_{jm}; \theta_e, \sigma_e)(x - \check{x}_{jm}(\theta_e, \sigma_e))]\mathbb{1}\{x \geq \check{x}_{jm}(\theta_e, \sigma_e)\} \qquad (21)\end{aligned}$$

where γ_{jm} denotes the probability density function (PDF) of the corresponding log-normal distribution. In the empirical implementation, I use,

$$\begin{aligned}\overline{\Gamma}_{jm}(x, \hat{x}; \theta_e, \sigma_e) &:= \overline{\Gamma}^1_{jm}(x; \theta_e, \sigma_e)\mathbb{1}\{\hat{x} < \check{x}_{jm}(\theta_e, \sigma_e)\} + \\ &\max\{\Gamma_{jm}(\hat{x}; \theta_e, \sigma_e) + \gamma_{jm}(\hat{x}; \theta_e, \sigma_e)(x - \hat{x}), \Gamma_{jm}(x; \theta_e, \sigma_e)\}\mathbb{1}\{\hat{x} \geq \check{x}_{jm}(\theta_e, \sigma_e)\}, \qquad (22)\end{aligned}$$

where $\hat{x}$ is a $\mathcal{I}$-measurable approximation point.

That is, if the $\mathcal{I}$-measurable approximation point lies below the inflection point, I use the convex upper bound given by equation (21). Otherwise, I use a linear approximation at the approximation point (which is on the concave part of the CDF), bounded below by the CDF, Γ_{jm}. Figure B.1 illustrates the convex upper bounds.[1] Given convex upper bound $\overline{\Gamma}_{jm}(x, \hat{x}_{jm}; \theta_e, \sigma_e)$, I derive,

$$\overline{\Gamma}_{jm}\big(\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f})|\mathcal{I}, \mathcal{G}^f], \hat{x}_{jm}; \theta_e, \sigma_e\big) \geq \mathbb{E}[O_{jm}|\mathcal{I}, \mathcal{G}^f].$$

[1]Note that these are global bounds and, unlike the CDF, take values larger than 1.

I now apply Jensen's inequality and obtain,

$$\mathbb{E}[\overline{\Gamma}_{jm}\big(\Delta_{jm}(\{j\},\Omega_m^{-f}),\hat{x}_{jm};\theta_e,\sigma_e\big)-O_{jm}|\mathcal{I},\mathcal{G}^f]\geq 0 \tag{23}$$

Inequality (23) is a conditional moment inequality that can be used for estimation. Given this conditional moment inequality, unconditional moment inequalities can be derived using positive functions of $\mathcal{I}$. Moreover, provided the approximation points $\hat{x}_{jm}$ are $\mathcal{I}$-measurable, one can use observation-specific upper bound functions $\overline{\Gamma}_{jm}$.[2]

To derive the lower bound inequality, I start in a manner similar to equation (16), by writing, for any bundle Ω_m^f such that $j\notin\Omega_m^f$,

$$\begin{aligned}&(\mathbb{1}_{\Omega_m^f\cup\{j\}}+\mathbb{1}_{\Omega_m^f})\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\Omega_m^f\cup\{j\},\Omega_m^{-f})-F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e)\leq 0\}\\&=(\mathbb{1}_{\Omega_m^f\cup\{j\}}+\mathbb{1}_{\Omega_m^f})\mathbb{1}_{\Omega_m^f}.\end{aligned} \tag{24}$$

Equation (24) says that conditional on firm f choosing either $\Omega_m^f\cup\{j\}$ or Ω_m^f, it will choose Ω_m^f if and only if it is preferred to $\Omega_m^f\cup\{j\}$.

Due to submodularity, the lowest possible expected change in profits from offering product j in market m is obtained whenever the firm is offering all products in its portfolio $\mathcal{G}^f$ in market m, which implies that,

$$\begin{aligned}(\mathbb{1}_{\Omega_m^f\cup\{j\}}+\mathbb{1}_{\Omega_m^f})\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\mathcal{G}^f,\Omega_m^{-f})|\{\nu_{jm}^e\}_{j\in\mathcal{G}^f,m},\mathcal{I},\mathcal{G}^f]-F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e)\leq 0\}\\\geq(\mathbb{1}_{\Omega_m^f\cup\{j\}}+\mathbb{1}_{\Omega_m^f})\mathbb{1}_{\Omega_m^f}.\end{aligned} \tag{25}$$

As with the upper bound inequality, I remove conditioning on $\{\nu_{jm}^e\}_{j\in\mathcal{G}^f,m}$ in the expectation in inequality (25) due to Assumption 1. Then, realizing that $\sum_{\Omega_m^f:j\notin\Omega_m^f}\mathbb{1}_{\Omega_m^f}=1-O_{jm}$ and $\sum_{\Omega_m^f:j\notin\Omega_m^f}\mathbb{1}_{\Omega_m^f\cup\{j\}}=O_{jm}$, I sum inequality (25) over all bundles Ω_m^f with $j\notin\Omega_m^f$ to obtain,

$$\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\mathcal{G}^f,\Omega_m^{-f})|\mathcal{I},\mathcal{G}^f]-F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e)\leq 0\}\geq 1-O_{jm}. \tag{26}$$

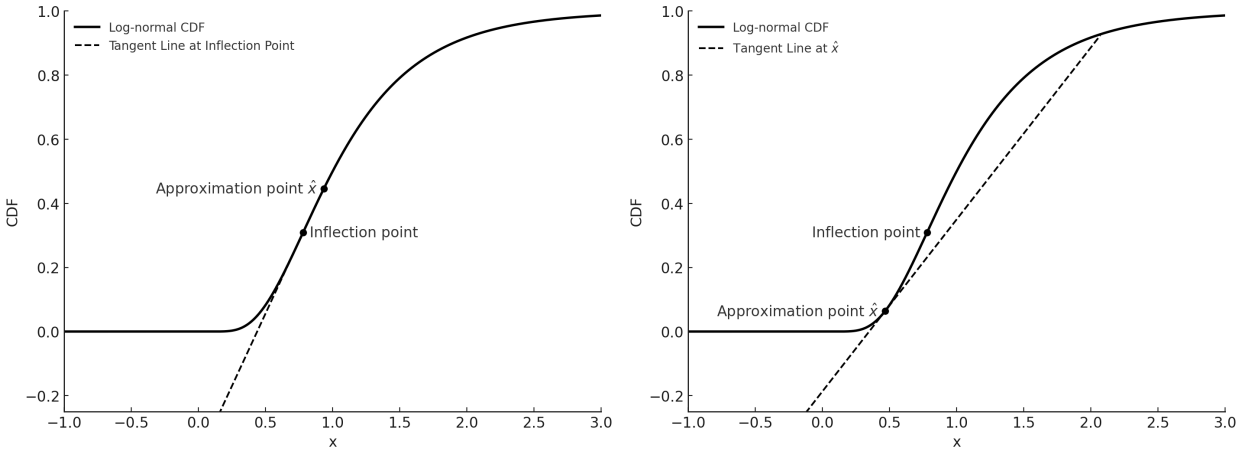
Taking expectations conditional on $\mathcal{I}$ and $\mathcal{G}^f$ on both sides of inequality (26) yields,

$$\Gamma_{jm}\big(\mathbb{E}[\Delta_{jm}(\mathcal{G}^f,\Omega_m^{-f})|\mathcal{I},\mathcal{G}^f];\theta_e,\sigma_e\big)\leq\mathbb{E}[O_{jm}|\mathcal{I},\mathcal{G}^f]. \tag{27}$$

Inequality (27) provides a lower bound on the probability that product j is offered in market m, conditional on $\mathcal{I}$ and $\mathcal{G}^f$. I now follow a very similar logic as with the upper

[2] Porcher et al. (2024) also employ observation-specific linear approximations to derive moment inequalities, though in a single-agent setting and using odds-based inequalities rather than bounding choice probabilities using convex upper and concave lower bounds of the CDF of the unobserved shock.

Figure B.2: Concave Lower Bounds of a Log-Normal CDF



bound and use a *concave lower bound* for Γ_{jm} before applying Jensen's inequality. A family of such concave lower bounds is given by,

$$\begin{aligned}\underline{\Gamma}^1_{jm}(x;\theta_e,\sigma_e) &:= \Gamma(x;\theta_e,\sigma_e)\mathbb{1}\{x \geq \check{x}_{jm}(\theta_e,\sigma_e)\} \\ &+ [\Gamma_{jm}(\check{x};\theta_e,\sigma_e) + \gamma_{jm}(\check{x}_{jm};\theta_e,\sigma_e)(x - \check{x}_{jm}(\theta_e,\sigma_e))]\mathbb{1}\{x < \check{x}_{jm}(\theta_e,\sigma_e)\} \qquad (28)\end{aligned}$$

where recall that the inflection point is given by, $\check{x}_{jm}(\theta_e,\sigma_e) = \exp(Z'_{jm}\theta_e - \sigma_e^2)$. In the empirical implementation, I use,

$$\begin{aligned}\underline{\Gamma}_{jm}(x,\hat{x};\theta_e,\sigma_e) &:= \underline{\Gamma}^1_{jm}(x;\theta_e,\sigma_e)\mathbb{1}\{\hat{x} \geq \check{x}_{jm}(\theta_e,\sigma_e)\} \\ &+ \min\{\Gamma_{jm}(\hat{x};\theta_e,\sigma_e) + \gamma_{jm}(\hat{x};\theta_e,\sigma_e)(x - \hat{x}), \Gamma_{jm}(x;\theta_e,\sigma_e)\}\mathbb{1}\{\hat{x} < \check{x}_{jm}(\theta_e,\sigma_e)\}. \qquad (29)\end{aligned}$$

That is, if the approximation point lies above the inflection point, I use the concave lower bound given in equation (28). Otherwise, I use a linear approximation at the approximation point (which is in the convex part of the CDF), bounded above by the CDF, Γ_{jm}. Figure B.2 illustrates the concave lower bound functions.

Given such a concave lower bound, together with Jensen's inequality, I obtain

$$\mathbb{E}[\underline{\Gamma}_{jm}(\Delta_{jm}(\mathcal{G}^f, \Omega_m^{-f}), \hat{x}_{jm};\theta_e,\sigma_e) - O_{jm}|\mathcal{I},\mathcal{G}^f] \leq 0. \qquad (30)$$

As with the upper bound, conditional moment inequality (30) can now be used for estimation given positive-valued functions of $\mathcal{I}$, which yield unconditional moment inequalities. Also, as before, provided $\hat{x}_{jm}$ is $\mathcal{I}$-measurable, one can use observation-dependent lower bound functions $\underline{\Gamma}_{jm}$.

Note that because by assumption, market entry fixed cost shocks $\{\nu^e_{jm}\}_{j\in\mathcal{G}^f,m}$ are unobserved at the time of choosing $\mathcal{G}^f$, I condition on the observed product portfolios when using conditional moment inequalities (23) and (30) for estimation.

This concludes the derivation of the inequalities used to estimate bounds on (θ_e, σ_e). In the next subsection, I derive inequalities that provide bounds on (θ_g, σ_g).

B.2 Stage 1: Global Portfolio Choice Inequalities

At this stage, firms choose their global product portfolios $\mathcal{G}^f$. The derivation of the inequalities at this stage follows a similar logic as the derivations in Stage 2. For the derivations that follow, I will use the following result:

Proposition 1 *The function $g : \mathbb{R} \to \mathbb{R}$ given by,*

$$g(y) = \mathbb{E}_X[\mathbb{1}(X \leq y)(y - X)] = \mathbb{E}_X[y - X|X \leq y]\mathbb{P}_X(X \leq y)$$

is convex in y for any continuous random variable X provided $F_X(y) := \mathbb{P}_X(X \leq y) > 0$ i.e., the conditional expectation is well defined.

Proof. Differentiating with respect to y, I obtain,

$$\begin{aligned} g'(y) &= F_X(y)\left[1 - \frac{f_X(y)}{F_X(y)}(y - \mathbb{E}[X|X \leq y])\right] + \mathbb{E}[y - X|X \leq y]f_X(y) \\ &= F_X(y) \end{aligned}$$

where f_X denotes the density of random variable X. But then, clearly, $g''(y) = f(y) > 0$, which proves convexity. ■

To derive the inequalities at this stage, I first define the value of a given portfolio for a firm f as,

$$\begin{aligned} \mathcal{V}_f(\mathcal{G}^f, \{\nu^g_j\}_{j\in\mathcal{A}^f}, \mathcal{I}) = \mathbb{E}\Bigg[&\sum_{m\in\mathcal{M}} \max_{\{\Omega^f_m \subseteq \mathcal{G}^f\}} \Pi^f_m(\Omega^f_m; \mathcal{G}^f, \{\nu^e_{jm}\}_{j\in\mathcal{G}^f,m}, \mathcal{I}) \Bigg| \mathcal{I} \Bigg] \\ &- \sum_{j\in\mathcal{A}^f} G_j F^g_j(\nu^g_j; \theta_g, \sigma_g), \end{aligned} \tag{31}$$

where

$$\Pi^f_m(\Omega^f_m; \mathcal{G}^f, \{\nu^e_{jm}\}_{j\in\mathcal{G}^f,m}, \mathcal{I})$$

$$= \mathbb{E}\left[\sum_{j \in \Omega_m^f} [\pi_{jm}^*(\Omega_m^f, \Omega_m^{-f}) - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e)] \middle| \mathcal{G}^f, \{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}, \mathcal{I} \right].$$

Given any chosen product portfolio $\mathcal{G}^f$ and realized global portfolio fixed cost shocks, firm f computes its Stage 1 value as the expected maximal profits it will obtain once it realizes its market entry fixed cost shocks for each product and chooses offerings in each market optimally. The expectation in equation (31) is only conditional on $\mathcal{I}$, so the firm must also integrate over its own market entry fixed costs shocks. At Stage 2, firm f's optimal entry decisions depend on the portfolio it chooses in Stage 1 and the market entry fixed cost shocks it realizes in Stage 2. I write the Stage-2 decision rule in each market as $\Omega_m^f(\mathcal{G}^f, \{\nu_{jm}^e\}_{j \in \mathcal{G}^f}, \mathcal{I})$.

Henceforth, I abuse notation and let $\Omega_m^{\mathcal{G}^f}$ denote the decision rule under portfolio $\mathcal{G}^f$ (at given market entry fixed cost shocks and at $\mathcal{I}$) and $O_{jm}^{\mathcal{G}^f}$ denote the optimal entry decision rule for product j in market m under $\mathcal{G}^f$ corresponding to $\Omega_m^{\mathcal{G}^f}$.

I start by deriving an upper bound inequality. First, note that,

$$\begin{aligned}
&(\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f \setminus \{j\}}) \mathbb{1}\left\{ \mathcal{V}_f(\mathcal{G}^f, \{\nu_j^g\}_{j \in \mathcal{A}^f}, \mathcal{I}) - \mathcal{V}_f(\mathcal{G}^f \setminus \{j\}, \{\nu_j^g\}_{j \in \mathcal{A}^f}, \mathcal{I}) \geq 0 \right\} \\
&= (\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f \setminus \{j\}}) \mathbb{1}_{\mathcal{G}^f}. \qquad (32)
\end{aligned}$$

That is, conditional on choosing either portfolio $\mathcal{G}^f$ or portfolio $\mathcal{G}^f \setminus \{j\}$, the firm chooses $\mathcal{G}^f$ if and only if it is preferred to $\mathcal{G}^f \setminus \{j\}$. A lower bound of $\mathcal{V}_f(\mathcal{G}^f \setminus \{j\}, \{\nu_j^g\}_{j \in \mathcal{A}^f}, \mathcal{I})$ can be obtained by using the fact that the entry decisions in all markets for products $j \in \mathcal{G}^f \setminus \{j\}$ must weakly dominate the optimal-under-$\mathcal{G}^f$ entry decision rules for such products. This means that the equality (32) implies the following inequality,

$$\begin{aligned}
&(\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f \setminus \{j\}}) \mathbb{1}\left\{ \sum_{m \in \mathcal{M}} \left(\mathbb{E}[\Pi_m^f(\Omega_m^{\mathcal{G}^f}) - \Pi_m^f(\Omega_m^{\mathcal{G}^f} \setminus \{j\}) | \mathcal{I}] \right) - F_j^g(\nu_j^g; \theta_g, \sigma_g) \geq 0 \right\} \\
&\qquad \geq (\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f \setminus \{j\}}) \mathbb{1}_{\mathcal{G}^f}. \qquad (33)
\end{aligned}$$

Equation (33) says that the change in value from including j into the portfolio starting from a suboptimal entry decision rule in the second stage of the game – which mandates choosing the *same* entry decisions in each market under $\mathcal{G}^f \setminus \{j\}$ as under $\mathcal{G}^f$ for all products $j' \neq j$ – must be higher than the actual change in expected value from introducing product j into the firm's portfolio at the best response. I re-write inequality (33) as,

$$(\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f \setminus \{j\}}) \times$$

$$\mathbb{1}\Big\{ \sum_{m\in\mathcal{M}} \big(\mathbb{E}\big[O_{jm}^{\mathcal{G}^f}\big[\mathbb{E}\big[\Delta_{jm}(\Omega_m^{\mathcal{G}^f}\setminus\{j\},\Omega_m^{-f})|\mathcal{I},\{\nu_{jm}^e\}_{j\in\mathcal{G}^f,m},\mathcal{G}^f\big] - F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e)\big|\mathcal{I}\big]\big)$$
$$- F_j^g(\nu_j^g;\theta_g,\sigma_g)\geq 0\Big\} \geq (\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f\setminus\{j\}})\mathbb{1}_{\mathcal{G}^f}. \tag{34}$$

Inequality (34) shows that if product j is introduced in firm f's global portfolio, then the change in profits from offering it in each market in which firm f chooses to offer product j (evaluated at the optimal entry decisions under $\mathcal{G}^f$) is weakly positive net of the portfolio fixed cost. Under Assumption 4, I bound the expression inside the indicator function in inequality (34) by above and make it independent of the optimal portfolio choice. I obtain,

$$(\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f\setminus\{j\}})\times$$
$$\mathbb{1}\Big\{ \sum_{m\in\mathcal{M}} \big(\mathbb{E}\big[O_{jm}^{\mathcal{G}^f}\big[\mathbb{E}\big[\Delta_{jm}(\{j\},\Omega_m^{-f})|\mathcal{I},\{\nu_{jm}^e\}_{j\in\mathcal{G}^f,m},\mathcal{G}^f\big] - F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e)\big|\mathcal{I}\big]\big)$$
$$- F_j^g(\nu_j^g;\theta_g,\sigma_g)\geq 0\Big\} \geq (\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f\setminus\{j\}})\mathbb{1}_{\mathcal{G}^f}. \tag{35}$$

Therefore,

$$(\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f\setminus\{j\}})\mathbb{1}\Big\{ \sum_{m\in\mathcal{M}} \big(\mathbb{E}\big[O_{jm}^{\{j\}}\big[\mathbb{E}\big[\Delta_{jm}(\{j\},\Omega_m^{-f})|\mathcal{I}\big] - F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e)\big|\mathcal{I}\big]\big)$$
$$- F_j^g(\nu_j^g;\theta_g,\sigma_g)\geq 0\Big\} \geq (\mathbb{1}_{\mathcal{G}^f} + \mathbb{1}_{\mathcal{G}^f\setminus\{j\}})\mathbb{1}_{\mathcal{G}^f}, \tag{36}$$

where

$$O_{jm}^{\{j\}} = \mathbb{1}\big\{\mathbb{E}[\Delta_{jm}(\{j\},\Omega_m^{-f})|\mathcal{I}] - F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e)\geq 0\big\}.$$

It follows that the contribution of product j into firm f's Stage 1 value has to be smaller than the contribution if it were the only product the firm sold in each market, conditional on positive profits from product j in each market net of the fixed market entry cost.

Summing all inequalities of the form of inequality (36) across all bundles containing product j yields,

$$\mathbb{1}\Big\{ \sum_{m\in\mathcal{M}} \big(\mathbb{E}\big[O_{jm}^{\{j\}}\big[\mathbb{E}\big[\Delta_{jm}(\{j\},\Omega_m^{-f})|\mathcal{I}\big] - F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e)\big|\mathcal{I}\big]\big)$$
$$- F_j^g(\nu_j^g;\theta_g,\sigma_g)\geq 0\Big\} \geq G_j. \tag{37}$$

The expectation inside the indicator function in inequality (37) is an expectation such as that in Proposition 1.[3] Applying Jensen's inequality gives,

$$\mathbb{1}\Big\{\mathbb{E}\Big[\sum_{m\in\mathcal{M}}\Gamma_{jm}\big(\Delta_{jm}(\{j\},\Omega_m^{-f})\big)\big[\Delta_{jm}(\{j\},\Omega_m^{-f}) \\ -\mathbb{E}[F^e_{jm}(\nu^e_{jm})|\mathcal{I},F^e_{jm}(\nu^e_{jm})\leq\Delta_{jm}(\{j\},\Omega_m^{-f})]\big]\big|\mathcal{I}\Big]-F^g_j(\nu^g_j;\theta_g,\sigma_g)\geq 0\Big\}\geq G_j. \tag{38}$$

Henceforth, the analysis is identical to the derivation of the Stage 2 inequalities. As before, ν^g_j is now independent of the conditional expectation inside the indicator function in inequality (38). I take expectations on both sides conditional on $\mathcal{I}$ to obtain,

$$\Lambda_j\Big(\mathbb{E}\Big[\sum_{m\in\mathcal{M}}\Gamma_{jm}\big(\Delta_{jm}(\{j\},\Omega_m^{-f})\big)\big[\Delta_{jm}(\{j\},\Omega_m^{-f}) \\ -\mathbb{E}[F^e_{jm}(\nu^e_{jm})|\mathcal{I},F^e_{jm}(\nu^e_{jm})\leq\Delta_{jm}(\{j\},\Omega_m^{-f})]\big]|\mathcal{I}\Big];\theta_g,\sigma_g\Big)\geq\mathbb{E}[G_j|\mathcal{I}]. \tag{39}$$

As with the Stage 2 inequalities, I use the convex upper bound of Λ_j of the form of that in equation (22), at an $\mathcal{I}$-measurable approximation point $\hat{x}_j$, such that, applying Jensen's inequality,

$$\mathbb{E}\Big[\overline{\Lambda}_j\Big(\sum_{m\in\mathcal{M}}\Gamma_{jm}\big(\Delta_{jm}(\{j\},\Omega_m^{-f})\big)\big[\Delta_{jm}(\{j\},\Omega_m^{-f}) \\ -\mathbb{E}[F^e_{jm}(\nu^e_{jm})|\mathcal{I},F^e_{jm}(\nu^e_{jm})\leq\Delta_{jm}(\{j\},\Omega_m^{-f})]\big],\hat{x}_j;\theta_g,\sigma_g\Big)-G_j|\mathcal{I}\Big]\geq 0. \tag{40}$$

Inequality (40) can now be used to estimate θ_g and σ_g. Note that $\mathbb{E}\big[F^e_{jm}(\nu^e_{jm})|\mathcal{I},F^e_{jm}(\nu^e_{jm})\leq\Delta_{jm}(\{j\},\Omega_m^{-f})\big]$ can be computed using numerical integration or Gaussian quadrature. I use the QuadGK Julia package to evaluate this conditional mean.

To derive a lower bound conditional moment inequality, I start from the fact that the following equality must hold at the best response for any $\mathcal{G}^f$ with $j\notin\mathcal{G}^f$,

$$(\mathbb{1}_{\mathcal{G}^f\cup\{j\}}+\mathbb{1}_{\mathcal{G}^f})\mathbb{1}\big\{\mathcal{V}_f(\mathcal{G}^f\cup\{j\},\{\nu^g_j\}_{j\in\mathcal{A}^f},\mathcal{I})-\mathcal{V}_f(\mathcal{G}^f,\{\nu^g_j\}_{j\in\mathcal{A}^f},\mathcal{I})\leq 0\big\} \\ =(\mathbb{1}_{\mathcal{G}^f\cup\{j\}}+\mathbb{1}_{\mathcal{G}^f})\mathbb{1}_{\mathcal{G}^f}. \tag{41}$$

Equation (41) says that conditional on firm f choosing either $\mathcal{G}^f\cup\{j\}$ or $\mathcal{G}^f$, it chooses $\mathcal{G}^f$ if and only if $\mathcal{G}^f$ is preferred to $\mathcal{G}^f\cup\{j\}$. Consider now the following sub-optimal second-stage decision rules under portfolio $\mathcal{G}^f\cup\{j\}$ in all markets m,

[3] I lighten notation by removing the (θ_e,σ_e) arguments from Γ_{jm} and F^e_{jm}.

$$\underline{\Omega}_m^{\mathcal{G}^f \cup \{j\}} = \begin{cases} O_{j'm}^{\mathcal{G}^f}, & j' \neq j \\ \underline{O}_{j'm}, & j' = j. \end{cases}$$

i.e., firm f chooses the decision rule used for products $\mathcal{G}^f$ under portfolio $\mathcal{G}^f \cup \{j\}$ for all non-j products and uses the $\underline{O}_{jm}$ decision rule for product j, where for now this decision rule is any arbitrary decision rule that must depend on $\mathcal{I}$ and any set of realized market entry fixed cost realizations. Combining these sub-optimal decision rules with equation (41),

$$\begin{aligned} &(\mathbb{1}_{\mathcal{G}^f \cup \{j\}} + \mathbb{1}_{\mathcal{G}^f}) \times \\ &\mathbb{1}\Big\{ \sum_{m \in \mathcal{M}} \mathbb{E}\big[\Pi_m^f\big(\underline{\Omega}_m^{\mathcal{G}^f \cup \{j\}}\big) - \Pi_m^f\big(\Omega_m^{\mathcal{G}^f}\big) | \mathcal{I}\big] - F_j^g(\nu_j^g; \theta_g, \sigma_g) < 0 \Big\} \geq (\mathbb{1}_{\mathcal{G}^f \cup \{j\}} + \mathbb{1}_{\mathcal{G}^f}) \mathbb{1}_{\mathcal{G}^f}, \end{aligned}$$

or

$$\begin{aligned} (\mathbb{1}_{\mathcal{G}^f \cup \{j\}} + \mathbb{1}_{\mathcal{G}^f}) \mathbb{1}\Big\{ \sum_{m \in \mathcal{M}} &\mathbb{E}\big[\underline{O}_{jm}\big[\mathbb{E}\big[\Delta_{jm}(\Omega_m^{\mathcal{G}^f}, \Omega_m^{-f}) | \mathcal{I}, \{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}, \mathcal{G}^f\big] - F_{jm}^e(\nu_{jm}^e) | \mathcal{I}\big] \\ &- F_j^g(\nu_j^g; \theta_g, \sigma_g) < 0 \Big\} \geq (\mathbb{1}_{\mathcal{G}^f \cup \{j\}} + \mathbb{1}_{\mathcal{G}^f}) \mathbb{1}_{\mathcal{G}^f}. \end{aligned} \tag{42}$$

As with the upper bound, I use submodularity of variable profits to further bound the expression inside the indicator function in inequality (42). In particular,

$$\begin{aligned} (\mathbb{1}_{\mathcal{G}^f \cup \{j\}} + \mathbb{1}_{\mathcal{G}^f}) \mathbb{1}\Big\{ \sum_{m \in \mathcal{M}} &\mathbb{E}\big[\underline{O}_{jm}\big[\mathbb{E}\big[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f}) | \mathcal{I}\big] - F_{jm}^e(\nu_{jm}^e) | \mathcal{I}\big] \\ &- F_j^g(\nu_j^g; \theta_g, \sigma_g) < 0 \Big\} \geq (\mathbb{1}_{\mathcal{G}^f \cup \{j\}} + \mathbb{1}_{\mathcal{G}^f}) \mathbb{1}_{\mathcal{G}^f}. \end{aligned} \tag{43}$$

Inequality (43) holds for all bundles $\mathcal{G}^f$ that do not contain product j. Thus, summing all such inequalities across all bundles $\mathcal{G}^f$ that do not contain product j yields,

$$\mathbb{1}\Big\{ \sum_{m \in \mathcal{M}} \mathbb{E}\big[\underline{O}_{jm}\big[\mathbb{E}\big[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f}) | \mathcal{I}\big] - F_{jm}^e(\nu_{jm}^e) | \mathcal{I}\big] - F_j^g(\nu_j^g; \theta_g, \sigma_g) \geq 0 \Big\} \leq G_j. \tag{44}$$

Next, let,

$$\underline{O}_{jm} = \mathbb{1}\big\{\mathbb{E}[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f}) | \mathcal{I}] - F_{jm}^e(\nu_{jm}^e) \geq 0\big\}.$$

Under convexity (Proposition 1), and letting $y := \Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})$, a lower bound of

$$g(\mathbb{E}(y|\mathcal{I})) := \mathbb{E}[\mathbb{1}\{\mathbb{E}(y|\mathcal{I}) - F_{jm}^e(\nu_{jm}^e) \geq 0\}[\mathbb{E}(y|\mathcal{I}) - F_{jm}^e(\nu_{jm}^e)]|\mathcal{I}]$$

is obtained by taking a first-order approximation around some $\mathcal{I}$-measurable point $\hat{x}_{jm}$. This yields,

$$\begin{aligned} g(\mathbb{E}(y|\mathcal{I})) &\geq \Gamma_{jm}(\hat{x}_{jm})[\mathbb{E}[y|\mathcal{I}] - \mathbb{E}[F^e_{jm}(\nu^e_{jm})|F^e_{jm} \leq \hat{x}_{jm}, \mathcal{I}]] \\ &= \Gamma_{jm}(\hat{x}_{jm})[\mathbb{E}[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})|\mathcal{I}] - \mathbb{E}[F^e_{jm}(\nu^e_{jm})|F^e_{jm}(\nu^e_{jm}) \leq \hat{x}_{jm}, \mathcal{I}]]. \end{aligned} \tag{45}$$

Plugging inequality (45) into inequality (44), I obtain,

$$\begin{aligned} \mathbb{1}\Big\{\mathbb{E}\Big[\sum_{m\in\mathcal{M}} \Gamma_{jm}(\hat{x}_{jm})\left[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f}) - \mathbb{E}[F^e_{jm}(\nu^e_{jm})|F^e_{jm}(\nu^e_{jm}) \leq \hat{x}_{jm}, \mathcal{I}]\right]|\mathcal{I}\Big] \\ - F^g_j(\nu^g_j; \theta_g, \sigma_g) \geq 0\Big\} \leq G_j \end{aligned} \tag{46}$$

where recall Γ_{jm} was the CDF of F^e_{jm} (given θ_e and σ_e). Given that the expectation term is now independent of ν^g_j conditional on $\mathcal{I}$, I take expectations conditional on $\mathcal{I}$ and obtain,

$$\begin{aligned} \Lambda_j\Big(\mathbb{E}\Big[\sum_{m\in\mathcal{M}} \Gamma_{jm}(\hat{x}_{jm})\left[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})\right. \\ \left. - \mathbb{E}[F^e_{jm}(\nu^e_{jm})|F^e_{jm}(\nu^e_{jm}) \leq \hat{x}_{jm}, \mathcal{I}]\right]|\mathcal{I}\Big]; \theta_g, \sigma_g\Big) \leq \mathbb{E}[G_j|\mathcal{I}]. \end{aligned} \tag{47}$$

Finally, I use the concave lower bound (as in equation (29)) $\underline{\Lambda}_j$ at an $\mathcal{I}$-measurable approximation point $\hat{x}_j$ and apply Jensen's inequality to obtain,

$$\begin{aligned} \mathbb{E}\Big[\underline{\Lambda}_j\Big(\sum_{m\in\mathcal{M}} \Gamma_{jm}(\hat{x}_{jm})\left[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})\right. \\ \left. - \mathbb{E}[F^e_{jm}(\nu^e_{jm})|F^e_{jm}(\nu^e_{jm}) \leq \hat{x}_{jm}, \mathcal{I}]\right], \hat{x}_j; \theta_g, \sigma_g\Big) - G_j|\mathcal{I}\Big] \leq 0. \end{aligned} \tag{48}$$

I have now shown how to derive the inequalities that can be used for estimation of parameters θ_g and σ_g. These are given by inequalities (40) and (48). In sum, I have proven Theorem 1.

C Solution Method Based on Inequalities

In this section, I provide a method to bound the equilibrium distribution of product offerings in each market. The method relies on first-order stochastic dominance for multivariate random vectors. Before defining first-order stochastic dominance, I define a partial order in $\mathbb{R}^n$. Throughout, I employ the standard partial order on $\mathbb{R}^n$: $\boldsymbol{x} \geq \boldsymbol{y}$ if and only if $x_i \geq y_i$ for

all $i \in \{1, ..., n\}$. An upper set in $\mathbb{R}^n$ is any set U of the form $U(\boldsymbol{y}) = \{\boldsymbol{x} : \boldsymbol{x} \geq \boldsymbol{y}\}$.

Definition 2 *Let $\boldsymbol{X}$ and $\boldsymbol{Y}$ be two random vectors in $\mathbb{R}^n$ such that,*

$$\mathbb{P}(\boldsymbol{X} \in U) \geq \mathbb{P}(\boldsymbol{Y} \in U) \qquad \text{for all upper sets } U \subseteq \mathbb{R}^n.$$

Then $\boldsymbol{X}$ is said to first-order stochastically dominate $\boldsymbol{Y}$.

With Bernoulli random vectors, $\boldsymbol{X}$ first order stochastically dominating $\boldsymbol{Y}$ amounts to,

$$\mathbb{P}\Big(\bigcap_{i\in\Omega}\{X_i = 1\}\Big) \geq \mathbb{P}\Big(\bigcap_{i\in\Omega}\{Y_i = 1\}\Big) \qquad \text{for any } \Omega \subseteq \{1, ..., n\}.$$

I will also employ the following result.

Theorem 3 (Shaked 2007) *$\boldsymbol{X}$ FOSD $\boldsymbol{Y}$ if and only if $\mathbb{E}[u(\boldsymbol{X})] \geq \mathbb{E}[u(\boldsymbol{Y})]$ for all non-decreasing functions u for which the expectations exist.*

In counterfactual exercises, I am interested in learning about changes in the equilibrium market structure and product offerings in response to policies subsumed in $\mathcal{I}$. Notice that a change in $\mathcal{I}$ could be a subsidy, a tax, a change in the ownership structure of firms, or any other change in market conditions that is known at the time of making portfolio and market offerings choices. To obtain counterfactual bounds to these entry probabilities, I use Bayes' rule to write,

$$\mathbb{P}(O_{jm} = 1|\mathcal{I}) = \mathbb{P}(O_{jm} = 1|\mathcal{I}, G_j = 1)\mathbb{P}(G_j = 1|\mathcal{I}).$$

In deriving the moment inequalities, I showed that inequalities (19) and (26) hold. These inequalities state that conditional on product $j \in \mathcal{G}^f$, $\underline{O}^*_{jm} \leq O^*_{jm} \leq \overline{O}^*_{jm}$, where

$$\overline{O}^*_{jm} := \mathbb{1}\{\mathbb{E}_{\boldsymbol{\mu}^*_m}[\Delta_{jm}(\{j\}, \Omega^{-f}_m)|\mathcal{I}] - F^e_{jm}(\nu^e_{jm}) \geq 0\} \tag{49}$$

$$\underline{O}^*_{jm} := \mathbb{1}\{\mathbb{E}_{\boldsymbol{\mu}^*_m}[\Delta_{jm}(\mathcal{A}^f, \Omega^{-f}_m)|\mathcal{I}] - F^e_{jm}(\nu^e_{jm}) \geq 0\}, \tag{50}$$

Moreover, recall from (37) and (44) that, $\underline{G}^*_j \leq G^*_j \leq \overline{G}^*_j$, where

$$\overline{G}^*_j := \mathbb{1}\Big\{\sum_{m\in\mathcal{M}} \mathbb{E}\big[\overline{O}^*_{jm}[\mathbb{E}_{\boldsymbol{\mu}^*_m}[\Delta_{jm}(\{j\}, \Omega^{-f}_m)|\mathcal{I}] - F^e_{jm}(\nu^e_{jm})|\mathcal{I}\big] - F^g_j(\nu^g_j) \geq 0\Big\}, \tag{51}$$

$$\underline{G}^*_j := \mathbb{1}\Big\{\sum_{m\in\mathcal{M}} \mathbb{E}\big[\underline{O}^*_{jm}\big[\mathbb{E}_{\boldsymbol{\mu}^*_m}[\Delta_{jm}(\mathcal{A}^f, \Omega^{-f}_m)|\mathcal{I}] - F^e_{jm}(\nu^e_{jm})|\mathcal{I}\big] - F^g_j(\nu^g_j) \geq 0\Big\}. \tag{52}$$

In equations (49)-(52), $\boldsymbol{\mu}_m^*$ denotes the distribution of firms' offerings decisions in market m at whichever equilibrium emerges under $\mathcal{I}$.[4]

Lemma 1 *Under Assumptions 1-4, for each market $m \in \mathcal{M}$, the sequences $\{\underline{O}_{jm}^*\}_{j\in\mathcal{A}}$, $\{\overline{O}_{jm}^*\}_{j\in\mathcal{A}}$, $\{\underline{G}_j^*\}_{j\in\mathcal{A}}$, and $\{\overline{G}_j^*\}_{j\in\mathcal{A}}$ are sequences of independent random variables conditional on $\mathcal{I}$ bounding the equilibrium binary decisions $\{O_{jm}^*\}_{j\in\mathcal{A}}$ and $\{G_j^*\}_{j\in\mathcal{A}}$.*

Proof. The fact that $\underline{G}_j^* \leq G_j^* \leq \overline{G}_j^*$ and $\underline{O}_{jm}^* \leq O_{jm}^* \leq \overline{O}_{jm}^*$ follow from the derivations leading to inequalities (37) and (36), and (19) and (26), respectively. Assumptions 2-3 impose conditional independence, so the independence result is immediate. ■

Importantly, the true equilibrium sequences are *not* independent sequences of random variables given that firms' product introduction choices are correlated due to interdependencies coming from cannibalization. The asterisks in the definitions of the bounding random variables in (49)-(52) denote that firms' expectations are with respect to the *true* joint distribution of entry decisions made by other firms. Firms care about which products are ultimately sold in each market and, therefore, about the joint distribution (across products and markets) of random variables $\{V_{jm}^*\}_{j\in\mathcal{A}}$, where $V_{jm}^* := O_{jm}^* G_j^* = O_{jm}^*$, that is, the event that product j is offered in market m, which requires it being developed and being offered. The latter equality holds because, in equilibrium, $O_{jm}^* = 1$ implies $G_j^* = 1$. Then, $\underline{V}_{jm}^* \leq V_{jm}^* \leq \overline{V}_{jm}^*$, where $\underline{V}_{jm}^* := \underline{O}_{jm}^* \underline{G}_j^*$ and $\overline{V}_{jm}^* := \overline{O}_{jm}^* \overline{G}_j^*$. The independence of first and second-stage shocks implies that $\overline{O}_{jm}^*$ and $\overline{G}_j^*$ as well as $\underline{O}_{jm}^*$ and $\underline{G}_j^*$, are independent conditional on $\mathcal{I}$. A corollary of Lemma 1 is that within each market m, $\{\underline{V}_{jm}^*\}_{j\in\mathcal{A}}$ and $\{\overline{V}_{jm}^*\}_{j\in\mathcal{A}}$ are conditionally (on $\mathcal{I}$) independent sequences of random variables. Therefore, conditional on $\mathcal{I}$, it follows that,

$$\begin{aligned}\mathbb{P}(\underline{O}_{jm}^* = 1|\mathcal{I})\mathbb{P}(\underline{G}_j^* = 1|\mathcal{I}) = \mathbb{P}(\underline{V}_{jm}^* = 1|\mathcal{I}) &\leq \mathbb{P}(O_{jm}^* = 1|\mathcal{I}) \\ &\leq \mathbb{P}(\overline{V}_{jm}^* = 1|\mathcal{I}) = \mathbb{P}(\overline{O}_{jm}^* = 1|\mathcal{I})\mathbb{P}(\overline{G}_j^* = 1|\mathcal{I}). \qquad (53)\end{aligned}$$

So, given any equilibrium distribution of product offerings, the inequalities (53) must hold. By Lemma 1, we also have that,

$$\mathbb{E}[\prod_{j\in\mathcal{W}} \underline{V}_{jm}^*|\mathcal{I}] \leq \mathbb{E}[\prod_{j\in\mathcal{W}} V_{jm}^*|\mathcal{I}] \leq \mathbb{E}[\prod_{j\in\mathcal{W}} \overline{V}_{jm}^*|\mathcal{I}], \qquad (54)$$

given any set of products $\mathcal{W}$. We will use this result to bound not only the marginal probabilities that products are offered, but also the joint distribution of product offerings in each market.

[4]Firm f only integrates over rival firms' offerings decisions. I simplify notation by using a $\boldsymbol{\mu}_m^*$ subscript, rather than a $\boldsymbol{\mu}_m^{*,-f}$ subscript.

C.1 Proof of Theorem 2

First, define:

$$\overline{O}_{jm}^{k} := \mathbb{1}\{\mathbb{E}_{\underline{\boldsymbol{\mu}}_m^{k-1}}[\Delta_{jm}(\{j\}, \Omega_m^{-f})|\mathcal{I}] - F_{jm}^{e}(\nu_{jm}^{e}) \geq 0\}, \tag{55}$$

$$\underline{O}_{jm}^{k} := \mathbb{1}\{\mathbb{E}_{\overline{\boldsymbol{\mu}}_m^{k-1}}[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})|\mathcal{I}] - F_{jm}^{e}(\nu_{jm}^{e}) \geq 0\}, \tag{56}$$

$$\overline{G}_{j}^{k} := \mathbb{1}\Big\{\sum_{m} \mathbb{E}\big[\overline{O}_{jm}^{k}\big[\mathbb{E}_{\underline{\boldsymbol{\mu}}_m^{k-1}}\big[\Delta_{jm}(\{j\}, \Omega_m^{-f})\big|\mathcal{I}\big] - F_{jm}^{e}(\nu_{jm}^{e})\big|\mathcal{I}\big] - F_{j}^{g}(\nu_{j}^{g}) \geq 0\Big\}, \tag{57}$$

$$\underline{G}_{j}^{k} := \mathbb{1}\Big\{\sum_{m} \mathbb{E}\big[\underline{O}_{jm}^{k}\big[\mathbb{E}_{\overline{\boldsymbol{\mu}}_m^{k-1}}\big[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})\big|\mathcal{I}\big] - F_{jm}^{e}(\nu_{jm}^{e})\big|\mathcal{I}\big] - F_{j}^{g}(\nu_{j}^{g}) \geq 0\Big\}, \tag{58}$$

To lighten notation in the proof, I define denote the probability that each of $\underline{O}_{jm}^{k}$, $\overline{O}_{jm}^{k}$, $\underline{G}_{j}^{k}$, $\overline{G}_{j}^{k}$, are equal to 1, respectively, as,

$$\underline{\mu}_{jm}^{cond}(\overline{\boldsymbol{\mu}}_m^{k-1}) := \Gamma_{jm}\big(\mathbb{E}_{\overline{\boldsymbol{\mu}}_m^{k-1}}[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})|\mathcal{I}]\big), \tag{59}$$

$$\overline{\mu}_{jm}^{cond}(\underline{\boldsymbol{\mu}}_m^{k-1}) := \Gamma_{jm}\big(\mathbb{E}_{\underline{\boldsymbol{\mu}}_m^{k-1}}[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})|\mathcal{I}]\big), \tag{60}$$

$$\underline{\mu}_{j}^{port}(\overline{\boldsymbol{\mu}}^{k-1}) := \Lambda_j\Big(\sum_{m\in\mathcal{M}} \mathbb{E}\big[\underline{O}_{jm}^{k}[\mathbb{E}_{\overline{\boldsymbol{\mu}}_m^{k-1}}\big[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})\big|\mathcal{I}\big] - F_{jm}^{e}(\nu_{jm}^{e})\big]\big|\mathcal{I}\big]\Big), \tag{61}$$

$$\overline{\mu}_{j}^{port}(\underline{\boldsymbol{\mu}}^{k-1}) := \Lambda_j\Big(\sum_{m\in\mathcal{M}} \mathbb{E}\big[\overline{O}_{jm}^{k}[\mathbb{E}_{\underline{\boldsymbol{\mu}}_m^{k-1}}\big[\Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f})\big|\mathcal{I}\big] - F_{jm}^{e}(\nu_{jm}^{e})\big]\big|\mathcal{I}\big]\Big), \tag{62}$$

where $\boldsymbol{\mu}^k := (\boldsymbol{\mu}_m^k)_{m\in\mathcal{M}}$. Note that $\{\underline{O}_{jm}^{k}\}_{j\in\mathcal{A}}$, $\{\overline{O}_{jm}^{k}\}_{j\in\mathcal{A}}$, $\{\underline{G}_{j}^{k}\}_{j\in\mathcal{A}}$, and $\{\overline{G}_{j}^{k}\}_{j\in\mathcal{A}}$ are sequences of independent random variables conditional on $\mathcal{I}$, and offering and portfolio choice bounds are conditionally independent, e.g., $\overline{O}_{jm}^{k}$ is independent of $\overline{G}_{j}^{k}$. **Proof.** I first prove that for any product j and any market m, $\{\overline{\mu}_{jm,k}\}_{k=0}^{\infty}$ is a decreasing sequence and $\{\underline{\mu}_{jm,k}\}_{k=0}^{\infty}$ is an increasing sequence. I prove this by induction. By construction $\underline{\mu}_{jm,1} > 0 = \underline{\mu}_{jm,0}$ and $\overline{\mu}_{jm,1} < 1 = \overline{\mu}_{jm,0}$. It follows that the hypothesis is true for $k = 1$. Assume that the hypothesis is true up to index $K > 1$. I show that (i) $\underline{\mu}_{jm,K+1} \geq \underline{\mu}_{jm,K}$ and (ii) $\overline{\mu}_{jm,K+1} \leq \overline{\mu}_{jm,K}$. To prove (i), note that by definition,

$$\underline{\mu}_{jm,K+1} = \underline{\mu}_{jm}^{cond}(\overline{\boldsymbol{\mu}}_m^K)\underline{\mu}_{j}^{port}(\overline{\boldsymbol{\mu}}^K).$$

By the inductive hypothesis, $\overline{\mu}_{jm,K} \leq \overline{\mu}_{jm,K-1}$ for each j and m. By definition of first-order stochastic dominance, and due to independence of the bounds across products within markets, this implies that $\overline{\boldsymbol{\mu}}_m^K \leq_{FOSD} \overline{\boldsymbol{\mu}}_m^{K-1}$ for all m. Moreover, due to Assumption 4 and

Theorem 3, both $\underline{\mu}_{jm}^{cond}$ and $\underline{\mu}_{j}^{port}$ are decreasing (in the FOSD sense). Therefore,

$$\underline{\mu}_{jm,K+1} = \underline{\mu}_{jm}^{cond}(\overline{\boldsymbol{\mu}}_m^K)\underline{\mu}_j^{port}(\overline{\boldsymbol{\mu}}^K) \geq \underline{\mu}_{jm}^{cond}(\overline{\boldsymbol{\mu}}_m^{K-1})\underline{\mu}_j^{port}(\overline{\boldsymbol{\mu}}^{K-1}) = \underline{\mu}_{jm,K}.$$

I have therefore proven that for each product j and each market m, the sequence of lower bound probabilities is increasing. Again, due to independence, this also implies that $\underline{\boldsymbol{\mu}}_m^{K+1} \geq_{FOSD} \underline{\boldsymbol{\mu}}_m^K$ for each K and each market m. I analogously prove (ii). First, write $\underline{\mu}_{jm,K} \geq \underline{\mu}_{jm,K-1}$ for each j and in each m. By definition of first-order stochastic dominance, and due to independence, this implies that $\underline{\boldsymbol{\mu}}_m^K \geq_{FOSD} \underline{\boldsymbol{\mu}}_m^{K-1}$. It follows from Assumption 4 and Theorem 3 that,

$$\overline{\mu}_{jm,K+1} = \overline{\mu}_{jm}^{cond}(\underline{\boldsymbol{\mu}}_m^K)\overline{\mu}_j^{port}(\underline{\boldsymbol{\mu}}^K) \leq \overline{\mu}_{jm}^{cond}(\underline{\boldsymbol{\mu}}_m^{K-1})\overline{\mu}_j^{port}(\underline{\boldsymbol{\mu}}^{K-1}) = \overline{\mu}_{jm,K}.$$

Thus, $\{\overline{\mu}_{jm,k}\}_{k=0}^{\infty}$ is decreasing and $\overline{\boldsymbol{\mu}}_m^K \leq_{FOSD} \overline{\boldsymbol{\mu}}_m^{K-1}$ due to independence. Next, I prove that for all k, necessarily,

$$\underline{\boldsymbol{\mu}}_m^k \leq_{FOSD} \boldsymbol{\mu}_m^* \leq_{FOSD} \overline{\boldsymbol{\mu}}_m^k \tag{63}$$

where asterisks denote any equilibrium distribution of product offering decisions under any arbitrary information set $\mathcal{I}$, for each $m \in \mathcal{M}$. I also prove this by induction. Inequalities (63) clearly hold for $k = 1$. Assume they hold for arbitrary $K \in \mathbb{N}$. I now prove that it must hold for $K+1$. To do so, take any arbitrary subset $\mathcal{W} \subseteq \mathcal{A}$. Note that $\mathcal{W} = \mathcal{W}^1 \cup \mathcal{W}^2 \cup ... \cup \mathcal{W}^F$, a partition across firms. Then,

$$\mathbb{P}\Big(\bigcap_{j\in\mathcal{W}}\{\overline{V}_{jm}^{K+1} = 1\}|\mathcal{I}\Big) = \prod_{f=1}^{F}\prod_{j\in\mathcal{W}^f}\mathbb{P}(\overline{V}_{jm}^{K+1} = 1|\mathcal{I}) \tag{64}$$

$$= \prod_{f=1}^{F}\prod_{j\in\mathcal{W}^f}\overline{\mu}_{jm}^{cond}(\underline{\boldsymbol{\mu}}_m^K)\overline{\mu}_j^{port}(\underline{\boldsymbol{\mu}}^K) \geq \prod_{f=1}^{F}\prod_{j\in\mathcal{W}^f}\overline{\mu}_{jm}^{cond}(\underline{\boldsymbol{\mu}}_m^*)\overline{\mu}_j^{port}(\underline{\boldsymbol{\mu}}^*) \tag{65}$$

$$= \prod_{f=1}^{F}\prod_{j\in\mathcal{W}^f}\mathbb{P}(\overline{O}_{jm}^* = 1|\mathcal{I})\mathbb{P}(\overline{G}_j^* = 1|\mathcal{I}) = \prod_{f=1}^{F}\prod_{j\in\mathcal{W}^f}\mathbb{P}(\overline{V}_{jm}^* = 1|\mathcal{I}) = \prod_{f=1}^{F}\prod_{j\in\mathcal{W}^f}\mathbb{E}(\overline{V}_{jm}^*|\mathcal{I})$$

$$= \mathbb{E}\Big[\prod_{f=1}^{F}\prod_{j\in\mathcal{W}^f}\overline{V}_{jm}^*|\mathcal{I}\Big] \geq \mathbb{E}\Big[\prod_{f=1}^{F}\prod_{j\in\mathcal{W}^f}V_{jm}^*|\mathcal{I}\Big] = \mathbb{P}\Big(\bigcap_{j\in\mathcal{W}}\{V_{jm}^* = 1\}|\mathcal{I}\Big), \tag{66}$$

where the inequalities in lines (64) and (66) follow from the (conditional) independence of the upper bounds and (65) follows from the inductive hypothesis and Theorem 3. This

proves $\overline{\boldsymbol{\mu}}^{K+1} \geq_{FOSD} \boldsymbol{\mu}^*$. I analogously prove the other direction. Indeed,

$$\mathbb{P}\Big(\bigcap_{j\in\mathcal{W}} \{\underline{V}_{jm}^{K+1} = 1\}|\mathcal{I}\Big) = \prod_{f=1}^{F} \prod_{j\in\mathcal{W}^f} \mathbb{P}\big(\underline{V}_{jm}^{K+1} = 1|\mathcal{I}\big) \tag{67}$$

$$= \prod_{f=1}^{F} \prod_{j\in\mathcal{W}^f} \underline{\mu}_{jm}^{cond}(\overline{\boldsymbol{\mu}}_m^K)\underline{\mu}_j^{port}(\overline{\boldsymbol{\mu}}^K) \leq \prod_{f=1}^{F} \prod_{j\in\mathcal{W}^f} \underline{\mu}_{jm}^{cond}(\overline{\boldsymbol{\mu}}_m^*)\underline{\mu}_j^{port}(\overline{\boldsymbol{\mu}}^*) \tag{68}$$

$$= \prod_{f=1}^{F} \prod_{j\in\mathcal{W}^f} \mathbb{P}(\underline{O}_{jm}^* = 1|\mathcal{I})\mathbb{P}(\underline{G}_j^* = 1|\mathcal{I}) = \prod_{f=1}^{F} \prod_{j\in\mathcal{W}^f} \mathbb{P}(\underline{V}_{jm}^* = 1|\mathcal{I}) = \prod_{f=1}^{F} \prod_{j\in\mathcal{W}^f} \mathbb{E}(\underline{V}_{jm}^*|\mathcal{I})$$

$$= \mathbb{E}\Big[\prod_{f=1}^{F} \prod_{j\in\mathcal{W}^f} \underline{V}_{jm}^*|\mathcal{I}\Big] \leq \mathbb{E}\Big[\prod_{f=1}^{F} \prod_{j\in\mathcal{W}^f} V_{jm}^*|\mathcal{I}\Big] = \mathbb{P}\Big(\bigcap_{j\in\mathcal{W}} \{V_{jm}^* = 1\}|\mathcal{I}\Big), \tag{69}$$

where (67) and (69) follow from (conditional) independence of the lower bounds and (68) follows from the inductive hypothesis and Theorem 3.

This concludes the proof. ■

Notice that this proof only relies on Assumption 1, Assumption 4, and conditional independence of fixed cost shocks. It does not rely on the log-normality assumption.

C.2 Implementation

To implement Algorithm 1 in practice, I need to compute expectations $\mathbb{E}_{\boldsymbol{\mu}_m}$, which are with respect to the distribution of rival firms' offerings decisions given $\boldsymbol{\mu}_m$ and over demand and marginal cost shocks. Thus, I first take 100 draws from the joint distribution of $(\boldsymbol{\xi}, \boldsymbol{\omega})$ across 9766 product markets, yielding a 9766×100 matrix of demand shocks and a similar matrix of marginal cost shocks. Second, I draw a 9766×100 matrix U^{cond} and a 1530×100 matrix U^{port} of uniformly distributed draws on (0, 1). I hold all draws fixed throughout all iterations.

The U^{port} matrix yields thresholds for the portfolio decisions given by probabilities $\overline{\mu}_{j,k}^{port}$ and $\underline{\mu}_{j,k}^{port}$. For instance, if product j has a draw smaller than $\overline{\mu}_{j,k}^{port}$, then that product is introduced in the firm's portfolio under that draw at iteration k at the upper bound. Otherwise, at iteration k and under such a draw, product j is not introduced in the firm's global product portfolio at the upper bound. The U^{cond} matrix yields thresholds for the offering choices given by probabilities $\overline{\mu}_{jm,k}^{cond}$ and $\underline{\mu}_{jm,k}^{cond}$. For instance, if for product j in market m the U^{cond} draw is less than the $\overline{\mu}_{jm,k}^{cond}$ probability at iteration k, then that draw corresponds to such a product being offered in such a market at iteration k (conditional on being in the firm's portfolio). Otherwise, that product is not offered in market m under iteration k and that draw. For a product to be offered in a market at a given iteration k, it

must satisfy both the offering and portfolio thresholds.

It follows that under each iteration k, I need to compute 9766×100 marginal values across all product-market pairs. Then, I average them across the 100 draws to obtain simulated values of the expected marginal values at the product-market level. This averaging yields the expectations that enter the fixed cost CDFs in equations (59)-(62). To compute the marginal values under any draw, I use the Morrow and Skerlos (2011) contraction mapping for the pricing equilibria with 500 income and normal draws for each market.

I run Algorithm 1 for 6 iterations under each parameter vector $(\theta_e, \sigma_e, \theta_g, \sigma_g)$ and each policy counterfactual. I find that after $k > 6$, the additional gains in informativeness are small in my application.

D Estimation Implementation

D.1 Demand and Marginal Cost Estimation

Demand estimation follows Petrin (2002). The assumption that $(\boldsymbol{\xi}, \boldsymbol{\omega})$ are realized after firms make offerings choices implies that $\mathbb{E}[(\xi_{jm}, \omega_{jm})|\mathcal{I}] = 0$. Under this assumption, standard Berry et al. (1995) or Gandhi and Houde (2019) instruments are valid. I use size, horsepower, and horsepower/weight to build Gandhi and Houde (2019) differentiation instruments. These instruments are constructed by using the characteristics of products that are "close" in the characteristics space. I use the PyBLP Python package to construct the differentiation instruments (Conlon and Gortmaker 2020). For each of the three characteristics, I build two instruments. Let $x_{jm\ell}$ denote the value of any such characteristic, indexed by ℓ. I construct, for each $\ell \in \{\text{size, horsepower, horsepower/weight}\}$:

$$z_{jm\ell}^{other} = \sum_{k \in \Omega_m^f \setminus \{j\}} \mathbb{1}\{|d_{jkm\ell}| < SD_\ell\}, \qquad z_{jm\ell}^{rival} = \sum_{k \in \Omega_m^{-f}} \mathbb{1}\{|d_{jkm\ell}| < SD_\ell\}$$

where $d_{jkm\ell} = x_{jm\ell} - x_{km\ell}$ and SD_ℓ denotes the standard deviation of all such pairwise differences computed across all markets. This yields 6 differentiation instruments.

To improve the precision of my estimates, I include 4 additional moments. First, I include (i) the log of distance to the brand's HQ country, which enters my marginal cost specification, and (ii) the average price of the same product in other markets / the average price of products of the same parent company in other markets (when a given product is only sold in one market). The latter set of instruments resembles those in Hausman (1996), which are valid provided demand and marginal cost shocks $(\boldsymbol{\xi}_m, \boldsymbol{\omega}_m)$ are uncorrelated across markets.

Second, I use micro-moments, similarly to Petrin (2002). I use micro-data from the 2019 MRI-Simmons Crosstab report to help pin down the heterogeneity in preferences for prices within countries. I match the probabilities that a consumer in the United States is within a given income group conditional on purchasing a vehicle in a given price range. More specifically, I match the following two moments:

$$\text{(a)} \qquad \mathbb{E}[income_i > \$100,000 | price_{jm} > \$50,000, m = \text{United States}]$$

$$\text{(b)} \qquad \mathbb{E}[income_i \in [\$60,000, \$100,000] | price_{jm} > \$50,000, m = \text{United States}].$$

I jointly estimate demand and marginal costs using Python's PyBLP package. To integrate over the distribution of income in each market/country, I take 20,000 simulation draws in each market m from a log-normal distribution with location and scale,

$$\mu_m = \log(GDP_per_capita_PPP_m) - \sigma_m^2/2, \qquad \sigma_m = \sqrt{2}\Phi^{-1}\left(\frac{Gini_m + 1}{2}\right),$$

where Φ is the standard normal CDF. This parametrization ensures that income is drawn from a log-normal distribution with mean and Gini coefficients equal to the observed values. I obtain the Gini coefficient and PPP GDP per capita in each market m from the World Bank. The main results from the estimation are in the main text, Table 1. Below, I report the matched micro-moments:

Table D.1: Micro-Moments

Moment	**Observed**	**Estimated**	**Difference**
(a)	0.631	0.612	0.0188
(b)	0.212	0.245	-0.0329

D.1.1 Distribution of ($\boldsymbol{\xi}$, $\boldsymbol{\omega}$)

I fit a bivariate normal distribution for the distribution of demand and marginal cost shocks, which I assume firms know at the time of making product portfolio and product offerings decisions. I estimate a variance of 2.36 for ξ, a variance of 0.024 for ω, and a covariance of 0.081.

D.2 Fixed Cost Estimation - Instruments

To estimate fixed costs, I implement the moment inequalities derived in Sections B.1 and B.2. As described in the main text, I first construct instruments using an approach

similar to two-stage least squares. I then use these instruments to build unconditional moment inequalities. I use the empirical counterparts of these moment inequalities to estimate $(\theta_e, \sigma_e, \theta_g, \sigma_g)$. The steps are described in the main text, Section 4.

To construct the instruments, I use the Stage 3 model for variable profits, together with the fitted distribution of $(\boldsymbol{\xi}, \boldsymbol{\omega})$ to compute the ingredients necessary to implement the inequalities in Theorem 1. I simulate $S = 200$ draws $\{\xi_{jm}^s, \omega_{jm}^s\}_{s=1, j\in\mathcal{A}, m\in\mathcal{M}}^{200}$ from the fitted bivariate normal distribution and construct, for each product j and each market m,

$$\widehat{\Delta}_{jm}(\{j\}, \Omega_m^{-f}) = \frac{1}{S}\sum_{s=1}^{S} \Delta_{jm}(\{j\}, \Omega_m^{-f}; \xi_{jm}^s, \omega_{jm}^s) \tag{70}$$

$$\widehat{\Delta}_{jm}(\mathcal{G}^f, \Omega_m^{-f}) = \frac{1}{S}\sum_{s=1}^{S} \Delta_{jm}(\mathcal{G}^f, \Omega_m^{-f}; \xi_{jm}^s, \omega_{jm}^s) \tag{71}$$

$$\widehat{\Delta}_{jm}(\mathcal{A}^f, \Omega_m^{-f}) = \frac{1}{S}\sum_{s=1}^{S} \Delta_{jm}(\mathcal{A}^f, \Omega_m^{-f}; \xi_{jm}^s, \omega_{jm}^s). \tag{72}$$

To obtain an exogenous ($\mathcal{I}$- measurable) predictor of the above marginal values, I project the log of the predicted values in (70)-(72) on the following objects: (i) the log of $\tilde{\delta}_{jm} = \beta_m + \beta_{b(j)} + \beta^x \boldsymbol{X}_{jm}$, the non-price mean utility of product j in market m net of unobserved heterogeneity ξ_{jm}, (ii) $\tilde{c}_{jm} = \log(c_{jm}) - \omega_{jm}$, the (log) marginal cost of supplying market m with product j net of the unobserved heterogeneity ω_{jm}, and (iii) $\log(\tilde{\delta}_{jm}) \times \tilde{c}_{jm}$. To carry out the projection, I estimate 4 PPML specifications. I first project (70)-(71), conditional on $j \in \mathcal{G}^f$, using, $\hat{y}_{jm} = \exp(\kappa^1 M_m + \kappa_m^2 \tilde{\delta}_{jm} + \kappa_m^3 \tilde{c}_{jm} + \kappa_m^4 \tilde{c}_{jm} \times \tilde{\delta}_{jm}) + \varepsilon_{jm}$. I then use the same specification to project (70) and (72), using the full sample of potential products. As described in the main text, I then use predicted values from these regressions to construct the unconditional moments that can be used for estimation.

To implement the Stage 2 inequalities that partially identify (θ_e, σ_e), I use tercile bin indicators for the predicted values in the PPML regressions that project (70)-(71), which I denote by $\hat{\bar{x}}_{jm}^g$ and $\hat{\underline{x}}_{jm}^h$, respectively. I also use the squares of such (positive) predicted values interacted with an indicator function denoting whether the log of the predicted value (in billions of USD) is greater than -2, which selects more profitable products. This helps to provide bounds on the scale parameter σ_e, as argued in Section 4.

To implement the Stage 1 inequalities that partially identify (θ_g, σ_g), I use quintile bin indicators (5 bins) for the *sum* across markets of the predicted values arising from PPML regressions that project (70) and (72) using the full sample of potential products – $\hat{\bar{x}}_{jm}^a$ and $\hat{\underline{x}}_{jm}^l$, respectively –, as well as squares of such sums of predicted values, interacted with an indicator function for whether the log of the sum (in billions of USD) is greater than 0.

D.2.1 Approximation Points Used for Convex and Convex CDF Bounds

To implement the moment inequality procedure, I use observation-specific upper and lower bounds of the fixed-cost CDFs. I use the convex and concave families of functions in equations (22) and (29), respectively. The approximation points I use for the Stage 2 upper and lower bound inequalities are $\hat{\overline{x}}^g_{jm}$ and $\hat{\underline{x}}^h_{jm}$, respectively. I use $\sum_m \hat{\overline{x}}^a_{jm}$ and $\sum_m \hat{\underline{x}}^l_{jm}$, respectively, for the Stage 1 upper and lower bound inequalities.

D.3 Fixed Cost Estimation - Empirical Analogues of Moments in Theorem 1

I construct positive-valued instruments and interact them with the conditional moments in Theorem 1 to construct empirical unconditional moment inequalities,

$$\frac{1}{J_g M}\sum_{j\in\mathcal{G},m\in\mathcal{M}}[\overline{\Gamma}_{jm}(\Delta_{jm}(\{j\},\Omega_m^{-f});\theta_e,\sigma_e)-O_{jm}]\mathbb{1}\{\hat{\overline{x}}^g_{jm}\in Q_\tau(\hat{\overline{x}}^g_{jm})\]\geq 0, \tag{73}$$

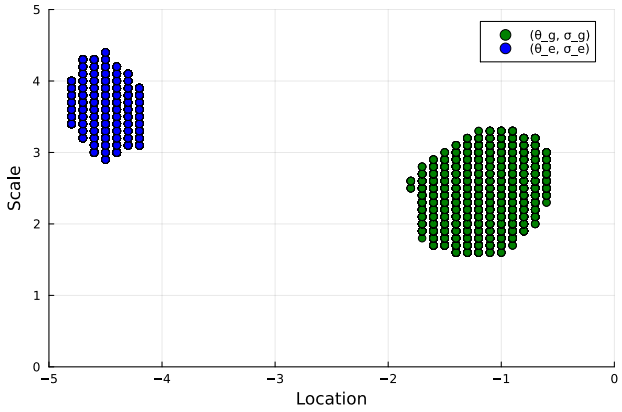
$$\frac{1}{J_g M}\sum_{j\in\mathcal{G},m\in\mathcal{M}}[\underline{\Gamma}_{jm}(\Delta_{jm}(\mathcal{G}^f,\Omega_m^{-f});\theta_e,\sigma_e)-O_{jm}]\mathbb{1}\{\hat{\underline{x}}^h_{jm}\in Q_\tau(\hat{\underline{x}}^h_{jm})\]\leq 0, \tag{74}$$

$$\begin{aligned}
&\frac{1}{J}\sum_{j\in\mathcal{A}}\Bigg[\overline{\Lambda}_j\Bigg(\sum_{m\in\mathcal{M}}\Gamma_{jm}\big(\Delta_{jm}(\{j\},\Omega_m^{-f})\big)\big[\Delta_{jm}(\{j\},\Omega_m^{-f})\\
&\quad-\mathbb{E}\big[F^e_{jm}(\nu^e_{jm})|\mathcal{I},F^e_{jm}(\nu^e_{jm})\leq\Delta_{jm}(\{j\},\Omega_m^{-f})\big]\big];\theta_g,\sigma_g\Bigg)-G_j\Bigg]\\
&\quad\times\mathbb{1}\Bigg\{\sum_{m\in\mathcal{M}}\hat{\overline{x}}^a_{jm}\in Q_\tau\Bigg(\sum_{m\in\mathcal{M}}\hat{\overline{x}}^a_{jm}\Bigg)\Bigg\}\geq 0,
\end{aligned} \tag{75}$$

$$\begin{aligned}
&\frac{1}{J}\sum_{j\in\mathcal{A}}\Bigg[\underline{\Lambda}_j\Bigg(\sum_{m\in\mathcal{M}}\Gamma_{jm}\big(\hat{\underline{x}}^l_{jm}\big)\big[\Delta_{jm}(\mathcal{A}^f,\Omega_m^{-f})\\
&\quad-\mathbb{E}\big[F^e_{jm}(\nu^e_{jm})|F^e_{jm}(\nu^e_{jm})\leq\hat{\underline{x}}^l_{jm},\mathcal{I}\big]\big];\theta_g,\sigma_g\Bigg)-G_j\Bigg]\\
&\quad\times\mathbb{1}\Bigg\{\sum_{m\in\mathcal{M}}\hat{\underline{x}}^l_{jm}\in Q_\tau\Bigg(\sum_{m\in\mathcal{M}}\hat{\underline{x}}^l_{jm}\Bigg)\Bigg\}\leq 0,
\end{aligned} \tag{76}$$

where the expectation of the truncated fixed costs in inequalities (75)-(76) are computed using the QuadGK Julia package (Gaussian quadrature). J_g denotes the number of products in the sample that are offered in at least one market.

Figure E.1: Baseline Fixed Cost Estimates - Grid Search



E Fixed Cost Estimation - Grid Search and Robustness

Table E.1: Stages 1 and 2 Parameter Confidence Set Limits

	95% Confidence Set Limits	
	More IV Bins	**Fewer IV Bins**
Stage 2: Market Entry Fixed Cost		
θ_e (Location)	[-4.9, -4.1]	[-4.8, -4.2]
σ_e (Scale)	[2.9, 4.4]	[2.8, 4.3]
Stage 1: Product Fixed Cost		
θ_g (Location)	[-1.9, -0.5]	[-1.9, -0.6]
σ_g (Scale)	[1.5, 3.6]	[1.5, 3.7]
Observations – Stage 2	3240	3240
Observations – Stage 1	739	739

Note: Confidence sets computed using Andrews and Soares (2010). First, I implement a grid search to compute a 97.5% confidence set for parameters (θ_e, σ_e) using the Stage 2 moment inequalities. Then, I use the Stage 1 moment inequalities to compute a 97.5% confidence set for (θ_g, σ_g), evaluating the moments at the accepted values of (θ_e, σ_e). The Bonferroni correction yields a 95% confidence set for all four parameters $(\theta_e, \sigma_e, \theta_g, \sigma_g)$. Marginal values are in billions of U.S. dollars.

Figure E.1 plots the grid search for the baseline fixed cost estimates.

Next, I report the estimation results using more instrument bins for both Stage 2 and Stage 1 inequalities. I construct 10 bins for the Stage 2 inequalities and 8 bins for the Stage 1 inequalities using the regressions from Section D.2. Recall that the baseline specification used 3 bins for Stage 2 and 5 bins for Stage 1. I keep the remaining polynomial-based instruments as they are in the main specification. I do not find that including more instrument bins leads to model rejection. Table E.1 shows that the confidence sets do not change very significantly

under more instrument bins. And, if anything, the limits of the confidence sets become wider. This is due to increased covariance across more instruments.

I also report the confidence sets under fewer instrument bins: 2 bins for both Stage 2 and Stage 1. Again, all remaining instruments are left unchanged. Table E.1 shows that the confidence sets do not change much relative to the main specification. This increases my confidence in the results and shows that they are robust to the number of instrument bins.

F Moment Inequalities Extension - Economies of Scope

The model described in the main text assumes that portfolio fixed costs and market entry fixed costs are independent of the firms' chosen portfolio and offerings bundles. In this subsection, I show how this assumption can be relaxed. I allow market entry fixed costs to exhibit a version of economies of scope. Similar arguments can also be used to allow for economies of scope in global portfolio fixed costs.

Suppose that the total market entry fixed costs paid by firm f if it offers products Ω_m^f are,

$$F_m^{e,f} = \theta_0 \mathbb{1}\{|\Omega_m^f| \geq 1\} + \sum_{j \in \Omega_m^f} F_{jm}^e.$$

That is, firms pay a constant amount θ_0 to enter the market, and then market entry fixed costs with a specification as in the main text, given by Assumption 2. Throughout, I assume that a natural lower bound on θ_0 is 0 to be consistent with the fact that fixed costs are positive under the log-normal assumption. Then, I apply similar arguments to those in the main text to derive a lower and an upper bound inequality. As before, I start with revealed-preference equality,

$$\begin{aligned}(\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}})[\mathbb{1}\mathbb{E}[\Delta_{jm}(\Omega_m^f, \Omega_m^{-f})|\mathcal{I}, \{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}, \mathcal{G}^f] \\ - \theta_0 \mathbb{1}\{\Omega_m^f = \{j\}\} - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0 - \mathbb{1}_{\Omega_m^f}] = 0,\end{aligned}$$

where I have now allowed the fixed cost to be higher if j is the first product to be introduced into market m. Under submodularity and $\theta_0 \geq 0$, one can derive upper bound inequality,

$$\begin{aligned}(\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f \setminus \{j\}})[\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\{j\}, \Omega_m^{-f})|\mathcal{I}, \{\nu_{jm}^e\}_{j \in \mathcal{G}^f, m}, \mathcal{G}^f] \\ - F_{jm}^e(\nu_{jm}^e; \theta_e, \sigma_e) \geq 0\} - \mathbb{1}_{\Omega_m^f}] \geq 0,\end{aligned}$$

which results from maximizing the marginal value and minimizing the marginal fixed cost.

For the upper bound, I then follow the same arguments as in the main text to derive upper bound inequality,

$$\mathbb{E}[\overline{\Gamma}_{jm}(\Delta_{jm}(\{j\},\Omega_m^{-f});\theta_e,\sigma_e) - O_{jm}|\mathcal{I},\mathcal{G}^f] \geq 0. \tag{77}$$

Inequality (77) does not provide any bound on θ_0. For the lower bound, I proceed in a similar fashion. Minimizing Δ_{jm} and maximizing marginal fixed costs, I derive inequality,

$$\begin{aligned}(\mathbb{1}_{\Omega_m^f} + \mathbb{1}_{\Omega_m^f\setminus\{j\}})[\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\mathcal{G}^f,\Omega_m^{-f})|\mathcal{I},\{\nu_{jm}^e\}_{j\in\mathcal{G}^f,m},\mathcal{G}^f] - \theta_0 \\ - F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e) \geq 0\} - \mathbb{1}_{\Omega_m^f}] \leq 0.\end{aligned}$$

Following similar steps as in section B.1, I obtain conditional moment inequality,

$$\mathbb{E}[\underline{\Gamma}_{jm}(\Delta_{jm}(\mathcal{G}^f,\Omega_m^{-f}) - \theta_0;\theta_e,\sigma_e) - O_{jm}|\mathcal{I}] \leq 0. \tag{78}$$

Inequality (78) provides a lower bound on θ_0. Therefore, the inequalities discussed so far only identify a lower bound and not an upper bound of the new parameter of interest θ_0. In the following subsection, I discuss additional inequalities that partially identify θ_0.

F.1 Additional Inequalities

To derive additional inequalities, I condition on either a single product or no product being offered in market m. That is, I now start with,

$$\begin{aligned}&(\mathbb{1}_{\{j\}_m} + \mathbb{1}_{\emptyset_m})\times \\ &[\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\{j\},\Omega_m^{-f})|\mathcal{I},\{\nu_{jm}^e\}_{j\in\mathcal{G}^f,m},\mathcal{G}^f] - \theta_0 - F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e) \geq 0\} - \mathbb{1}_{\{j\}_m}] = 0.\end{aligned}$$

To deal with the issue of selection, I first re-write the above inequality as,

$$\begin{aligned}&(\mathbb{1}_{\{j\}_m} + \mathbb{1}_{\emptyset_m})\times \\ &\quad\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\{j\},\Omega_m^{-f})|\mathcal{I},\{\nu_{jm}^e\}_{j\in\mathcal{G}^f,m},\mathcal{G}^f] - \theta_0 - F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e) \geq 0\} = \mathbb{1}_{\{j\}_m}.\end{aligned} \tag{79}$$

Since $(\mathbb{1}_{\{j\}_m} + \mathbb{1}_{\emptyset_m}) \leq 1$ and due to Assumption 1, I obtain,

$$\mathbb{1}\{\mathbb{E}[\Delta_{jm}(\{j\},\Omega_m^{-f})|\mathcal{I},\mathcal{G}^f] - \theta_0 - F_{jm}^e(\nu_{jm}^e;\theta_e,\sigma_e) \geq 0\} - \mathbb{1}_{\{j\}_m} \geq 0. \tag{80}$$

Inequality (80) says that if only a single product j is offered in market m, then the change in profits from doing so must necessarily be weakly positive. Letting O_{jm}^1 denote the event

that only product j is sold in market m by firm f, I follow steps similar to those in the main text to derive inequality,

$$\mathbb{E}[\overline{\Gamma}_{jm}(\Delta_{jm}(\{j\}, \Omega_m^{-f}) - \theta_0; \theta_e, \sigma_e) - O_{jm}^1 | \mathcal{I}] \geq 0.$$

To implement this moment inequality in practice, the empirical counterpart of $\mathbb{E}[O_{jm}^1|\mathcal{I}]$ averages across firm-markets a dummy equal to 1 if a given firm introduces a singleton bundle in a given market. An upper bound on θ_0 is identified because a value of θ_0 that is too large would imply too few observations of singleton product entry across markets.

In practice, I do not observe sufficient instances of singleton product entry to obtain informative bounds on θ_0.

G Simulating the Method

In this section, I use simulation to:

1. Test the behavior of the moment inequalities proposed in Section B under varying data-generating processes (DGPs),

2. Test the performance of the inference methods used in this paper, based on Andrews and Soares (2010), when a single realization of the product entry game is observed.

To simulate data, I need to solve the model fully. Thus, I simulate a solvable version of my product entry model, which I describe in Section G.1.

G.1 Fully Solvable Version of Global Product Entry Game

The solvable model features N symmetric J-product firms competing in M markets. I set $J = 3$ so that firms have 3 potential products that they can introduce in their global product portfolios and across markets. Both firms and products are symmetric in their profit shifters, but markets are allowed to be heterogeneous.

Profits for firm f from selling N_m^f products in market m take the form, $\Pi_m^f(N_m^f, N_m^{-f}) = A_m \frac{N_m^f}{1+(N_m^f)^{\kappa_o}(N_m^{-f})^{\kappa_r}}$, where $\kappa_o \in (0,1)$ regulates substitutability across the firm's own products, and $\kappa_r \in (0,1)$ regulates substitution across rival firms' products. A_m is an exogenous market-level profit shifter. Firms draw log-normal fixed costs, as in the main text. I abuse notation and index all firms' potential products as 1, 2, or 3 (even though they are different products).

Timing: The timing is exactly as in the model in the main text. In the first stage, each firm f draws private fixed portfolio cost shocks $\{\nu_j^g\}$ for each of their 3 potential products. Upon observing this private information, firms choose which products to introduce in their portfolio. Next, they draw private fixed market entry shocks $\{\nu_{jm}^e\}_{j\in\mathcal{G}^f,m}$ for each product and choose how many of the products in their portfolio to offer in each market.

Best response: Firm f's best response in Stage 2 depends on portfolio size N_p^f. Let $FE_{mn}^{(k)}$ denote the k^{th} order statistic among the n potential entry cost draws in market m, and let $FD^{(k)}$ denote the k^{th} order statistic among 3 product development cost draws.

If $N_p^f = 1$, the firm sells in market m iff $\mathbb{E}[\Pi_m^f(1, N_m^{-f})|\mathcal{I}] - FE_{m1}^{(1)} \geq 0$.

If $N_p^f = 2$,

$$\begin{cases} \text{2 products iff } \mathbb{E}[\Pi_m^f(2, N_m^{-f})|\mathcal{I}] - \mathbb{E}[\Pi_m^f(1, N_m^{-f})|\mathcal{I}] - FE_{m2}^{(2)} \geq 0; \\ \text{1 product iff } \mathbb{E}[\Pi_m^f(1, N_m^{-f})|\mathcal{I}] - FE_{m2}^{(1)} \geq 0, \\ \quad \text{and } \mathbb{E}[\Pi_m^f(2, N_m^{-f})|\mathcal{I}] - \mathbb{E}[\Pi_m^f(1, N_m^{-f})|\mathcal{I}] - FE_{m2}^{(2)} < 0. \end{cases}$$

If $N_p^f = 3$,

$$\begin{cases} \text{3 products iff } \mathbb{E}[\Pi_m^f(3, N_m^{-f})|\mathcal{I}] - \mathbb{E}[\Pi_m^f(2, N_m^{-f})|\mathcal{I}] - FE_{m3}^{(3)} \geq 0; \\ \text{2 products iff } \mathbb{E}[\Pi_m^f(3, N_m^{-f})|\mathcal{I}] - \mathbb{E}[\Pi_m^f(2, N_m^{-f})|\mathcal{I}] - FE_{m3}^{(3)} < 0, \\ \quad \text{and } \mathbb{E}[\Pi_m^f(2, N_m^{-f})|\mathcal{I}] - \mathbb{E}[\Pi_m^f(1, N_m^{-f})|\mathcal{I}] - FE_{m3}^{(2)} \geq 0; \\ \text{1 product iff } \mathbb{E}[\Pi_m^f(2, N_m^{-f})|\mathcal{I}] - \mathbb{E}[\Pi_m^f(1, N_m^{-f})|\mathcal{I}] - FE_{m3}^{(2)} < 0, \\ \quad \text{and } \mathbb{E}[\Pi_m^f(1, N_m^{-f})|\mathcal{I}] - FE_{m3}^{(1)} \geq 0. \end{cases}$$

In Stage 1, define

$$\mathcal{V}_f(N_p^f) = \mathbb{E}\Big[\sum_m \max_{N_m^f \leq N_p^f} \mathbb{E}[\overline{\Pi}_m^f(N_m^f, N_m^{-f})|\mathcal{I}]\Big|\mathcal{I}\Big].$$

The firm develops:

$$\begin{cases} 3 \text{ products iff } \mathcal{V}_f(3) - \mathcal{V}_f(2) - FD^{(3)} \geq 0; \\ 2 \text{ products iff } \mathcal{V}_f(3) - \mathcal{V}_f(2) - FD^{(3)} < 0, \\ \quad \text{and } \mathcal{V}_f(2) - \mathcal{V}_f(1) - FD^{(2)} \geq 0; \\ 1 \text{ product iff } \mathcal{V}_f(2) - \mathcal{V}_f(1) - FD^{(2)} < 0, \\ \quad \text{and } \mathcal{V}_f(1) - FD^{(1)} \geq 0. \end{cases}$$

The equilibrium solution is summarized by three market thresholds:

$$t^e_{1m} = \mathbb{E}[\Pi^f_m(1,N^{-f}_m)|\mathcal{I}], \quad t^e_{2m} = \mathbb{E}[\Pi^f_m(2,N^{-f}_m)|\mathcal{I}] - \mathbb{E}[\Pi^f_m(1,N^{-f}_m)|\mathcal{I}],$$
$$t^e_{3m} = \mathbb{E}[\Pi^f_m(3,N^{-f}_m)|\mathcal{I}] - \mathbb{E}[\Pi^f_m(2,N^{-f}_m)|\mathcal{I}] \tag{81}$$

and three product development thresholds,

$$t^g_1 = \mathcal{V}_f(1), \quad t^g_2 = \mathcal{V}_f(2) - \mathcal{V}_f(1), \quad t^g_3 = \mathcal{V}_f(3) - \mathcal{V}_f(2), \tag{82}$$

where $\mathcal{V}_f$ integrates over rival and own product choices, determined by $(t^e_{1m}, t^e_{2m}, t^e_{3m})$ in equilibrium. Since firms are symmetric, note that the thresholds are not firm-specific.

Computing equilibrium profits given threshold strategies: The probability that any firm offers n products, for $n \in \{1,2,3\}$, given the portfolio and market entry strategies, are given by,

$$\begin{aligned} p_1(\boldsymbol{t}^e, \boldsymbol{t}^g) &= \mathbb{P}(FD^{(3)} \leq t^g_3)[\mathbb{P}(FE^{(1)}_3 \leq t^e_1) - \mathbb{P}(FE^{(2)}_3 \leq t^e_2)] \\ &\quad + [\mathbb{P}(FD^{(2)} \leq t^g_2) - \mathbb{P}(FD^{(3)} \leq t^g_3)][\mathbb{P}(FE^{(2)}_2 \leq t^e_2) - \mathbb{P}(FE^{(1)}_2 \leq t^e_1)] \\ &\quad + [\mathbb{P}(FD^{(1)} \leq t^g_1) - \mathbb{P}(FD^{(2)} \leq t^g_2)][\mathbb{P}(FE^{(1)}_1 \leq t^e_1)], \qquad (83) \\ p_2(\boldsymbol{t}^e, \boldsymbol{t}^g) &= \mathbb{P}(FD^{(3)} \leq t^g_3)[\mathbb{P}(FE^{(2)}_3 \leq t^e_2) - \mathbb{P}(FE^{(3)}_3 \leq t^e_3)] \\ &\quad + [\mathbb{P}(FD^{(2)} \leq t^g_2) - \mathbb{P}(FD^{(3)} \leq t^g_3)]\mathbb{P}[FE^{(2)}_2 \leq t^e_2], \qquad (84) \\ p_3(\boldsymbol{t}^e, \boldsymbol{t}^g) &= \mathbb{P}(FD^{(3)} \leq t^g_3)\mathbb{P}(FE^{(3)}_3 \leq t^e_3). \qquad (85) \end{aligned}$$

To compute expected profits in the first stage given any N^f_p, first realize that $\mathcal{V}_f$ takes the

following form,

$$\mathcal{V}_f(N_p^f) = \sum_m \sum_{i=1}^{N_p^f} \mathbb{P}(N_m^f = i|N_p^f)[\mathbb{E}[\Pi_m^f(i, N_m^{-f})|\mathcal{I}] - i\mathbb{E}[F_{jm}^e|F_{jm}^e \leq \mathbb{E}[\Pi_m^f(i, N_m^{-f})|\mathcal{I}]]].$$

Equations (83)-(85) characterize $\mathbb{P}(N_m^f = i|N_p^f)$ as a function of $(\boldsymbol{t}^e, \boldsymbol{t}^p)$. Given expected variable profits, $\mathbb{E}[\Pi_m^f(i, N_m^{-f})|\mathcal{I}]$, the term $\mathbb{E}[F_{jm}^e|F_{jm}^e \leq \mathbb{E}[\Pi_m^f(i, N_m^{-f})|\mathcal{I}]]$ is simple to compute numerically using Gaussian quadrature. I use the QuadGK package in Julia to compute this expectation.

To compute $\mathbb{E}[\Pi_m^f(i, N_m^{-f})|\mathcal{I}]$ for each i given $(\boldsymbol{t}^e, \boldsymbol{t}^g)$, I use the DSP package in Julia to perform a convolution which gives the probability distribution of the number of rival product offerings given the number of firms N and $p_n(\boldsymbol{t}^e, \boldsymbol{t}^g)$ for $n \in \{1, 2, 3\}$.

Solving for $(\boldsymbol{t}^e, \boldsymbol{t}^g)$: Given $(\boldsymbol{t}^e, \boldsymbol{t}^g)$, I showed how to evaluate $\mathcal{V}_f$ and $\mathbb{E}[\Pi_m^f(i, N_m^{-f})|\mathcal{I}]$ for any $i \in \{1, 2, 3\}$. Thus, I solve for $(\boldsymbol{t}^e, \boldsymbol{t}^g)$ by solving the non-linear system of equations (81)-(82). To do so, I use the NLsolve Julia package.

G.2 Behavior of Moment Inequalities Across DGPs

In this section, I study how the informativeness of the moment inequalities varies with parameters κ_o and κ_r and with the number of firms N. The informativeness of the inequalities depends on the loss from bounding the marginal value of introducing a product by evaluating it at extreme bundles and on the loss from applying Jensen's inequality to average out firms' expectational errors.

Implementation: I simulate $S = 500$ different realizations for each of $T = 12$ different "types" of global product and market entry game described in the previous section. In each game (s, t), there are N firms competing in 12 markets. I hold fixed the market-level profit shifters and set them at $A_m^{(s,t)} = 0.2mt$ for $m \in \{1, 2, ..., 12\}$, $t \in \{1, 2, ..., 12\}$, and $s \in \{1, 2, ..., 500\}$. Different values of t generate variation in profitability across different game types. In this subsection, I perform valid asymptotics as $ST \to \infty$. With $ST = 6000$, simulation noise is small, so I report the identified sets.

True parameters: I set the true parameters to be $(\theta_g, \sigma_g) = (3, 1)$ and $(\theta_e, \sigma_e) = (1, 1)$.

Instruments: For the Stage 2 inequalities, I condition on the realized set of portfolio decisions and construct instruments following the PPML procedure described in the main text. More precisely, I project the minimal (maximal) marginal values at the product-market level – evaluated at all (no) products in the portfolio being introduced in the market and at the realized set of rival offerings decisions – using specification $\gamma_0 \exp(\gamma_m^t A_m^{(s,t)})$. I then compute the predicted values of both regressions, sort them, and define 4 percentile

categories and bins associated with these categories. This gives 8 instruments in total – 4 for the upper bound inequality and 4 for the lower bound inequality.

For the Stage 1 inequalities, I implement an equivalent procedure as with the Stage 2 inequalities, but without conditioning on the realized product portfolios. I project the realized maximal and minimal (where now the minimal marginal value does not condition on the observed portfolio, so is evaluated at all 3 products being introduced in each market) marginal values on $\gamma_0 \exp(\gamma_m^t A_m^{(s,t)})$. I then compute the predicted values of both regressions and sum them across markets. Then, I sort the sums and define 4 percentile categories for both the lower and the upper bound and define 4 bins associated with these categories, yielding 8 instrument bins - 4 for the lower bound and 4 for the upper bound.

G.2.1 Baseline: $N = 10$, $\kappa_o = 0.1$, $\kappa_r = 0.1$

In this baseline case, I set the number of firms to $N = 10$ and set $\kappa_o = 0.1$ and $\kappa_r = 0.1$.

Figure G.1: Identified Sets

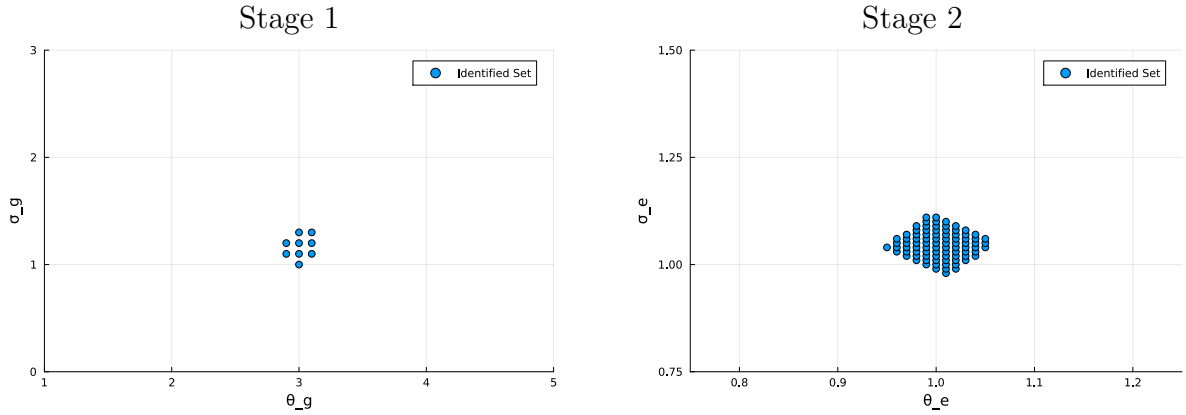


Figure G.1 shows the identified sets for (θ_g, σ_g) and (θ_e, σ_e), which I obtained via grid search. As seen in this figure, the identified set is quite informative both for the Stage 1 and the Stage 2 fixed cost parameters (note: the scale is not the same for the Stage 1 and the Stage 2 parameters).[5]

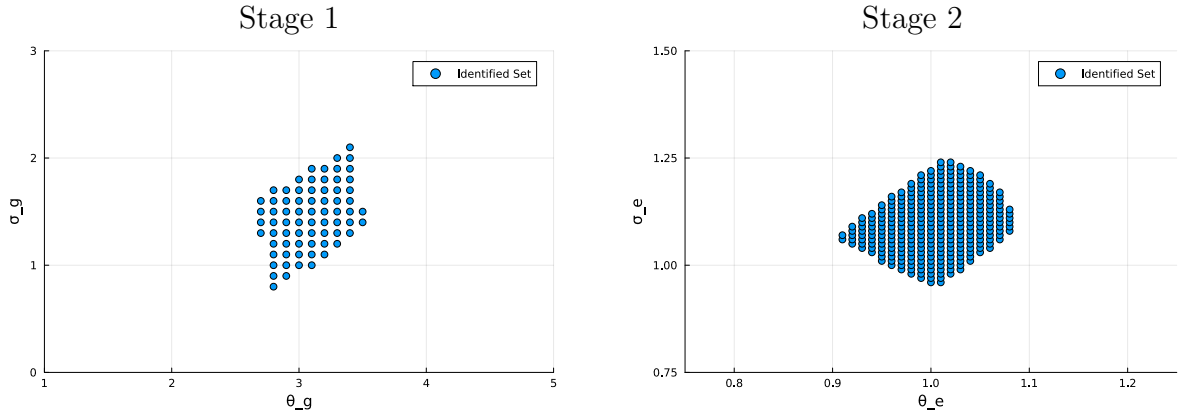
G.2.2 High Substitutability Within the Firm: $N = 10$, $\kappa_o = 0.25$, $\kappa_r = 0.1$

I now study a deviation from the baseline case in which the parameter determining substitution within the firm is larger.

Figure G.2 shows the results. Compared to the baseline cases, high substitution within the firm reduces the informativeness of the fixed cost parameter bounds, both for the Stage

[5]In Additional Materials C I also report the lower contours for the moment inequalities.

Figure G.2: Identified Sets



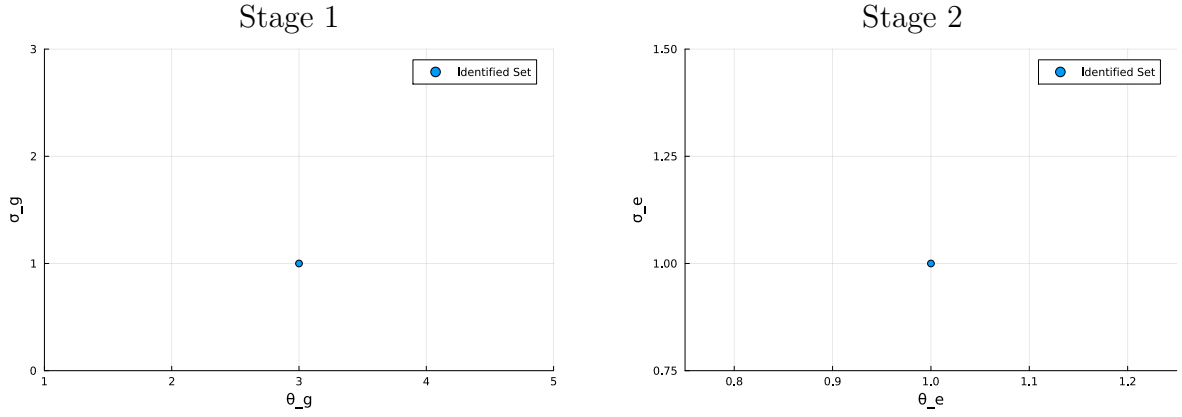
1 and for the Stage 2 fixed cost parameters. This is expected given that the loss from bounding the marginal value of introducing a product at extreme bundles (all other products vs. no other products offered) is greater when such products are highly substitutable. That is, greater cannibalization within the firm makes the moment inequalities less informative about the true parameters.

G.2.3 Low Substitutability within the Firm: $N = 10$, $\kappa_o = 0.01$, $\kappa_r = 0.1$

Figure G.3 show what happens when there is very low substitutability across the firms' own products. Interestingly, the tightness of the inequalities increases greatly so that only the true fixed cost parameters both in Stage 1 and Stage 2 are accepted (given my grid).

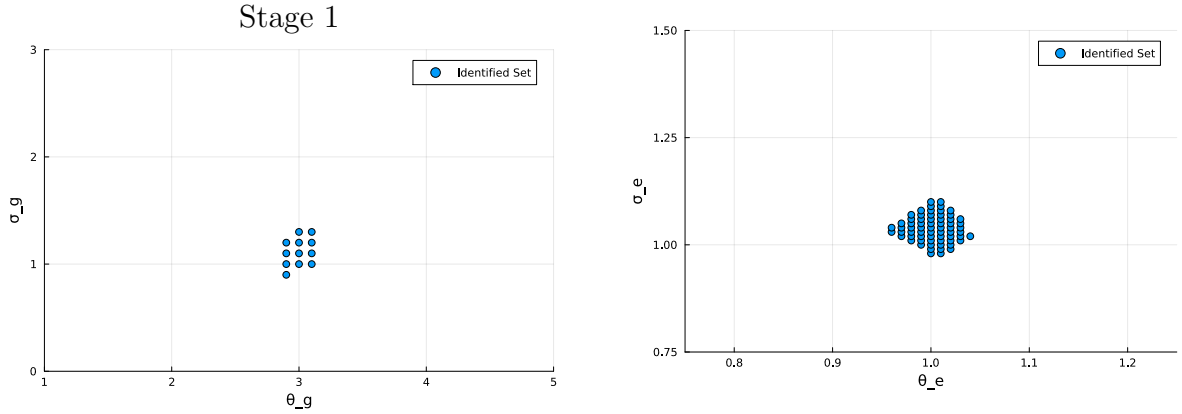
This result is of special significance for anyone wanting to use this estimation approach in a setting where firms are single-product (rather than multi-product). Indeed, Figure G.3 shows that in the absence of cannibalization within the firm, my estimation procedure is highly informative.

Figure G.3: Identified Sets



G.2.4 High Substitutability Across Firms: $N = 10$, $\kappa_o = 0.1$, $\kappa_r = 0.25$

Figure G.4: Identified Sets



Interestingly, high substitution across firms does not seem to change the tightness of the inequalities much relative to the baseline case. While it seems to slightly reduce the informativeness of the Stage 1 inequalities, it does not seem to lead to reduced informativeness in Stage 2. If anything, in this case, the Stage 2 inequalities yield tighter bounds on θ_e. The impact of higher κ_r on the informativeness of the inequalities is therefore not as quantitatively important as the impact of κ_o. This makes sense because, by virtue of the unobservability of rivals' fixed cost shocks, I showed that one can use the realized set of rival offerings decisions to construct the moment inequalities. Greater substitution across firms reduces the rate of product introduction in equilibrium, but firms expect this, and the informativeness of the moment inequalities is not significantly affected.

G.2.5 Low Substitutability Across Firms $N = 10$, $\kappa_o = 0.1$, $\kappa_r = 0.01$

For completeness, Figure G.5 reports the identified sets for the case in which there is low substitution across rival firms' products. As expected, there are no substantial differences in the informativeness of the moment inequalities relative to the baseline case for the same reasons as in Section G.2.4.

G.2.6 Large Number of Firms: $N = 20$, $\kappa_o = 0.1$, $\kappa_r = 0.1$

Figure G.6 illustrates the effect of having more firms participating in the global entry game. The identified set becomes slightly smaller when there are more firms. This is due to two effects. First, when there are more firms, each firm makes smaller expectational errors. Second, when there are more firms, more products may be offered on average, which reduces the loss from bounding marginal values with extreme bundles due to submodularity.

G.2.7 Small Number of Firms: $N = 2$, $\kappa_o = 0.1$, $\kappa_r = 0.1$

Figure G.7 illustrate the effect of having only 2 firms competing globally. In this case, the inequalities become less informative for the same reasons as argued in Section G.2.6.

G.2.8 Main Takeaways

While the simulations abstract away from heterogeneity across products and across firms, they are still useful for understanding some of the key properties of the moment inequalities. I have established that much of the informativeness of the moment inequalities relies on the extent to which products *within* the firm are substitutable. High degrees of substitutability within the firm render the moment inequalities less informative, while lower degrees of substitutability make them tighter. Moreover, the number of firms (relative to the number of products per firm) matters. With smaller firms relative to the overall market, expectational errors are smaller. If the set of potential products of a firm is small relative to the set of products that are offered in the market, the loss from bounding a product's marginal value with extreme bundles within the firm is smaller. Finally, I showed that all else equal,

Figure G.5: Identified Sets

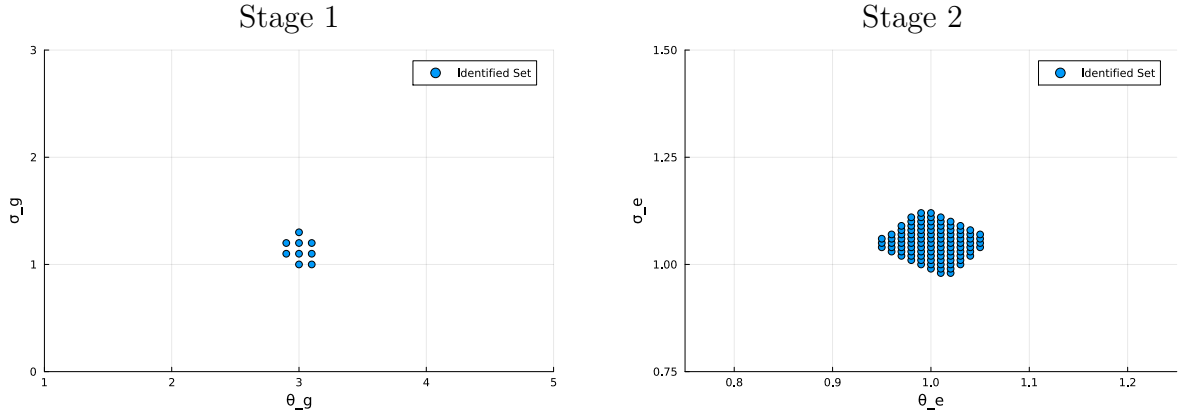


Figure G.6: Identified Sets

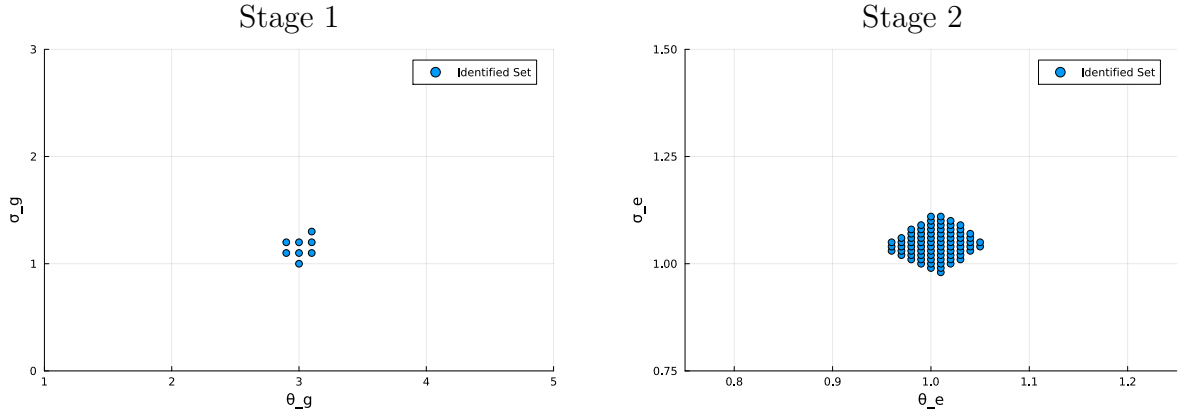
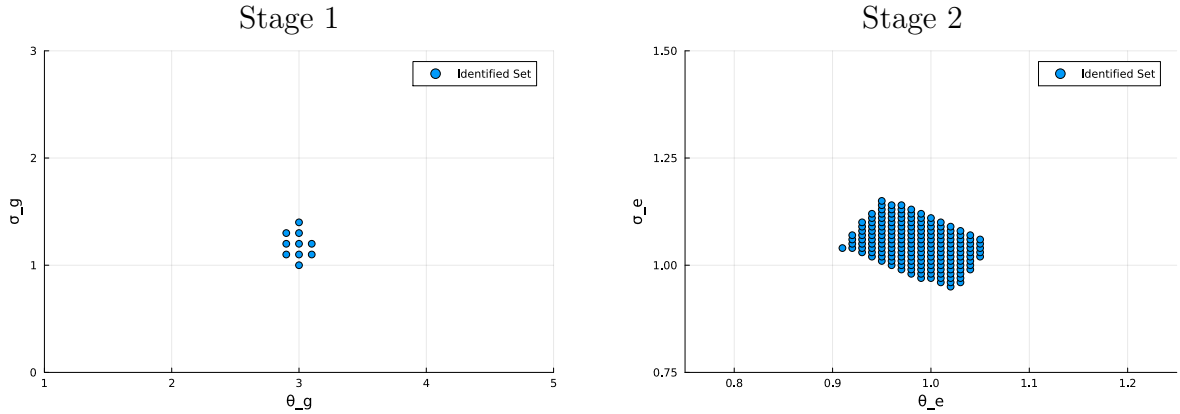


Figure G.7: Identified Sets



substitution across firms does not have much of an effect on the tightness of the moment inequalities.

G.3 Inference Under a Single Realization of Global Product Entry Game

In this section, I use the fully solvable version of the model to assess the properties of Andrews and Soares (2010) confidence sets in my setting. I simulate $S = 100$ realizations for each of $T = 12$ different "types" of global product and market entry games, just as in Section G.2. In each game (s, t), there are N firms competing in 12 markets. I hold fixed the market-level profit shifters and set them at $A_m^{(s,t)} = 0.2mt$ for $m \in \{1, 2, ..., 12\}$, $t \in \{1, 2, ..., 12\}$ and $s \in \{1, 2, ..., 100\}$. Different values of t generate variation in profitability across different game types.

True parameters: I set the true parameters to be $(\theta_g, \sigma_g) = (3, 1)$ and $(\theta_e, \sigma_e) = (1, 1)$.

Instruments: I construct the instruments as described in Section G.2. However, I run the PPML regressions at the (s, t)-level to mimic the actual implementation in the main text, where I only observe a single realization of the cross-section. Thus, the variation used to construct confidence sets in this section is across product-market pairs.

For each (s, t) pair, I construct confidence sets for parameters (θ_e, σ_e) using the procedure in Andrews and Soares (2010). For each of the $S \times T$ 95% confidence sets, I record: (i) whether the true values $(1, 1)$ are included in the confidence set (coverage), (ii) the length of the confidence set along the θ_e dimension (holding $\sigma_e = 1$ at the truth), (iii) the length of the confidence set along the σ_e dimension (holding $\theta_e = 1$ at the truth).

Panel A in Table G.1 reports the average coverage of the confidence set across all $S \times T$ realizations of the global product entry game. I do this both under a robust variance-covariance matrix and a clustered (market-level) variance-covariance matrix. Undercoverage

can occur, particularly when the number of firms is relatively small. As the number of firms increases, the coverage of the Andrews and Soares (2010) confidence sets tends to increase. Note that clustering does not necessarily yield higher coverage, though this is the case when the number of firms is large ($N = 75$).

Table G.1: Confidence Set Properties for (θ_e, σ_e)

SE Type	$N = 5$	$N = 25$	$N = 50$	$N = 75$
Panel A: Coverage of True Parameters (%)				
Robust	92.9	93.8	94.1	93.7
Clustered	92.4	94.1	94.1	95.0
Panel B: Median Length of Confidence Set				
Robust	(0.8, 2.8)	(0.4, 0.9)	(0.2, 0.6)	(0.2, 0.5)
Clustered	(0.9, 4.4)	(0.4, 1.6)	(0.3, 1.1)	(0.3, 0.9)

Note: Panel A reports the average coverage across all $S \times T$ simulations of the confidence sets for (θ_e, σ_e). Panel B reports the median length of these confidence sets along each dimension. The first coordinate corresponds to the θ_e dimension conditional on σ_e being at its true value, and the second to the σ_e dimension conditional on θ_e being at its true value. Confidence sets are computed at the (s, t) level, as in the empirical application.

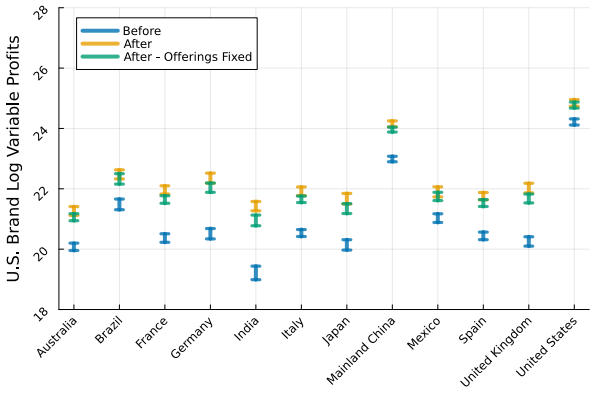
Panel B in Table G.1 reports the median length of the confidence set along each of the dimensions of (θ_e, σ_e), conditional on the confidence set not being empty. The confidence sets are smaller whenever there are more firms and without clustering. In the empirical implementation, I report confidence sets using a robust variance-covariance matrix because (i) Table G.1 (Panel A) shows that clustering does not necessarily improve coverage and undercoverage is not severe, and (ii) Table G.1 (Panel B) shows that clustering leads to a larger confidence set.

A caveat in this exercise is that the simulations require symmetry across firms and products within a (s, t) pair. While at Stage 2, even in the simulations, firms are not fully symmetric due to variation in how many products they have in their global product portfolio, greater symmetry relative to reality should worsen correlation across expectational errors since the realization of the market structure affects firms in (almost) the exact same way. However, greater symmetry also means that there are no "large" firms in the sample. Large firms with a lot of market power can deteriorate the asymptotic properties of the estimator.

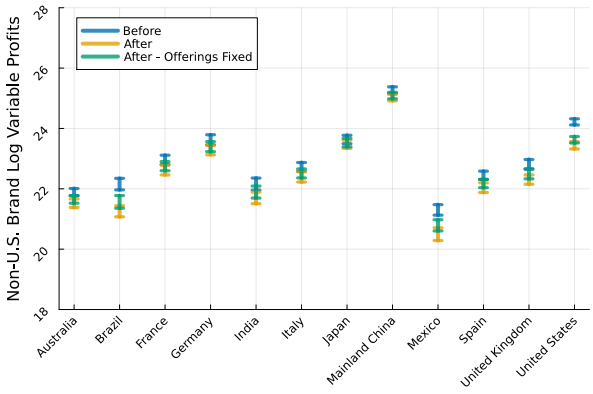
H Counterfactual Exercises - Variable Profit Plots

Figure H.1: Variable Profits under U.S. Policy Interventions

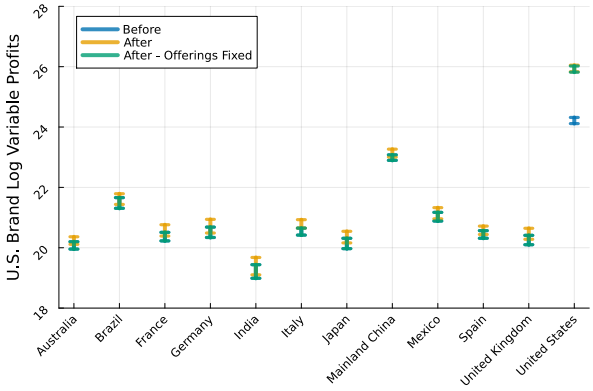
Panel A: U.S. Products – Marginal Cost Subsidy



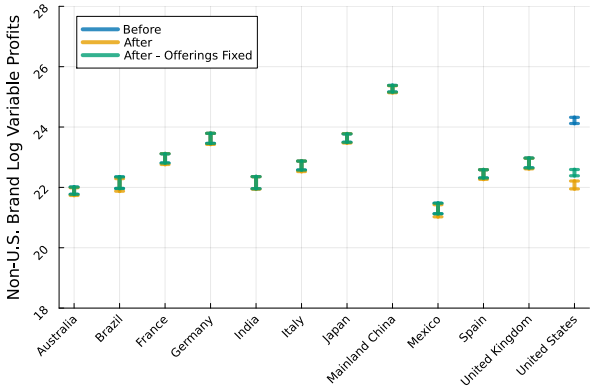
Panel B: Non-U.S. Products – Marginal Cost Subsidy



Panel C: U.S. Products – Consumer Subsidy



Panel D: Non-U.S. Products – Consumer Subsidy



Note: Panels A–B display bounds on the expected (log) total variable profits of U.S. and non-U.S. brands across countries before (blue) and after (orange) a 20% reduction in U.S. brands' marginal costs. Panels C–D show analogous bounds for a 50% consumer subsidy on U.S.-branded products in the United States. Expectations are taken with respect to bounds on the probability distribution of firms' offerings and demand and marginal cost shocks. The green intervals use the bounds on the distribution of product offerings in each market before the policy is implemented.

References

Donald W. K. Andrews and Gustavo Soares. Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection. *Econometrica*, 78(1):119–157, 2010.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile Prices in Market Equilibrium. *Econometrica*, 63(4):841–890, 1995.

Christopher Conlon and Jeff Gortmaker. Best Practices for Differentiated Products Demand Estimation with PyBLP. *The RAND Journal of Economics*, 51(4):1108–1161, 2020.

Amit Gandhi and Jean-François Houde. Measuring Substitution Patterns in Differentiated-Products Industries. *Mimeo*, 2019.

Paul L E Grieco, Charles Murry, and Ali Yurukoglu. The Evolution of Market Power in the U.S. Automobile Industry. *The Quarterly Journal of Economics*, page qjad047, 09 2023.

Jerry Hausman. Valuation of New Goods under Perfect and Imperfect Competition. In *The Economics of New Goods*, pages 207–248. National Bureau of Economic Research, Inc, 1996.

W. Ross Morrow and Steven J. Skerlos. Fixed-Point Approaches to Computing Bertrand-Nash Equilibrium Prices Under Mixed-Logit Demand. *Operations Research*, 59(2):328–345, April 2011.

Amil Petrin. Quantifying the Benefits of New Products: The Case of the Minivan. *Journal of Political Economy*, 110(4):705–729, 2002.

Charly Porcher, Eduardo Morales, and Thomas Fujiwara. Measuring Information Frictions in Migration Decisions: A Revealed-Preference Approach. *Mimeo*, May 2024.

Moshe Shaked. *Stochastic Orders*. Springer Science & Business Media, 2007.