Market Entry and Plant Location in Multi-Product Firms

Juanma Castro-Vincenzi
University of Chicago

Eugenia Menaguale
Princeton University

Eduardo Morales
Princeton University

Alejandro Sabal
Yale University*

March 25, 2026

Abstract

Multinational firms offer different bundles of products across countries due to variation in consumer tastes and production costs, and business-stealing externalities across products. Using car-industry data, we estimate a quantitative model in which manufacturers decide where to produce and sell their products. Firms face both complementary and substitutable decisions across locations and products due to business stealing and minimum-cost sourcing. We provide a new solution algorithm for binary integer programming problems featuring both pairwise complementarities and substitutabilities. We study the effect of tariffs and consumer and producer subsidies on firms' optimal production location and product offerings decisions. Our quantitative exercises show that U.S. tariffs are ineffective at inducing production relocation to the United States, and that U.S. consumer subsidies and producer subsidies affect the global structure of production and offerings.

JEL Classifications: F13, F23, L21, L23, L62.

Keywords: multinationals, multiproduct firms, discrete choice methods, moment inequalities, entry.

*We thank Pol Antràs, Elhanan Helpman, Michal Kolesár, Ellen McGrattan, Marc Melitz, and Ariel Pakes for insightful conversations, and seminar participants at the AEA Meetings, Chicago, Harvard, Minnesota Macro, NYU, Princeton, and SED, for helpful comments. This research was supported by the International Economics Section (IES) at Princeton University and Fundación Ramón Areces. All errors are our own. Emails: Juanma Castro-Vincenzi (castrovincenzi@uchicago.edu), Eugenia Menaguale (eugenia.menaguale@princeton.edu), Eduardo Morales (ecmorale@princeton.edu), Alejandro Sabal (alejandro.sabal@yale.edu).

1 Introduction

Multinational, multiproduct firms play an outsized role in shaping the global economy and account for a large share of global manufacturing output and employment. Although the literature on international trade often treats these two facets—their multiproduct nature and multinational presence—in isolation, the choices of the group of markets to serve, the set of products to sell and to export in each of those markets, and the location of production of these products are inherently interdependent.

In this paper, we study how policies influence the global production decisions of multiproduct firms, highlighting the interactions between country-specific demand, supply, and trade policies and shocks. As firms optimize their production and market entry decisions jointly, policy changes in one country can trigger non-trivial adjustments in product availability, sourcing, pricing, and trade patterns across the global economy.

To address these questions, we integrate a canonical export platform model, akin to Tintelnot (2017), with a standard multiproduct firm framework inspired by Hottman et al. (2016). In our model, a firm with a given product portfolio makes two key sets of binary decisions: (1) selecting the markets in which to sell each product and (2) choosing the locations where production takes place. On the demand side, firms face a nested constant elasticity of substitution (CES) utility structure, which allows the elasticity of demand across a firm's own products to differ from the elasticity of substitution across varieties supplied by different firms. On the production side, firms can operate as multinationals, producing in multiple countries and potentially serving multiple markets from a single production location, though they must incur trade costs. To either produce or sell a product in a country, firms must incur fixed costs.

In this setup, firms face a number of tradeoffs. First, the model features cannibalization across the products of a firm in a given market, i.e. when deciding to introduce a product in a given market a firm internalizes that a part of the market share of that product is going to come at the expense of its other products. Second, the model features minimum-cost sourcing. This assumption implies that production plants for the same product cannibalize each other. It also implies that producing a product and selling it are necessarily complementary decisions. These decisions richly interact with trade costs, as there is a proximity-concentration tradeoff by which product entry and plant location decisions are co-determined; the benefits of serving a market depend on where the product is produced, and the value of a plant on the markets it serves. Due to the existence of fixed costs for both sets of discrete decisions, the problem inevitably delves into a binary integer programming problem.

The solution to this problem requires addressing the computational challenge of solving a model in which firms make discrete, interdependent choices, with some pairs of decisions being complements, and others substitutes. To circumvent these difficulties, we develop a novel algorithm for high-dimensional single-agent discrete choice problems that operates on the space of all sublattices of the original decision lattice. By defining a set-increasing mapping on these sublattices, we leverage Tarski's fixed point theorem to obtain bounds on the optimal solution even when the objective

function is neither globally supermodular nor submodular. The key condition required to apply the algorithm is that the sign of the interdependence between any two decisions is known and is invariant to other choices. In an extension, we relax this assumption and prove that when the payoff function is additively separable, the property of constant-signed interdependence is required only for each of the additive components of the payoff function.

The mapping works by operating on the set of sublattices of the J-dimensional set of possible decision vectors. For each sublattice and each decision or coordinate, we define a "best-case" scenario by choosing the upper-bound value for coordinates that are complementary and the lower-bound value for those that are substitutes. Conversely, the "worst-case" scenario uses the lower-bound for complements and the upper-bound for substitutes. Decisions with negative marginal benefits—even in their best case—are fixed at 0, while those with nonnegative marginal benefits—even in their worst case—are fixed at 1. Iterating this set-increasing mapping yields convergence, by Tarski's fixed point theorem, to an upper bound on the set of possible solution vectors.

We examine these margins of participation in the global car industry. This industry provides an ideal setting for studying these questions, as it comprises multinational, multiproduct firms producing tradable goods. To this end, we leverage a rich global dataset that links quantities from the production location of individual car models to their country of sale. We complement this dataset with data on car prices and characteristics at the car model level for a subset of countries that account for more than 80% of global car sales.

The global car industry exemplifies the defining features of multinational, multiproduct firms, with automakers managing extensive product portfolios and deciding over production and sales across markets. Brands offer an average of 3.56 models per country and segment, with some featuring as many as 23, and the typical model is sold in 12 markets, sometimes reaching up to 75. Production is similarly dispersed, with models manufactured in an average of 1.5 locations and roughly 44.5% of production exported when serving multiple markets. These patterns underscore the industry's rich structure, where firms must navigate trade-offs between market access, and production, fixed and trade costs.

To estimate the parameters of the demand equation, as well as the determinants of trade and marginal production costs, we use our merged dataset on prices and quantities. To estimate fixed costs, we introduce a a novel moment inequality approach, constructing moments that are monotonic in firms' binary choices. We use our solution algorithm to obtain computationally feasible upper and lower bounds for the firms' optimal choices given a candidate parameter vector, which we use to build our moment inequalities. Using these bounds for estimation instead of the model-implied optimal choices, which are generally infeasible to compute, we avoid using heuristic algorithms that approximate the solution of the problem and can introduce biases on estimated parameters.

With our estimated model, we examine how production subsidies, consumption subsidies, and tariffs favoring electric vehicles could shape the global automobile industry. For each policy, we assess the impact of the United States enacting each policy at different values. The goal of our

analysis is to shed light on the effectiveness of these policies in achieving their intended goals, rather than to judge whether there are normative reasons to implement these policies.

Our findings indicate that demand-side EV subsidies in the U.S. favor production in the United States, as well as origins that can serve the U.S. market competitively, such as Mexico. Demand-side subsidies also encourage EV market entry in the United States. Marginal cost subsidies that are large enough are effective at shifting production toward the United States, and lead to greater EV offerings globally. Both of these policies reduce the share of non-electric alternative being offered, not only in the U.S., but also in other markets. Tariffs, in contrast, have limited effects on global production structures due to the high fixed costs associated with plant relocation. These policy tools induce production relocation of some products – those that are very appealing to U.S. customers (yet are not produced in the United States at baseline), but lead to reductions in overall variety offered for consumers. Additionally, U.S. policies generate spillover effects on other countries by influencing firms' production location decisions, and consequently their sourcing and offering choices.

Our paper closely relates to different branches of the literature. First, it relates to papers that quantitatively study multinational production such as Ramondo and Rodríguez-Clare (2013), Tintelnot (2017), Arkolakis et al. (2018), Arkolakis et al. (2023), Castro-Vincenzi (2024) and Head et al. (2025)). In particular, our paper is related to those that allow for export platforms, whereby firms can operate plants that serve multiple markets, and increasing returns at the plant level. Second, this paper connects with the literature that studies multiproduct firms in international markets, such as Eckel and Neary (2010), Dhingra (2013), Mayer et al. (2014), Hottman et al. (2016), Head and Mayer (2019) and Arkolakis et al. (2021). Our paper merges these two literatures by studying the problem of jointly deciding in which locations to produce and in which markets to sell its goods for multinational, multiproduct firms in a quantitative framework.[1]

Third, our paper relates to the methodological literature concerned with solving models with multiple discrete choices. The methods developed in this literature have exclusively focused on problems in which the sign of the interdependencies is common across all the potential choices of the economic agent of interest, by devising algorithms that leverage the super- or sub-modularity of the agents' objective function (e.g. Jia 2008, Antràs et al. 2017, Arkolakis et al. 2023 or Alfaro-Ureña et al. 2024, Sabal 2025). Our contribution to the literature is to provide a solution method for solving single-agent static problems featuring simultaneous positive and negative interdependencies. We show that our method can be implemented under restrictions on the sign of the interdependence across pairs of binary choices.

Fourth, our paper contributes to the literature on estimating combinatorial discrete choice models. Previous approaches (e.g., Antràs et al. 2017, Alfaro-Ureña et al. 2024) have used algorithms that bound the solution of the model, together with heuristic algorithms to obtain a full solution, to implement estimation techniques that impose point-identification (e.g., MLE, GMM).

[1]It is worth noting that two papers, Bernard et al. (2018) and Helpman and Niswonger (2023), theoretically explore the multiproduct and multinational margins for large firms.

Our approach differs in that we use moment inequalities for estimation that exempt us from using heuristic methods. We build moment inequalities that are valid whenever the moments are monotone in firms' discrete choices. In contrast to previous moment inequality procedures in single-agent settings (e.g., Holmes 2011, Morales et al. 2019), we allow for unobserved heterogeneity in firms' payoff function. Our approach also differs from moment inequality procedures in strategic settings (e.g., Fan and Yang 2020, Sabal 2025) in that we use our bounding solution procedure to build moments that are monotone in the firms' entry decisions. Sabal (2025) provides a bounding solution procedure for entry games, but implementing it in estimation is computationally infeasible due to the pricing game, whereas it is feasible in our more tractable single-agent setting.

Closest to our paper is Head et al. (2025), who also study multinational production decisions in the automotive sector, focusing on multi-stage global supply chains in the production of battery-electric vehicles (BEVs). Their setting similarly features complementarities across stages of the supply chain and substitutabilities within each stage, generating a combinatorial problem akin to ours. Methodologically, they introduce a novel approach by recasting the problem as a mixed-integer linear programming (MILP) model, allowing them to leverage fast commercial solvers specifically designed for this class of problems. However, their method cannot be directly applied to our context, as the interactions among firms' product-entry decisions—specifically, cannibalization effects—non-linearly affect the variable profits of each product. They use their model to explore how policies such as subsidies, tariffs, and local production requirements can have complex and sometimes unintended impacts on firms' production-location decisions and overall battery production and EV adoption.

Finally, the paper belongs to the literature that studies multinational production in the automotive industry and the analysis of government policy effects in the sector. It builds on work examining multinational production decisions in the automotive industry (e.g., Coşar et al. 2016, Head and Mayer 2019, Castro-Vincenzi 2024, Head et al. 2025). While these prior studies analyze bilateral frictions, home market advantages, or production location choices, they do not fully integrate firms' endogenous decisions on both production and product entry. This paper also extends research on government policies in the automobile market (e.g., Goldberg 1995, Barwick et al. 2024, Allcott et al. 2024, Sabal 2025, Head et al. 2025) by studying their effect on the global production decisions of these multiproduct firms, although abstracting from strategic interactions across firms.

The rest of the paper is structured as follows. In Section 2, we present a quantitative model of multinational, multiproduct firms. Section 3 describes the solution algorithm. Section 4 introduces the data and explains the estimation strategy. In Section 5, we present the model estimates, and we discuss counterfactual results in Section 6. Section 8 concludes.

2 Model

In this section, we develop a quantitative model of plant location decisions for multiproduct firms. Brand $b \in \mathcal{B}$ competing in segment, s, with a given set of products, Ω^{bs}, decides in which

subset of potential countries, $\mathcal{N} = \{1, 2, ..., N\}$, to produce and sell each of its products. Each product, $m \in \Omega^{bs}$, is characterized by a set of demand and marginal cost shifters that potentially vary across destination and origin countries. We assume that brands are atomistic and do not internalize the effect of their decisions on segment-country expenditure.

We model brands' portfolio and market entry choices in two stages. In Stage 1, each brand b competing in a segment s draws for each product $m \in \Omega^{bs}$ a country in which it can produce and sell at zero fixed costs, a set of plant location fixed cost shocks $\{\nu^p_{mo}\}_{m\in\Omega^{b,s},o}$, and set of market entry fixed cost shocks $\{\nu^e_{mn}\}_{m\in\Omega^{b,s},n}$.[2] Upon observing the realizations of these unobservables (to the econometrician), brand b simultaneously chooses in which countries to produce and in which countries to sell each of its products. In Stage 2, brand b chooses prices and sourcing strategy for each product it chose to sell in each country.

2.1 Demand, Marginal Costs, and Prices

In each country $n \in \mathcal{N}$, the representative consumer has nested CES preferences.[3] Consumers in country n spend an exogenous amount E_{sn} in segment s, which we define by body type categories (e.g., SUV).[4] The consumption index of segment s in country n is given by,

$$C_{sn} = \left(\int_{b\in\Omega_{sn}} (\psi_{bsn} C_{bsn})^{\frac{\sigma^B - 1}{\sigma^B}} db \right)^{\frac{\sigma^B}{\sigma^B - 1}}, \tag{1}$$

where ψ_{bsn} is an exogenous brand-segment-country demand shifter and σ^B denotes the elasticity of substitution across brands within a segment. We parameterize ψ_{bsn} as,

$$\log(\psi_{bsn}) = \phi_b + \beta_1 \mathbb{1}\{h(b) = n\} + \beta_2 \log dist_{nh(b)} + u_{bsn},$$

where $h(b)$ denotes the headquarters country of brand b and u_{bsn} is an unobserved (to the econometrician) demand shock.

Brand b is unable to affect segment-country expenditure because it is atomistic. However, brand b internalizes the effect of its plant location and market entry decisions on its consumption index C_{bsn}, which is given by,

$$C_{bsn} = \left(\sum_{m=1}^{M_{bs}} I_{mn} (\psi_{mn} C_{mn})^{\frac{\sigma^M - 1}{\sigma^M}} \right)^{\frac{\sigma^M}{\sigma^M - 1}}, \tag{2}$$

[2] We denote collections of variables $\{X_m\}_{m\in\Omega^{b,s}}$ or $\{Y_n\}_{n\in\mathcal{N}}$ by $\{X_m\}_m$ and $\{Y_n\}_n$, respectively, for ease of notation.

[3] We recognize that our demand model does not capture rich substitution patterns, as in approaches that use mixed logit demand systems with random coefficients e.g., Berry et al. (1995), Grieco et al. (2023). Instead, we provide a micro-foundation that justifies why firms do not internalize their impact of market aggregates when making their choices (given our focus on solving a single firm's problem), and that allows us to sign theoretically the interdependence across any pair of choices.

[4] A possible micro-foundation is a Cobb-Douglas nest over automobiles and an outside good, followed by a Cobb-Douglas nest over a discrete set of possible segments (body types).

where ψ_{mn} is an exogenous product-country demand shifter and σ^M denotes the elasticity of substitution across models within a brand-segment. If the brand offers product m in country n, $I_{mn} = 1$; otherwise, $I_{mn} = 0$. We parametrize demand shifter, ψ_{mn} , as,

$$\log(\psi_{mn}) = \phi_m + \beta_3 \log(HP_m) \log(GDPpc_n) + \beta_4 \log(Weight_m) \log(GDPpc_n) + u_{mn}, \quad (3)$$

where $GDPpc_n$ denotes the GDP per capita of country n, HP_m denotes the horsepower of product m and u_{mn} denotes an unobserved (to the econometrician) demand shock.

The revenue share of brand b within segment s in country n is given by,

$$S_{bsn} = \frac{(P_{bsn}/\psi_{bsn})^{1-\sigma^B}}{P_{sn}^{1-\sigma^B}}, \quad (4)$$

where P_{sn} is the price index of segment s in country n and the brand-segment-country exact price index is given by,

$$P_{bsn} = \left(\sum_{m=1}^{M_{bs}} I_{mn} \left(\frac{P_{mn}}{\psi_{mn}} \right)^{1-\sigma^M} \right)^{\frac{1}{1-\sigma^M}}. \quad (5)$$

The price index P_{bsn} is a sufficient statistic capturing all endogenous factors that affect revenues and variable profits for brand b in segment-country sn.

The brand chooses which of its M_{bs} products to sell in each country, knowing that the demand equation for product m in country n is,

$$C_{mn} = A_{sn} (\psi_{bsn})^{\sigma^B - 1} (\psi_{mn})^{\sigma^M - 1} (P_{bsn})^{\sigma^M - \sigma^B} (P_{mn})^{-\sigma^M}, \quad (6)$$

where A_{sn} is an exogenous segment-country demand shifter and the price of product m in country n is denoted by P_{mn}. Thus, the revenue share of product m *within* brand-segment bs in country n is,

$$S_{mn} = \frac{(P_{mn}/\psi_{mn})^{1-\sigma^M}}{\sum_{k=1}^{M_{bs}} I_{kn} (P_{kn}/\psi_{kn})^{1-\sigma^M}}. \quad (7)$$

Under the empirically relevant case in which $\sigma^M > \sigma^B$, products offered by the same brand are more substitutable than products offered by different brands.[5] If brand b offers product m in country n, this leads to positive demand for product m in country n, as per equation (6), but also to a decrease in the the brand-segment price index, as observed in equation (5). The decline in the brand-segment price index reflects the greater availability of products within a brand, which causes a decline in the demand for other products offered by the same brand in country n. We refer to this phenomenon as *cannibalization*.

[5]The solution method we propose in Section 3 can also be applied when $\sigma^B > \sigma^M$. In this case, sales and production decisions across products would be complements rather than substitutes.

We assume that marginal production costs of product m in origin country o are given by,

$$c_{mo} = \kappa_b \kappa_o \kappa_s \exp\Big(\kappa_1 dist_{oh(b)} + \kappa_2 dist_{oh(b)}^2 + \kappa_3 \mathbb{1}\{o = h(b)\}\Big) \text{HP}_m^{\kappa_4} \text{Weight}_m^{\kappa_5}, \tag{8}$$

where $dist_{oh(b)}$ denotes the distance between the origin country and the headquarters country of brand b. The marginal cost of supplying a unit of product m in country n when produced in country o depends additionally on iceberg trade costs,

$$\tau_{on} = tariff_{on}^{\kappa_6} \exp(\kappa_7 dist_{on} + \kappa_8 dist_{on}^2 + \kappa_9 \mathbb{1}\{n = o\}). \tag{9}$$

At Stage 2, firms choose prices for products in each country given product offerings $\boldsymbol{I}_n^{bs} := (I_{1n}, I_{2n}, ..., I_{M_{bs}n})$ and production locations $\boldsymbol{D}^{bs} := [D_{mo}]_{m \in \Omega_{bs}, o}$. The optimal pricing decision is given by a constant markup rule,

$$P_{mn}(\boldsymbol{D}_m) = \frac{\sigma^B}{\sigma^B - 1} \times \min_{o: D_{mo}=1} \{\tau_{on} c_{mo}\}, \tag{10}$$

where $\boldsymbol{D}_m = \{D_{mo}\}_{o \in \mathcal{N}}$ and $D_{mo} = 1$ if product m is produced in o (and is zero otherwise). A tractable feature of this model is that the price of each product is independent of the set of products that the brand sells in country n. Brand b's production location choices induces a minimal marginal cost for each product m in each destination country n.[6] At the pricing stage, firms perfectly know marginal production and sourcing costs and choose to source from a unique production location.

2.2 Variable Profits

Given production location and market entry decisions, as well as optimal prices, the variable profits for product m in country n are given by,

$$\pi_{mn}(\boldsymbol{I}_n^{bs}, \boldsymbol{D}^{bs}) = A_{bsn} \big(P_{bsn}(\boldsymbol{I}_n^{bs}, \boldsymbol{D}^{bs})\big)^{\sigma^M - \sigma^B} (P_{mn}(\boldsymbol{D}_m)/\psi_{mn})^{1-\sigma^M}, \tag{11}$$

where

$$P_{bsn}(\boldsymbol{I}_n^{bs}, \boldsymbol{D}^{bs}) = \Bigg(\sum_{m=1}^{M_{bs}} I_{mn} \bigg(\frac{P_{mn}(\boldsymbol{D}_m)}{\psi_{mn}} \bigg)^{1-\sigma^M} \Bigg)^{\frac{1}{1-\sigma^M}}. \tag{12}$$

The demand shifter $A_{bsn} := A_{sn}(\psi_{bsn})^{\sigma^B - 1}$ summarizes the exogenous market conditions for brand b in segment s and country n.

[6] Minimum-cost sourcing is optimal in a large class of models where firm's profits are decreasing in marginal costs. Models with origin-destination fixed selling costs may not satisfy this property.

2.3 Plant Location and Market Entry

Brand b in segment s simultaneously chooses its production location and market entry decisions for all products Ω^{bs} so as to maximize total profits. We assume the fixed cost of offering product m in market n is denoted by,

$$F^e_{mn} = (1 - d_{mn}) \exp(\gamma_1 + \nu^e_{mn}) \tag{13}$$

and the fixed cost of producing product m in origin o is,

$$F^p_{mo} = (1 - d_{mo}) \exp(\gamma_2 + \nu^p_{mo}), \tag{14}$$

where d_{mo} follows a multinomial distribution with N possible realizations (as many as the number of origin countries), one trial, and success probability $p_{mo} = \frac{e^{\gamma_3 \mathbf{1}\{o=h(b)\}}}{\sum_{k=1}^{N} e^{\gamma_3 \mathbf{1}\{k=h(b)\}}}$.[7] Thus, we assume that brands are more likely to draw free locations that are geographically nearby their headquarters country. Fixed cost unobservables (to the econometrician) are iid and normally distributed, i.e., $\nu^e_{mn} \sim_{iid} \mathbb{N}(0, \sigma_e)$ and $\nu^p_{mo} \sim_{iid} \mathbb{N}(0, \sigma_p)$. In the data, we only observe products that are produced and sold in at least one country. By assuming that each product can be produced in some country at zero fixed cost, we directly match this feature of the data.[8]

Given fixed costs and variable profits, brand b in segment s solves:

$$\max_{\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}} \Pi_{bs}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) = \sum_{n=1}^{N} \sum_{m=1}^{M_{bs}} \{I_{mn}(\pi_{mn}(\boldsymbol{I}^{bs}_n, \boldsymbol{D}^{bs}) - F^e_{mn}) - D_{mn} F^p_{mn}\}. \tag{15}$$

Brand b solves for the $M_{bs} \times N$ matrices $\boldsymbol{I}^{bs}$ and $\boldsymbol{D}^{bs}$ to maximize its profits, internalizing the interdependence across each of the multiple binary choices it faces. A feature of this problem is that brand b must choose among many binary choices, some of which are pairwise complements and some pairwise substitutes.[9] For example, in this model, producing and selling a given product are complementary decisions. Indeed, producing product m can only reduce the marginal cost of supplying the product in any country, which increases the returns from offering it. However, selling product m in country n decreases the returns from offering product $m' \neq m$ in country n, due to cannibalization. Therefore, such binary choices are substitutes.

In what follows, we leverage such properties of the model to tractably study this problem. A brute force solution approach is clearly infeasible in light of large number of potential combinations that must be evaluated to find the solution.[10]

[7] By definition of the multinomial distribution, $\sum_{o \in \mathcal{N}} p_{mo} = 1$.

[8] Absent this assumption, the model only matches this feature of the data in certain regions of the support of the fixed cost distribution, causing a sample selection problem.

[9] We discuss formally what we mean by "complements" and "substitutes" in Section 3 and prove properties of the model in Appendix A.

[10] With only two products and 15 countries, a brute-force solution approach requires evaluating the profit function $2^{2 \times 2 \times 15} \approx 10^{18}$ times, which is computationally infeasible.

3 Solution Method

In this section, we provide a method for solving binary integer programming problems with both complementary and substitutable choices. We first provide an approach that bounds the solution of the problem. Then, we develop a novel branching algorithm akin to that in Arkolakis et al. (2023) to obtain a full solution of the problem, starting from the bounds.

3.1 First-Step Algorithm

Brands solve a problem of the form,

$$C^* = \text{argmax}_{C \in \mathcal{B}^J} \pi(C) \tag{16}$$

where $\mathcal{B}^J$ is the set of all 2^J J-dimensional binary vectors, i.e., the J-dimensional Boolean set.[11] To formalize the notion of interdependence across binary choices, we start by defining the marginal value of any discrete choice i under C.

Definition 1 (Marginal Value) *The marginal value of binary choice i under C is given by,*

$$\Delta_i^\pi(C) := \pi(C^{i\to 1}) - \pi(C^{i\to 0}) \tag{17}$$

where $C^{i\to 1}$ is the vector C with the i^{th} coordinate set to 1 (analogously for $C^{i\to 0}$).

Thus, $\Delta_i^\pi(C)$ is the change in $\pi(C)$ when changing the i^{th} coordinate from zero to one holding all other coordinates constant. We also define the cross-partial of i and j under C.

Definition 2 (Cross-Partial) *For any two choices $i, j \in \{1, 2, ..., J\}$, we define the cross-partial of i and j under C as*

$$\Delta_{ij}^\pi(C) := \Delta_j^\pi(C^{i\to 1}) - \Delta_j^\pi(C^{i\to 0}). \tag{18}$$

Thus, $\Delta_{ij}^\pi(C)$ is the difference in marginal value of coordinate i when j is set to 1 relative to when j is set to 0.

Our solution approach relies on the assumption that for all pairs of discrete choices i and j, the *sign* of $\Delta_{ij}^\pi(C)$ is known and independent of C. We state this assumption formally below.

Assumption 1 *For all $C \in \mathcal{B}^J$, and any $i, j \in \{1, 2, ..., J\}$, the sign of $\Delta_{ij}^\pi(C)$ is known and independent of C.*

In practice, Assumption 1 requires that, given model parameters, any two choices are known to be either complements or substitutes, independently of third choices. We refer to two binary choices i and j as being *complements* when $\Delta_{ij}^\pi > 0$ and as *substitutes* when $\Delta_{ij}^\pi < 0$.

[11] The solution to this problem is unique almost surely under continuously distributed determinants of payoffs, such as fixed-cost shocks.

To formalize the properties of our solution method, we exploit Tarski's fixed-point theorem (Tarski 1955), stated below.

Theorem 1 (Tarski 1955) *Let $(L, \leq)$ be a complete lattice. Suppose $F : L \to L$ is monotone increasing; i.e., for all $x, y \in L$, $x \leq y \implies F(x) \leq F(y)$. The set of fixed points of F is then a complete lattice with respect to $\leq$.*

While previous work (e.g., Jia 2008, Arkolakis et al. 2023) has defined monotone mappings on complete lattices to exploit Tarski's fixed-point theorem, our contribution is to show that we can define a monotone mapping on a space of subsets of $\mathcal{B}^J$ to allow for flexible interdependence across choices. Indeed, by defining a monotone mapping F on lattice $\mathcal{B}^J$, previous work has implicitly limited applications of Tarski's fixed-point theorem to problems featuring only either complementarities or substitutabilities. The reason is that the infimum of $\mathcal{B}^J$ is the zero vector, while the supremum is the vector of ones. As such, marginal profit evaluations for any binary choice at these extremes are only informative about the solution of the problem in cases with either complements or substitutes, but not both.

To allow for flexible interdependence patterns, we instead work on a lattice of sublattices of $\mathcal{B}^J$. Denote as $S(\mathcal{B}^J)$ the set containing the empty set and all interval-sublattices of $\mathcal{B}^J$. That is, any sublattice $A \subseteq \mathcal{B}^J$ such that for every $x, y \in A$, it is the case that $\{z \in A : x \wedge y \leq z \leq x \vee y\} \subseteq A$, where $x \vee y$ is the infimum (meet) of x and y in A and $x \wedge y$ is the supremum (join) of x and y in A.[12] If $J = 2$, then $\mathcal{B}^J = \{(0,0), (1,0), (0,1), (1,1)\}$ and

$$S(\mathcal{B}^J) = \big\{\emptyset, \{(0,0)\}, \{(1,0)\}, \{(0,1)\}, \{(1,1)\}, \{(0,0),(1,0)\}, \{(0,0),(0,1)\}, \\ \{(1,0),(1,1)\}, \{(0,1),(1,1)\}, \ldots, \{(0,0),(1,0),(0,1),(1,1)\}\big\}.$$

For instance, a sublattice of $\mathcal{B}^J$ that is not contained in $S(\mathcal{B}^J)$ is $\{(0,0,0),(1,1,1)\}$, since it does include all elements between $(0,0,0)$ and $(1,1,1)$ according to the partial order $\leq$.

The set $S(\mathcal{B}^J)$ together with the partial order $\subseteq$ yields lattice $(S(\mathcal{B}^J), \subseteq)$ with supremum $\mathcal{B}^J$ and infimum $\emptyset$. As a non-empty and finite lattice, it is a complete lattice. Thus, Tarski (1955) applies to any $F : S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ that is monotone increasing.[13]

The key methodological insight is that, by operating on $S(\mathcal{B}^J)$ rather than $\mathcal{B}^J$, we can construct a set-increasing mapping that flexibly allows for both complements and substitutes in the firm's payoff function. Indeed, define for any $\boldsymbol{C} \in S(\mathcal{B}^J)$,

$$\overline{\Omega}(\boldsymbol{C}) = \{i \in \{1, 2, ..., J\} : \Delta_i^{\pi}(\sup_i(\boldsymbol{C})) < 0\} \tag{19}$$

$$\underline{\Omega}(\boldsymbol{C}) = \{i \in \{1, 2, ..., J\} : \Delta_i^{\pi}(\inf_i(\boldsymbol{C})) \geq 0\} \tag{20}$$

[12] A lattice is a partially ordered set $(L, \leq)$ such that for any $\{a, b\} \subseteq L$, the infimum and the supremum are both in L. A complete lattice is a lattice $(L, \leq)$ such that any subset $S \subseteq L$ has a supremum and infimum in L. Note that any finite lattice is complete.

[13] Arkolakis et al. (2023) show in Appendix OA.7.2 how to write down their mapping as a function defined on a similar set of sublattices.

where, for each $i \in \{1, 2, ..., J\}$ and $j \in \{1, 2, ..., J\}$,

$$[\sup_i(\boldsymbol{C})]_j := \mathbb{1}[\Delta_{ij}^{\pi} \geq 0][\sup(\boldsymbol{C})]_j + \mathbb{1}[\Delta_{ij}^{\pi} < 0][\inf(\boldsymbol{C})]_j \tag{21}$$

$$[\inf_i(\boldsymbol{C})]_j := \mathbb{1}[\Delta_{ij}^{\pi} \geq 0][\inf(\boldsymbol{C})]_j + \mathbb{1}[\Delta_{ij}^{\pi} < 0][\sup(\boldsymbol{C})]_j. \tag{22}$$

We define mapping $F : S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ as,

$$F(\boldsymbol{C}) = \{C \in \boldsymbol{C} \colon C_i = 0 \text{ and } C_{i'} = 1 \, \forall i \in \overline{\Omega}(\boldsymbol{C}) \text{ and } \forall i' \in \underline{\Omega}(\boldsymbol{C})\}. \tag{23}$$

Intuitively, by Assumption 1, it is always possible to define a coordinate-specific supremum and infimum. Since the sign of all pairwise interdependencies are known and independent of third choices, one can always define the vectors $\sup_i(\boldsymbol{C})$ and $\inf_i(\boldsymbol{C})$ that maximize and minimize the marginal value of i, respectively. If at the "worst-case scenario" for i, $\inf_i(\boldsymbol{C})$, the marginal value of i is positive, then $i \in \underline{\Omega}(\boldsymbol{C})$ and $F(\boldsymbol{C})$ returns only vectors C with $C_i = 1$. If at "best-case scenario" for i, $\sup_i(\boldsymbol{C})$, the marginal value of i is negative, then $i \in \underline{\Omega}(\boldsymbol{C})$ and $F(\boldsymbol{C})$ returns only vectors C with $C_i = 0$.

We prove that the mapping F is set-increasing.

Theorem 2 (Set-Increasing F) *The mapping $F : S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ is set-increasing. That is, $\boldsymbol{C} \subseteq \boldsymbol{C}'$ implies that $F(\boldsymbol{C}) \subseteq F(\boldsymbol{C}')$.*

Proof. See Appendix C.1. ■

Theorem 2 follows from the fact that the minimum marginal value and the maximal marginal value for any coordinate i in any sublattice $\boldsymbol{C}$ are decreasing and increasing in the sublattice $\boldsymbol{C}$.[14] Intuitively, starting at worse "worst-case scenarios" and better "best-case scenarios" solves fewer coordinates, leading F to rule out fewer vectors.

We describe the first-step procedure in Algorithm 1 for future reference.

Algorithm 1 (First-Step Algorithm) *We implement the algorithm as follows:*

1. *Start with $\boldsymbol{C}_0 = \mathcal{B}^J$.*
2. *Apply F to obtain $\boldsymbol{C}_1 = F(\boldsymbol{C}_0)$; note that $\boldsymbol{C}_1 \in S(\mathcal{B}^J)$.*
3. *Obtain $\boldsymbol{C}_k = F(\boldsymbol{C}_{k-1})$ for $k = 1, 2, ...$ until convergence.*

Proposition 1 states that Algorithm 1 converges monotonically to a set $\boldsymbol{C}^f$ such that $C^* \in \boldsymbol{C}^f$.

Proposition 1 *Algorithm 1 has the following properties:*

1. *$\boldsymbol{C}_k = F(\boldsymbol{C}_{k-1}) \subseteq \boldsymbol{C}_{k-1}$ for all $k \in \mathbb{N}$.*

[14] More generally, $\inf_{x \in Z} f(x) \leq \inf_{x \in Y} f(x)$ and $\sup_{x \in Z} f(x) \geq \sup_{x \in Y} f(x)$ for $Y \subseteq Z$.

2. *It converges to a fixed point* $\boldsymbol{C}^f$ *in a finite number of iterations* K *i.e.,* $F(\boldsymbol{C}_T) = \boldsymbol{C}^f$ *for all* $T \geq K$*.*

3. *The solution to the problem* C^* *is an element of* $\boldsymbol{C}^f$*.*

Proof. Due to Tarski (1955), we can start at $\boldsymbol{C}_0 = \mathcal{B}^J$ and iteratively apply F until convergence to fixed point $\boldsymbol{C}^f$, which is the supremum of the set of fixed points. Convergence occurs monotonically because F is set-increasing, so that $\boldsymbol{C}_1 = F(\boldsymbol{C}_0) \subseteq \boldsymbol{C}_0$ implies that $F(\boldsymbol{C}_1) \subseteq F(\boldsymbol{C}_0)$. The algorithm convergences in a finite number of iterations, because the number of iterations is bounded above by J, the number of binary decisions that are to be solved. The sublattice $\{C^*\}$ containing the unique (almost surely) solution to the brand's problem is also a fixed point, since $C_i^* = 1$ if and only if the marginal value of i is positive i.e., $\Delta_i(C^*) \geq 0$. Thus, $C^* \in \boldsymbol{C}^f$. ■

We have proven that our iterative solution procedure bounds the argmax of any binary integer programming problem that features simultaneous complementarities and substitutabilities, provided Assumption 1 is satisfied.

3.1.1 The Multinational Brand's Interdependent Decisions

The multionational brand faces interdependent production and selling decisions across products and locations. We show that the sign of the interdependence across any pair of binary choices is independent of third choices, and satisfies Assumption 1. That is, any pair of binary choices are always either complements of substitutes regardless of the value of $(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$. Below, we list the the sign of the cross-partials that arise in our model under the empirically relevant case that $\sigma^M > \sigma^B$:

1. D_{mo} and I_{mn} are complementary across markets for any m and any $o, n \in \mathcal{N}$ due to minimum-cost sourcing. Indeed, producing product m in more markets can only decrease the marginal cost of selling product m in any given country, which increases product m sales.

2. D_{mo} and $I_{m'n}$ for $m \neq s$, and for any $o, n \in \mathcal{N}$, are substitutes. This is due to cannibalization. Indeed, producing model m in any country o can only decrease the price of model m in country n. This has the effect of attracting sales towards product m and away from other products $m' \neq m$ sold in the same country.

3. I_{mn} and $I_{m'n}$ for $m \neq s$ are substitutes due to cannibalization within markets.

4. I_{mn} and $I_{m't}$ are independent for $n \neq t$ and for all $m, m' \in \Omega^{bs}$ because demand across markets is independent.

5. D_{mn} and D_{mt} are substitutes for $n \neq t$. The marginal gains (from reduced trade costs) of producing a product in a given plant can only decrease as the product is produced in more plants, since fewer countries will find it optimal to source from the given plant.

6. D_{mn} and $D_{m'o}$ are substitutes for $m \neq s$. This substitutability is potentially strict only in markets where m and m' are both offered. The intuition in these cases is that the weak

decline in the price of some model m caused by an additional production plant leads to a lower increase in sales whenever product m' is also offered at lower prices.

In Appendix A, we provide the formal proof. We exploit this property of the brand's payoff function to apply Algorithm 1. We use this algorithm both at the estimation stage and for performing counterfactual exercises.

3.1.2 Simple Example

We provide a simple example to illustrate how Algorithm 1 works. Consider a case with three choices such that $\Delta_{12} > 0$, $\Delta_{13} < 0$, and $\Delta_{23} < 0$. A payoff function that yields this interdependence structure is $\pi(C) = \sum_{i=1}^{3} C_i[\pi_i(C) - F_i]$, $\pi_i(C) = A_i + \sum_{j \neq i} \gamma_{ij} C_j$, and $\gamma_{12} = 1/4$, $\gamma_{13} = -2$, $\gamma_{23} = -1$, $(A_1, A_2, A_3) = (1/2, 1, 4)$, and $F_i = 1$. Then,

$$\boldsymbol{C}_0 = \{(0,0,0), (1,0,0), (0,1,0), (0,0,1), (1,1,0), (1,0,1), (0,1,1), (1,1,1)\}.$$

As $\Delta_{12} > 0$, $\Delta_{13} < 0$, $\sup(\boldsymbol{C}_0) = (1,1,1)$, and $\inf(\boldsymbol{C}_0) = (0,0,0)$, it holds that

$$\sup_1(\boldsymbol{C}_0) = [\cdot, [\sup(\boldsymbol{C}_0)]_2, [\inf(\boldsymbol{C}_0)]_3] = [\cdot, 1, 0].$$

Similarly,

$$\sup_2(\boldsymbol{C}_0) = [1, \cdot, 0] \qquad \sup_3(\boldsymbol{C}_0) = [0, 0, \cdot]$$

and

$$\inf_1(\boldsymbol{C}_0) = [\cdot, 0, 1] \qquad \inf_2(\boldsymbol{C}_0) = [0, \cdot, 1] \qquad \inf_3(\boldsymbol{C}_0) = [1, 1, \cdot].$$

Under this payoff function,

$$\{\Delta_1(\pi(\sup_1(\boldsymbol{C}_0))) = -1/4 < 0, \Delta_2(\pi(\sup_2(\boldsymbol{C}_0))) = 1/4 \geq 0, \Delta_3(\pi(\sup_3(\boldsymbol{C}_0))) = 13/4 \geq 0\}$$
$$\{\Delta_1(\pi(\inf_1(\boldsymbol{C}_0))) < 0, \Delta_2(\pi(\inf_2(\boldsymbol{C}_0))) = -3 < 0, \Delta_3(\pi(\inf_3(\boldsymbol{C}_0))) \geq 0\}$$

then

$$\overline{\Omega}(\boldsymbol{C}_0) = 1 \qquad \text{and} \qquad \underline{\Omega}(\boldsymbol{C}_0) = 3.$$

Then, applying mapping F, we obtain,

$$F(\boldsymbol{C}_0) = \{C \in \boldsymbol{C}_0 : C_1 = 1, C_3 = 0\} = \{(0,1,1), (0,0,1)\} = \boldsymbol{C}_1.$$

The algorithm proceeds by applying F to $\boldsymbol{C}_1$. In effect, the algorithm has already solved for decisions 1 and 3. In the next iteration, the algorithm checks,

$$\Delta_2(\pi(\sup_2(\boldsymbol{C}_1))) = A_2 - F_2 + \gamma_{23} = 1 - 1 - 1 = -1 < 0,$$

and so $F(\boldsymbol{C}_1) = \{(0, 0, 1)\}$.

3.2 Extensions

The first extension is that different mappings $\mu \neq \pi$ may be used to solve for bounds of a subset of decisions $\{1, 2, ..., L\} \subseteq \{1, 2, ..., J\}$. In Appendix B.1, we argue that a natural such mapping μ is the value-function mapping, $\mu(B) = \max_{A \in \mathcal{B}^{J-L}} \pi(A, B)$. This mapping is useful in cases where conditional on B, solving for A is not a combinatorial problem and is computationally feasible. If one knows the sign of the interdependence across any pair of choices according to value-function mapping μ, then one can apply our first-step algorithm to find bounds for B. In Appendix B.2, we show that we can then compose different mappings μ to bound not only B, but also A.

In our setting, a mapping μ with such properties is the value of production locations for a single-product firm,

$$V(\boldsymbol{D}) := \max_{\boldsymbol{I}} \Pi(\boldsymbol{I}, \boldsymbol{D}) = \max_{\boldsymbol{I}} \sum_{n=1}^{N} I_n \Delta\pi_n(\boldsymbol{D}) - \sum_{n=1}^{N} D_n F_n^p, \tag{24}$$

where $\Delta\pi_n$ is the marginal value of selling in country n. Conditional on production locations $\boldsymbol{D}$, sales decisions are independent across countries, so this program is not combinatorial.

In Appendix A, we show that this mapping satisfies Assumption 1 – the sign of the interdependence across any pair of binary choices in the domain of V is known to be negative. Therefore, we can apply our first-step algorithm using V to bound the solution to the production location decisions, and then apply our baseline first-step method (Algorithm 1) using Π to bound decisions $\boldsymbol{I}$ (and also $\boldsymbol{D}$).[15] In Appendix B.2, we develop Algorithm 3, a more general version of equation (24) that composes multiple mappings to bound production location decisions of multi-product firms. We apply this more general version both for estimating the model and for performing counterfactual exercises.

Such value-function mappings are generally useful for two reasons. First, note that in equation (24), only the sign across production location choices is required. By optimizing over the other choices, it becomes unnecessary to know the sign of the cross-partial across sales decisions, or across production and sales decisions. Of course, it comes at the cost of only allowing the researcher (absent any other mappings) to learn about the bounds of $\boldsymbol{D}$. Second, by solving for a subset of the decisions, this mapping might be more informative. To see this, recall that because production location and sales decisions for the same product are complements, the worst-case marginal value

[15] In this particular case with a single-product firm, because V exhibits substitutabilites only, our first-step algorithm applied to V is equivalent to the first-step algorithm in Arkolakis et al. (2023) applied to V.

of a production location decision under Π at the first iteration is evaluated under no sales countries (other than the free location). At this extreme, it is likely that the worst-case marginal value is negative, and therefore we cannot solve for a 1. Under the mapping that uses equation (24), this is no longer the case: the marginal benefit of a production location decision under V allows optimal sales decisions to adjust when reconfiguring production choices. Thus, mappings as those in equation (24) serve to potentially relax knowledge of cross-partials, and to improve the informativeness of the bounds.

The second extension, discussed in detail in Appendix B.3, relaxes Assumption 1 under a restriction on the structure of π. If π can be written as a sum of K components i.e., $\pi(C) = \sum_{k=1}^{K} \pi_k(C)$, what is required is not that the sign of the cross-partial across any pair of choices is constant given any C, but that it is constant for each of the K components of π. For instance, suppose that two binary choices are substitutes under variable profits (e.g., due to cannibalization forces) but complements under fixed costs (e.g., due to economies of scope). This case can be handled by slightly modifying mapping F to account for decision-*component*-specific infima and suprema. Under economies of scope, the worst-case marginal value for a production location evaluates its marginal contribution under maximal cannibalization through variable profits (assuming that production locations in other countries for that product are set to the upper bound), and under maximal fixed costs (assuming production locations in other countries are set to the lower bound). The best-case marginal value is obtained under minimized cannibalization and fixed costs.

3.3 Second-Step Algorithm

Algorithms 1 and the extensions converge to sets of possible vectors (or matrices) containing the solution. To obtain a full solution, we extend and refine the branching algorithm in Arkolakis et al. (2023) in the context of binary integer programming problems with both complements and substitutes.

Intuitively, the first-step algorithms can be seen as solving for a subset of the binary decisions of the brand e.g., if $\underline{D}_{mo} = 1$, this is equivalent to stating that $D^*_{mo} = 1$; if $\overline{D}_{mo} = 0$, the algorithm states that $D^*_{mo} = 0$. For those binary choices that remain unsolved, the branching algorithm picks one, conditions on values of 0 and 1 in separate "branches" and then solves conditional on each of these values by applying in each "branch" the first-step procedure. This process is repeated until solutions are found in each of the branches. Given a finite number of possible solutions, these are compared by brute-force to find the maximum. As discussed in Arkolakis et al. (2023), a limitation of this procedure is that there is no guarantee that it will converge quickly to the solution. Indeed, in some cases, we find that the branching algorithm performs poorly and does not yield a solution in a reasonable amount of time. In practice, we implement the branching procedure but limit the number of possible profit evaluations – if that limit is surpassed, we simply use the bounds obtained from the first-step procedure.

In Appendix C.3, we characterize branching also as a monotone-increasing function from the power set of $S(\mathcal{B}^J)$ (which is a complete lattice) to itself. We prove that B converges to a set of

singleton sublattices, one of which must include the solution to the problem. Our characterization also enables us to refine the procedure in Arkolakis et al. (2023) by discarding branches at which the first-step mapping returns the empty set. We find that discarding such branches yields important computational gains. The intuition is as follows. Suppose that we are in a branch where we have conditioned on production in Spain, and France, and not in Italy. Importantly, these choices have been fixed by branching, and not by the first-step algorithm. When applying our first-step mapping F at this branch, we get the empty set if, for instance, conditional on producing in France, the firm finds it optimal *not* to produce in Spain. That is, mapping F is not consistent with this branch, and therefore we know that the solution will never contain production in Spain and France, but not in Italy. In sum, our computational refinement is effectively to continuously use F to check whether the coordinates fixed by the branching procedure are consistent with payoff maximization.

4 Estimation

We estimate the model in two steps. In the first step, we estimate demand and marginal cost parameters, including the demand elasticities for each nest and product-, brand- and segment-specific demand shifters, and trade and marginal production costs. In the second stage, we estimate the determinants of the distributions of fixed costs associated with selling a product in a country and producing a product in a given location. We develop a moment inequality approach based on our solution method and moments that are monotone in brands' production and offering decisions.

4.1 Data

Our analysis uses information on global car sales, production, prices and product characteristics, as well as data on measures of trade costs (namely geographic distances and tariffs), and country demographics.

Our global car sales and production data corresponds to 2019 and is compiled by IHS Markit using new car registrations of individual car products identifying the plant of assembly and the country of sale. The dataset contains information for 1,130 car products, 124 car brands, 56 countries of sale, and 30 production countries. A product is defined as a parent-brand-body type-nameplate (e.g. Toyota Corolla Sedan).[16] We allow for 7 market segments defined by body type categories: convertible, coupe, hatchback, MPV, SUV, Sedan and Wagon. In total, our dataset contains information on production and sales for 328 brand-segments. The average brand-segment has 3.5 car models in its portfolio, has production facilities in 2.2 countries and sells in 18.9 markets. The average model is produced in 1.4 countries and sold in 12.2 markets.

[16] In our sample, we consider seven bodytypes (convertible, coupe, hatchback, MPV, SUV, sedan, wagon); We drop ultra–luxury brands. We remove production/sales countries without full segment coverage; very small origin–destination flows (< 6 units) are set to zero. We retain models with ≥ 30 global units, brand–segment cells with $\geq$ 1,000 units, and production origins with at least 0.1% of world output ($\approx$ 5,000 units). Price/characteristics (12 countries, 2018–2019) are aggregated to quantity–weighted annual model–market–origin means; list prices are trimmed at \$150,000 and missing technical attributes are imputed using within–model/segment/brand averages.

Table 1: Summary Statistics

Number of:	Mean (unweigh.)	Mean (weight.)	p25	p50	p75	p90	Max
Models (per brand-segment)	3.5	8.4	1	2	4	7	19
Sales countries (per model)	12.2	25.0	1	3	21	42	56
Production countries (per model)	1.4	2.2	1	1	1	2	8
Models sold (per market)	35.8	128.9	10	23	50	88	332
Sales countries (per brand-segment)	18.9	49.8	1	5	40	62	77
Production countries (per brand-segment)	2.2	5.2	1	1	3	5	10
Share exported (by model-prod. loc.)	37.5%	40.0%	0%	9.1%	89.6%	100%	100%
Share produced in HQs (by brand-segment)	77.7%	43.4%	55.6%	100%	100%	100%	100%
Share sold in HQs (by brand-segment)	61.3%	29.9%	20.3%	74.6%	100%	100%	100%

In addition to the global production data, we use information on car prices and characteristics at the product level for Australia, Brazil, China, Germany, France, India, Italy, Japan, Mexico, Spain, the United Kingdom and the United States for 2018 and 2019. This dataset contains information about the manufacturer-suggested retail price (MSRP) and a host of characteristics of the car product, among which we use horsepower, car weight, and car length. For these car products and countries, we use the global car production dataset to identify their assembly location.

We measure trade costs using bilateral applied tariffs for each pair of countries, sourced from ITC (2022). Additionally, we utilize standard bilateral geographical distance measures from CEPII (Mayer and Zignago 2011). Finally, we gather data on country characteristics, including GDP, population, average family size, and urbanization rates, from the World Bank, to project demand shifters on country-level characteristics.

4.2 Demand and Marginal Costs

First, we describe how we estimate trade costs and marginal production costs. From equations (8)-(10), we write estimating equation,

$$\log P_{mn} = \alpha_n + \kappa_1 \log \text{tariff}_{on} + \kappa_2 \text{dist}_{on} + \kappa_3 \text{dist}_{on}^2 + \kappa_4 \mathbb{1}\{o = n\} + \alpha_{mo} + \varepsilon_{mn}. \tag{25}$$

Relative to equations (8)-(10), equation (25) also includes an error term ε_{mn}. We include a destination fixed effect, α_n, to control for destination country-specific unobserved factors and policies, such as VAT or sales taxes, that determine observed final consumer prices (MSRP). We interpret ε_{mn} as a measurement error and assume it is unobserved at the time of making offering, production, or sourcing decisions. Under this assumption, equation (25) can be estimated via OLS.

We identify trade costs from differences in prices of the same car model built in the same location

across destination markets. Model–origin fixed effects, α_{mo}, account for the marginal production cost for model m at origin o, so residual price variation across destinations for a given model-origin pair is attributed to trade costs. This regression's estimates allow us to construct trade costs as

$$\hat{\tau}_{on} = tariff_{on}^{\hat{\kappa}_1} \exp(\hat{\kappa}_2 dist_{on} + \hat{\kappa}_3 dist_{on}^2 + \hat{\kappa}_4 \mathbb{1}\{n = o\}). \tag{26}$$

Following equation (8), we interpret the estimated model-origin fixed effect, $\hat{\alpha}_{mo}$, as the marginal cost of producing model m in production origin o, and regress this estimate on origin, brand and segment fixed effects, a quadratic function of distance between the headquarters' location for the brand that produces model m and origin o, a home market effect, and model characteristics.

$$\hat{\alpha}_{mo} = \kappa_b + \kappa_o + \kappa_s + \kappa_5 dist_{oh(b)} + \kappa_6 dist_{oh(b)}^2 + \kappa_7 \mathbb{1}\{o = h(b)\} + \kappa_8 \log HP_m + \kappa_9 \log Weight_m. \tag{27}$$

Using the estimates of this regression, we construct predicted marginal costs, c_{mo}, for every potential model and origin.

To estimate the demand shifters and elasticities, we leverage information on both quantities and prices to construct revenue shares, which permits exploiting equations (4) and (7) for estimation. Based on such equations, we follow a two-step estimation procedure.

First, we estimate σ^M and ψ_{mn} using equation (7) and exploiting variation in product-country revenue shares within brands, segments, and destination countries. Combining equation (3) for ψ_{mn} with equation (7), we obtain estimating equation,

$$\begin{aligned} \log S_{mn} = {} & (1 - \sigma^M) \log P_{mn} + \alpha_{bsn} \\ & + \tilde{\phi}_m + \tilde{\beta}_3 \log(HP_m) \log(GDPpc_n) + \tilde{\beta}_4 \log(Weight_m) \log(GDPpc_n) + \tilde{u}_{mn}, \end{aligned} \tag{28}$$

where $\tilde{\theta} := (\sigma^M - 1)\theta$. We estimate this regression by instrumenting the price of each product, $\log(P_{mn})$, with $\log \text{dist}_{on}$ and $\log \text{tariff}_{on}$ to deal with potential endogeneity issues arising from measurement error on prices.[17]

To estimate σ^B and ψ_{bsn}, we regress the log of S_{bsn} (given in equation (4)) on the log prices of that brand-segment and characteristics of the brand. We construct these prices by using the regression equation (28) and noticing that the structure of our model implies that $\log(P_{bsn}) \equiv \frac{\alpha_{bsn}}{(\sigma^M - 1)}$. Thus, we estimate

$$\log(S_{bsn}) = (1 - \sigma^B) \log(P_{bsn}) + \phi_{sn} + \phi_b + \check{\beta}_1 \mathbb{1}\{h(b) = n\} + \check{\beta}_2 dist_{nh(b)} + \check{\beta}_3 dist_{nh(b)}^2 + \check{u}_{bsn}, \tag{29}$$

where $\check{\theta} := (\sigma^B - 1)\theta$ and $\phi_{sn} = (\sigma^B - 1) \log(P_{sn})$. In the regression above, we instrument $\log(P_{bsn})$,

[17] Note that, in the model, u_{mn} is not correlated with prices by assumption.

with the following distance and tariff shifters,

$$\log \check{\text{dist}}_{bsn} = \log\left(\sum_{m=1}^{M_{bs}}\left(\frac{\text{dist}_{on}}{\hat{\psi}_{mn}}\right)^{1-\hat{\sigma}^M}\right)^{\frac{1}{1-\hat{\sigma}^M}}, \quad \log \check{\text{tariff}}_{bsn} = \log\left(\sum_{m=1}^{M_{bs}}\left(\frac{\text{tariff}_{on}}{\hat{\psi}_{mn}}\right)^{1-\hat{\sigma}^M}\right)^{\frac{1}{1-\hat{\sigma}^M}}$$

Finally, we recover segment-country demand shifter, A_{sn}, by exploiting that equation (6) implies,

$$\sum_b \sum_{m=1}^{M_{bs}} P_{mn}C_{mn} = A_{sn}\sum_b \sum_{m=1}^{M_{bs}} (P_{bsn})^{\sigma^M-\sigma^B}(\psi_{bsn})^{\sigma^B-1}(\psi_{mn})^{\sigma^M-1}(P_{mn})^{1-\sigma^M} = A_{sn}P_{sn}^{1-\sigma^B},$$

with $P_{sn}^{1-\sigma^B} = \sum_b \left(\frac{P_{bsn}}{\psi_{bsn}}\right)^{1-\sigma^B}$.

The main challenge for the estimation of this equation is that we only observe prices in a subset of destination countries. To overcome this shortcoming of our data, we leverage our previous estimates for the demand shifters ψ_{mn} and ψ_{bsn}, and our estimates of marginal and trade costs, to construct a predicted price $\hat{P}_{mn}$ for product m in destination country n being supplied from the observed production location o, and we construct the related brand-segment-market level price index P_{bsn} using its formula (given in equation (5)).

We then recover the segment-country demand shifter, A_{sn}, as

$$\hat{A}_{sn} = \sum_b \sum_{m=1}^{M_{bs}} \hat{P}_{mn}C_{mn}\hat{P}_{sn}^{1-\sigma^B}. \tag{30}$$

This concludes our estimation strategy for the demand and marginal cost parameters of the model.

4.3 Fixed Costs: Moment Inequality Estimation

Given the demand and marginal cost parameters, we use moment inequalities to estimate fixed-cost parameters. Given the large number of interdependent binary choices, implementing (simulated) GMM or MLE is generally infeasible, as such methods require computing the full solution of the brands' problem under multiple parameter vectors and draws of unobserved shocks.[18] To construct moment inequalities, we combine Algorithm 1 and 3 to obtain computationally feasible upper and lower bounds on brands' optimal choices. We combine these bounds on the firms' optimum with moment functions that are weakly increasing in the choices, thereby formulating valid moment inequalities.[19]

Formally, conditional on the demand and marginal cost estimates, a realization of brand-

[18] As discussed in Section 3, while our two-step solution procedure is guaranteed to find the solution, only the first step of the algorithm converges rapidly. The branching algorithm has the disadvantage of not being guaranteed to converge quickly to the solution.

[19] We do not need moments to be increasing, we need them to be monotonic.

segment fixed cost shocks $\boldsymbol{\nu}_{bs} \equiv \{\nu^e_{mn}, \nu^p_{mn}\}_{m\in\Omega^{b,s},n}$ and a realization of the free plant for each product, our first-step solution procedure yields lower bounds $\{\underline{I}_{mn}(\boldsymbol{\nu}_{bs}), \underline{D}_{mn}(\boldsymbol{\nu}_{bs})\}_{m,n}$ and upper bounds $\{\overline{I}_{mn}(\boldsymbol{\nu}_{bs}), \overline{D}_{mn}(\boldsymbol{\nu}_{bs})\}_{m,n}$. Consider monotone moment functions $g(I, D; \boldsymbol{\gamma})$. We construct moment inequalities of the form,

$$\frac{1}{M}\sum_{m=1}^{M} \mathbb{E}_\nu[g_m(\underline{\boldsymbol{I}}_m(\boldsymbol{\nu}_{bs}), \underline{\boldsymbol{D}}_m(\boldsymbol{\nu}_{bs}); \boldsymbol{\gamma})] \leq \frac{1}{M}\sum_{m=1}^{M} g_m(\boldsymbol{I}_m, \boldsymbol{D}_m; \boldsymbol{\gamma}) \leq \frac{1}{M}\sum_{m=1}^{M} \mathbb{E}_\nu[g_m(\overline{\boldsymbol{I}}_m(\boldsymbol{\nu}_{bs}), \overline{\boldsymbol{D}}_m(\boldsymbol{\nu}_{bs}); \boldsymbol{\gamma})],$$

where $\mathbb{E}_\nu$ denotes that the expectation is with respect to random draws from the distribution of fixed cost unobservables, and M denotes the total number of products in our sample. To implement these moment inequalities, we simulate 10 draws from the distribution of fixed cost shocks, apply our solution method for each of the draws under each parameter vector $\boldsymbol{\gamma}$, and implement the procedure from Andrews and Soares (2010) to determine whether such a parameter vector $\boldsymbol{\gamma}$ is included in the 95% confidence set.

The procedure in Andrews and Soares (2010) is applicable to a set of unconditional moment inequalities. We employ the following monotone moment functions to derive bounds on $\boldsymbol{\gamma}$.

1. Number of destination countries in which product m is offered:

$$g^1_m(\boldsymbol{I}_m, \boldsymbol{D}_m) = \sum_{n=1}^{N} I_{mn}.$$

   This moment intuitively bounds parameter γ_0 (the mean fixed cost of offering product m in country n). If the simulated lower bound of this moment is larger than the observed value at $\boldsymbol{\gamma}$, then γ_1 is too small; if the simulated upper bound of this moment is smaller than the observed value, then γ_1 is too large.

2. Sold in destination in profit bin k:

$$g^{2,k}_m(\boldsymbol{I}_m, \boldsymbol{D}_m) = \sum_{n=1}^{N} I_{mn}\mathbb{1}\{n \in B^k(\boldsymbol{\gamma})\} \text{ for } k \in \{2,3,4,5,6,7,8\},$$

   where $B^k(\boldsymbol{\gamma})$ are percentile categories of profitability index,

$$\pi^{IV}_{mh(b)n} = \frac{1}{\sigma^B} A_{sn}\psi^{\sigma^B-1}_{bsn}\Big(\frac{\sigma^B}{\sigma^B-1}\tau_{h(b)n}c_{mh(b)}/\psi_{mn}\Big)^{1-\sigma^B} - \gamma_0 - \gamma_1 dist_{nh(b)},$$

   which are the profits that the brand earns in market n if it were only selling model m and sourcing it from its headquarters country. The model predicts that more profitable products are more likely to be offered. However, the extent to which this is true also depends on the variance of the fixed cost distribution. If the variance of unobserved factors affecting fixed costs (σ_e) is very large, product offerings should appear less related to demand and marginal cost fundamentals. Therefore, the moments $g^{2,k}_m$ (and the bounds thereof) provide

information about σ_e.

3. Number of production locations entered:

$$g_m^4(\boldsymbol{I}_m, \boldsymbol{D}_m) = \Big(\sum_{o=1}^{N} D_{mo} \mathbb{1}\{o \neq h(b)\} \Big) \mathbb{1}\{ \sum_{o=1}^{N} D_{mo} > 1 \}.$$

This moment provides information on γ_2, the mean fixed cost of setting up a (non-free) production plant for product m. At low production plant fixed costs, it is more likely that the brand chooses to set up a plant for a product. The opposite is true at high fixed costs.

4. Number of production locations entered in a given origin profit bin:

$$g_m^{5,k}(\boldsymbol{I}_m, \boldsymbol{D}_m) = \Big(\sum_{o=1}^{N} D_{mo} \mathbb{1}\big\{o \in W^k(\boldsymbol{\gamma})\big\} \mathbb{1}\{o \neq h(b)\} \Big) \mathbb{1}\{ \sum_{o=1}^{N} D_{mo} > 1 \}$$

for $k \in \{2, 3, 4, 5, 6, 7, 8\}$, where origin profit bins are defined using percentile categories of,

$$\sum_{n=1}^{N} \Big[\mathbb{1}\{0 \leq \pi_{mh(b)n}^{IV} < \pi_{mon}^{IV}\}(\pi_{mon}^{IV} - \pi_{mh(b)n}^{IV}) + \mathbb{1}\{\pi_{mh(b)n}^{IV} < 0 \leq \pi_{mon}^{IV}\}\pi_{mon}^{IV} \Big] - \gamma_2 \qquad (31)$$

which is the marginal value of producing model m in origin o assuming it is the only model that the brand produces and that there is only a free plant in the brand's headquarters country. This is a proxy for the marginal value of producing model m in origin o. To rationalize plant entry in locations with very small marginal values, the model requires a higher variance of the production plant fixed cost, σ^p.

5. Produced in the HQ country:

$$g_m^{3,k}(\boldsymbol{I}_m, \boldsymbol{D}_m) = \sum_{n=1}^{N} D_{mo} \mathbb{1}\big\{dist_{oh(b)} \in B_{dist}^k\big\} \text{ for } k \in \{2, 3, 4, 5, 6, 7, 8\},$$

The probability that a product is produced in the brand's headquarters country provides information on γ_3, the parameter determining the relationship between distance to the headquarters country and the probability that any given country is the zero fixed cost location for that product.

5 Estimation Results

In this section, we present our estimates of the demand, trade cost, marginal production cost, and fixed cost parameters. Table 2 summarizes the demand and marginal cost estimates. We first discuss trade and production cost estimates, then turn to demand and its implications for cannibalization, and conclude with fixed cost estimates obtained via moment inequalities.

Table 2: Demand and Marginal Cost Parameter Estimates

	Parameter	Standard Error
Demand (Inner Nest) ($N = 7,116$)		
σ^M	8.22	1.26
(log) Horsepower × (log) GDPpc ($\tilde{\beta}_4$)	0.203	0.060
(log) Weight × (log) GDPpc ($\tilde{\beta}_5$)	0.387	0.124
Demand (Upper Nest) ($N = 3,002$)		
σ^B	3.52	0.155
Home Bias ($\check{\beta}_1$)	0.835	0.114
Distance to Brand HQ ($\check{\beta}_2$)	-0.092	0.025
Squared Distance to Brand HQ ($\check{\beta}_3$)	0.002	0.002
Marginal Costs (log) ($N = 7,405$)		
Trade Costs		
(log) Tariff (κ_1)	0.708	0.068
Distance to Production Country (κ_2)	0.015	0.002
Squared Distance to Production Country (κ_3)	-0.000	0.000
$\mathbb{1}\{o = n\}$ (κ_4)	-0.022	0.007
Production Costs		
Distance to Brand HQ (κ_5)	0.021	0.004
Squared Distance to Brand HQ (κ_6)	-0.001	0.000
$\mathbb{1}\{o = h(b)\}$ (κ_7)	0.001	0.021
(log) Horsepower (κ_8)	0.660	0.037
(log) Weight (κ_9)	0.911	0.060

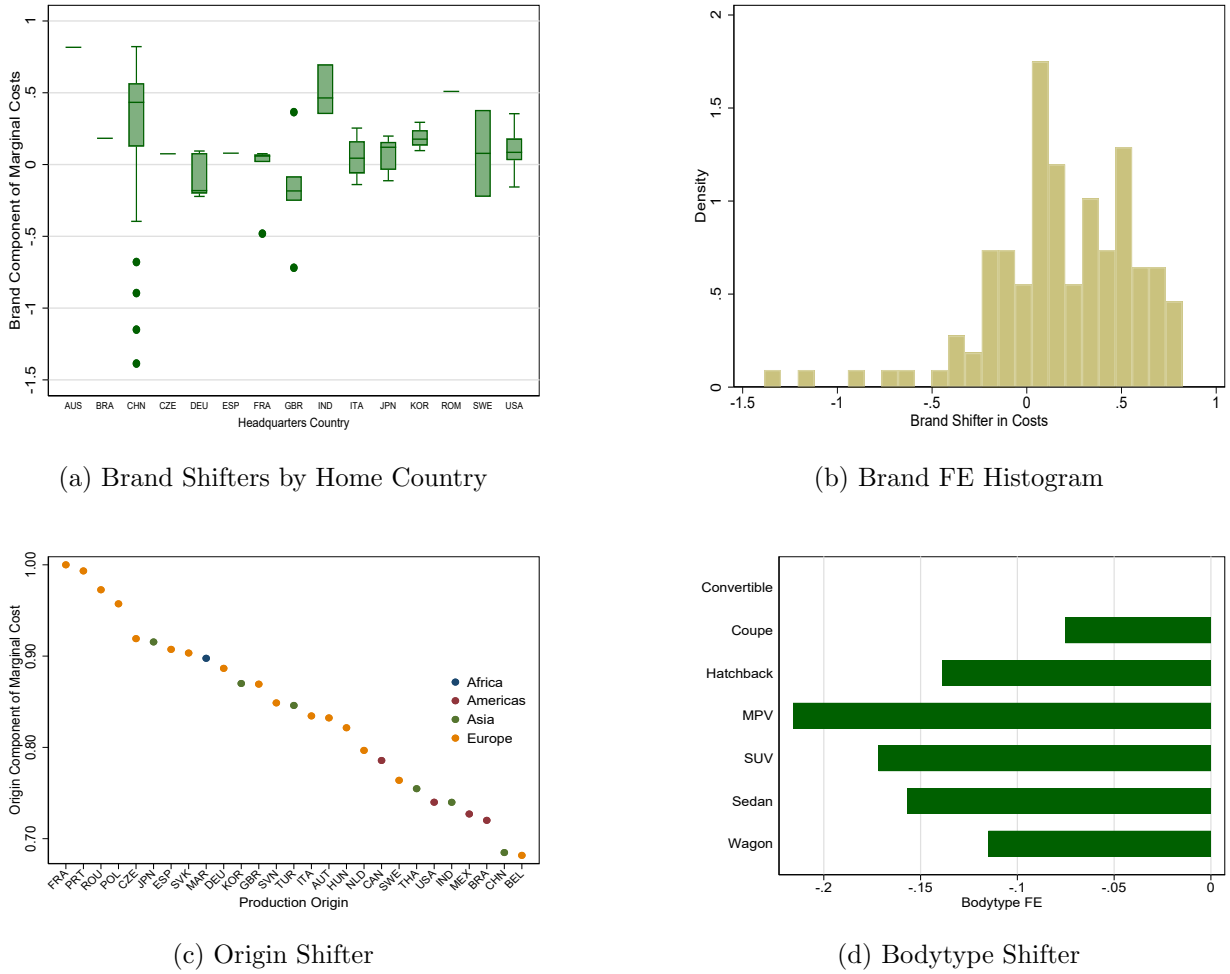
Notes: The inner nest demand estimation is conducted via IV-2SLS, instrumenting prices with tariffs and distance between the origin and destination country. Product and brand-segment-destination country fixed effects are included in the specification. Robust standard errors are reported. The upper nest demand estimation is conducted via OLS. We include segment-destination country fixed effects and report robust standard errors. The marginal cost specification is estimated via IV-2SLS, instrumenting brand-segment prices with a weighted index of tariffs and distance between the origin and destination country. The specification includes destination country fixed effects, and robust standard errors are reported.

5.1 Trade and Production Cost Estimates

We estimate trade and marginal cost parameters using 7,684 list prices for the products in the 12 countries for which we have data. We report the estimates in Table 2. The estimated elasticity of trade costs with respect to tariffs, κ_1, is 0.708. We find that trade costs increase concavely in geographic distance, with $\hat{\kappa}_2 > 0$ and statistically significant, but $\hat{\kappa}_3 < 0$ and statistically significant. We also find that conditional on geographic distance and tariffs, producing in the same country as the destination country reduces transport costs, by approximately 2%.

Turning to the marginal cost decomposition, we find that production costs increase concavely in the distance between the production location and the brand's headquarters ($\hat{\kappa}_5 > 0$, $\hat{\kappa}_6 < 0$), while the indicator for producing in the headquarters country (κ_3) is economically and statistically insignificant, suggesting that the quadratic distance function fully captures proximity advantages. We estimate statistically significant elasticities of marginal costs to horsepower (κ_4) and weight

Figure 1: Estimates of Marginal Cost Function



(a) Brand Shifters by Home Country

(b) Brand FE Histogram

(c) Origin Shifter

(d) Bodytype Shifter

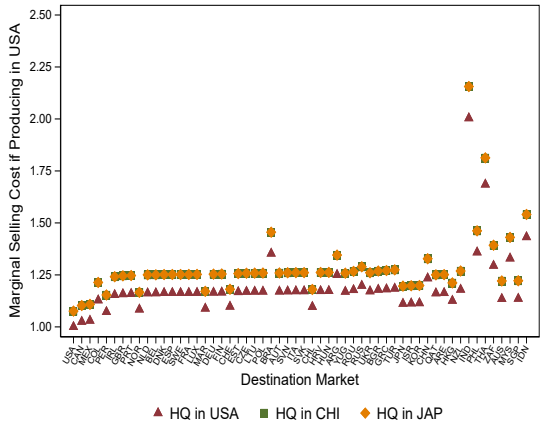
Notes: In this figure, we graphically represent the estimates for the different components in Equation 27. In Panel (a), we present averages of κ_b by the home market of the brands. In Panel (b), we present a histogram for the estimated κ_b. In Panel (c), we present the values of the origin fixed effects, κ_o. In Panel (d), we present the values of the segment shifters, κ_s.

(κ_5) equal to 0.660 and 0.911, respectively.

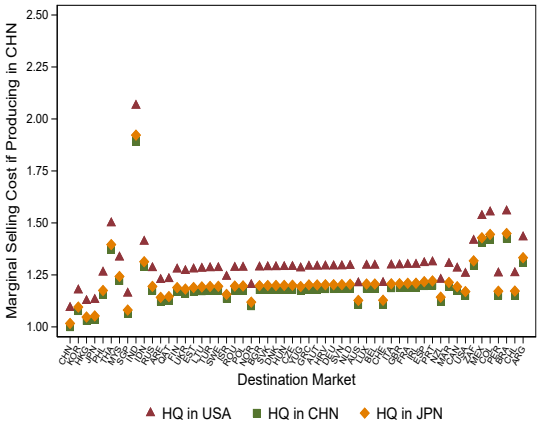
In Figure 1, we present the estimates of the fixed effects for brands, production locations, headquarters locations, and body type. As shown in Figures 1a and 1b, there is substantial heterogeneity in brand productivity, that can be partially explained by the headquarters location of those brands, with China and Western Europe accounting for the majority of low-cost brands. Figure 1c highlights that production locations differ quite substantially in the origin-specific component of marginal costs. Countries such as France, Portugal, and several Eastern European countries have the largest origin shifters, while countries such as China, Belgium, Mexico, and Brazil have the smallest, with substantial heterogeneity within regions. Figure 1d shows that sports cars are substantially costlier to produce.

In Figure 2, we translate our estimated marginal costs into the delivered costs of producing

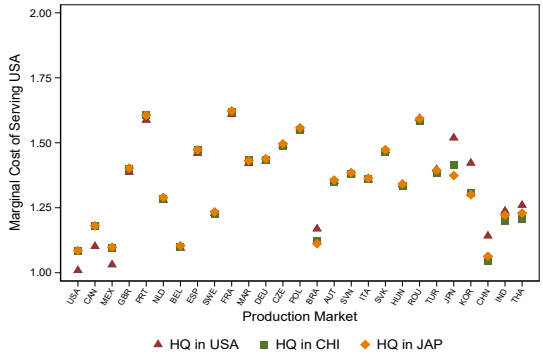
Figure 2: Predicted Marginal Costs by Origin and Destination



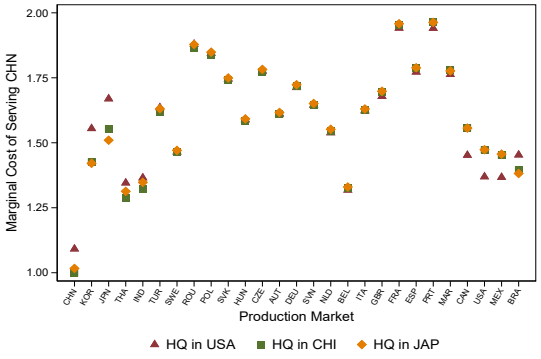
(a) Marginal Costs of Producing in the US



(b) Marginal Costs of Producing in China



(c) Marginal Costs of Selling in the US

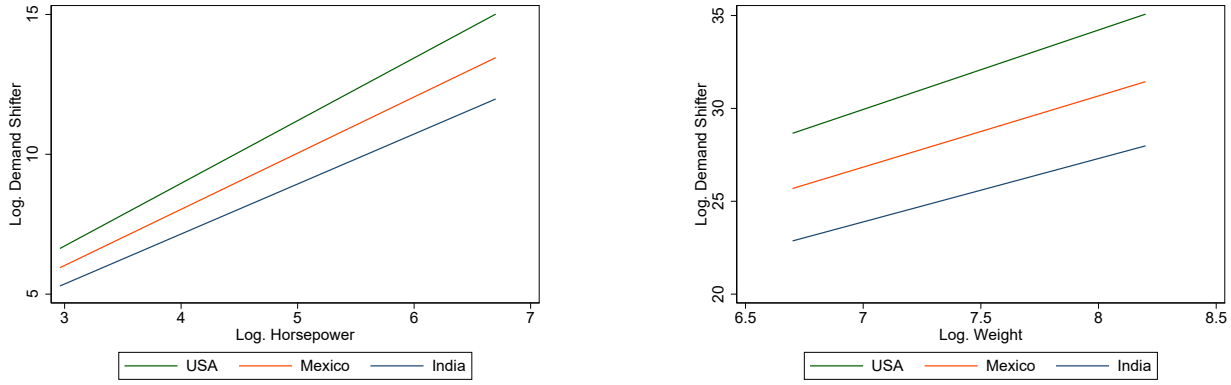


(d) Marginal Costs of Selling in China

Notes: In Panels (a) and (b) we plot the delivered marginal cost of producing in the United States and China and selling in each destination market, i.e., $\tau_{USA,n}\hat{c}_{b,USA}$ and $\tau_{CHN,n}\hat{c}_{b,CHN}$ for all destinations n. In Panels (c) and (d), we show the marginal cost of selling in the United States and China from each production origin, i.e., $\tau_{o,USA}\hat{c}_{b,o}$ and $\tau_{o,CHN}\hat{c}_{b,o}$ for all origins o. In computing the marginal costs, we condition on firms headquartered in the United States, China, and Japan, and abstract from the model-specific components.

in the United States and China and selling in each destination market, in Panels (a) and (b). Markets are sorted by distance to the production country, and we consider firms with headquarters in the United States, China, and Japan. Delivered costs increase as one moves farther from the headquarters country, and rise notably with tariffs, as highlighted by cases such as Brazil, India, and Thailand; by contrast, the effect of distance on trade costs is small. In Panels (c) and (d), we show the marginal costs of selling in the US and China from each production origin. These estimates show that countries like China, Mexico, and the United States are particularly productive.

Figure 3: Heterogeneous Preferences for Car Characteristics

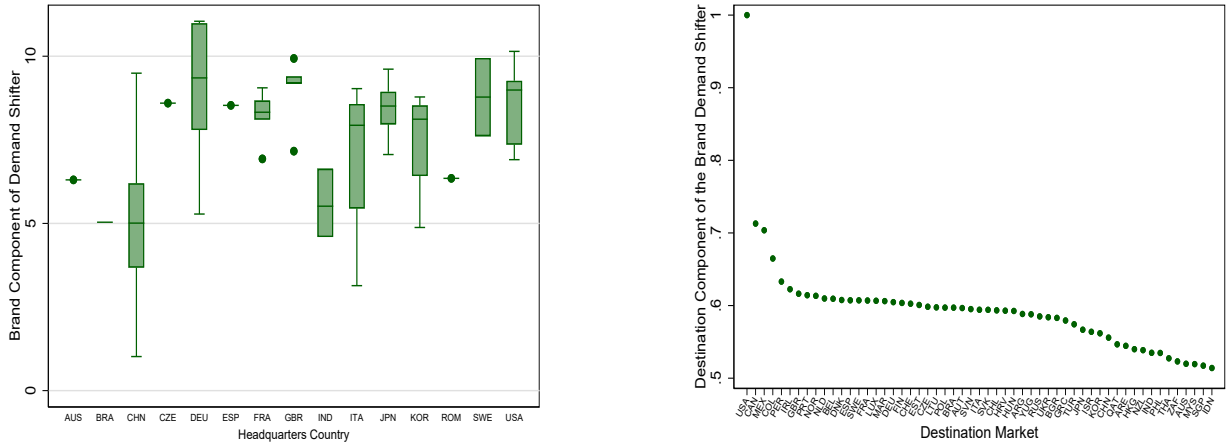


(a) Preference for Horsepower

(b) Preference for Weight

Notes: In this figure, we present the estimated effects of horsepower and weight on the product demand shifters from Equation 28, conditioning on three values of GDP per capita (USA, Mexico, and India). In Panel (a), we show the relationship for horsepower, and in Panel (b), for weight.

Figure 4: Estimates of Brand Fixed Effects in Demand



(a) Brand FE by Home Country

(b) Destination-Specific Component of the Brand Demand Shifter

Notes: In this figure, we present the values of the demand brand shifters, $\check{\phi}_b$, introduced in Equation 29. In Panel (a), we show the values of these demand shifters by the home country of the brands. In Panel (b), we present the value of $\check{\beta}_1 \mathbb{1}\{h(b) = n\} + \check{\beta}_2 dist_{nh(b)} + \check{\beta}_3 dist^2_{nh(b)}$ for a firm headquartered in the US normalized by the value in this country.

5.2 Demand Estimates

We report demand estimates in Table 2. The estimated within-brand-segment elasticity, σ^M, is 8.22, and the across-brand elasticity within market-segments, σ^B, is 3.52. Our estimates imply $\hat{\sigma}^M > \hat{\sigma}^B$ (statistically significantly), which has two implications. First, there is sizable within-brand cannibalization: introducing a new product in a country significantly erodes the sales of the brand's other products in that market. Second, the price-marginal cost markup, $\sigma^B/(\sigma^B - 1) \approx$

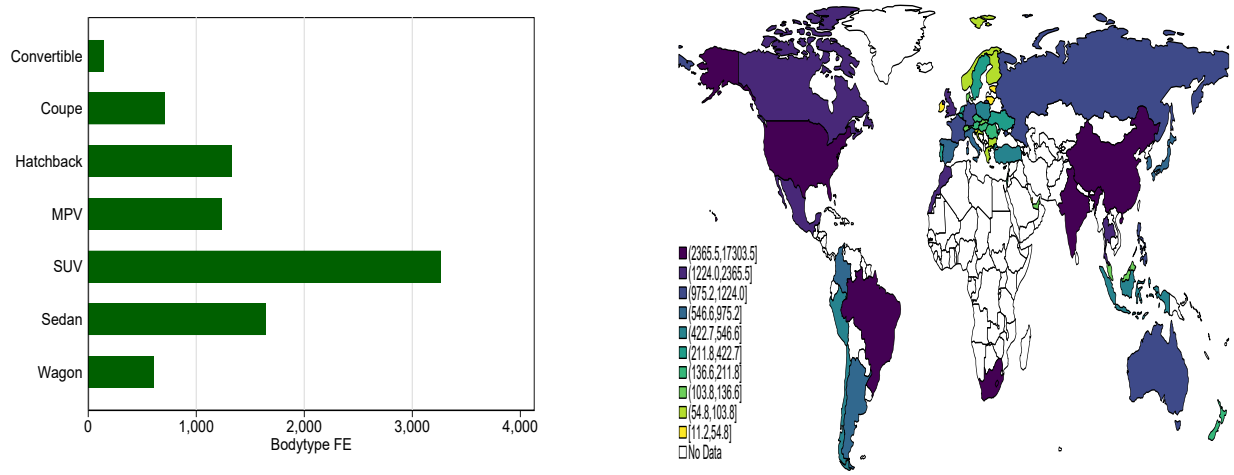
1.40, implies that firms charge approximately 40% above marginal cost.

Regarding the product-destination demand shifters, we find that, conditional on product fixed effects, products with both more horsepower and more weight, are preferred in countries with higher GDP per capita, as displayed in Figure 3.

At the upper nest, we estimate a home bias effect in demand of 0.84, in line with Coşar et al. (2016). In this nest, we also estimate brand fixed effects that capture vertical quality differences across brands. In Panel (a) of Figure 4, we plot these brand fixed effects. Comparing with Figure 1a, we find a positive relationship between cost and quality: expensive brands from countries like the United States, United Kingdom, Italy, or Germany tend to have large quality shifters, whereas brands from China and India are perceived as relatively lower quality. In Panel (b) of Figure 4, we plot the destination-specific component of the brand demand shifter for a brand headquartered in the United States, normalized by the shifter in its home market. The home market effect is substantial and declines monotonically with distance.

Averages across destination markets and across segments of the estimates for the market-segment shifters, A_{sn}, are shown in Figure 5. On average, SUVs are the segment with the largest demand shifters, whereas convertibles are the body type with the smallest shifters. These demand shifters are largest in larger and/or richer countries, such as the United States, China, India, United Kingdom, or Canada.

Figure 5: Estimates of Market-Segment Shifters



(a) Average A_{sn} by Segment

(b) Average A_{sn} by Destination Market

Notes: In this figure, we present the estimates of A_{sn} from Equation (30). In Panel (a), we present averages of A_{sn} by segment. In Panel (b), we present averages of A_{sn} by destination market.

With our demand estimates, we can translate within–brand revenue shares into an economically interpretable cannibalization rate. Let $\kappa \equiv (\sigma^M - \sigma^B)/(\sigma^M - 1)$. Using our baseline elasticities, $\kappa \approx 0.65$. Then

$$\log \Delta^c_{mn} = \kappa \log S_{mn} \quad \Longrightarrow \quad \Delta^c_{mn} = S^\kappa_{mn},$$

so that $r_{mn} \equiv 1 - \Delta^c_{mn} = 1 - S^\kappa_{mn}$ is the fraction of the sales that model m would obtain if

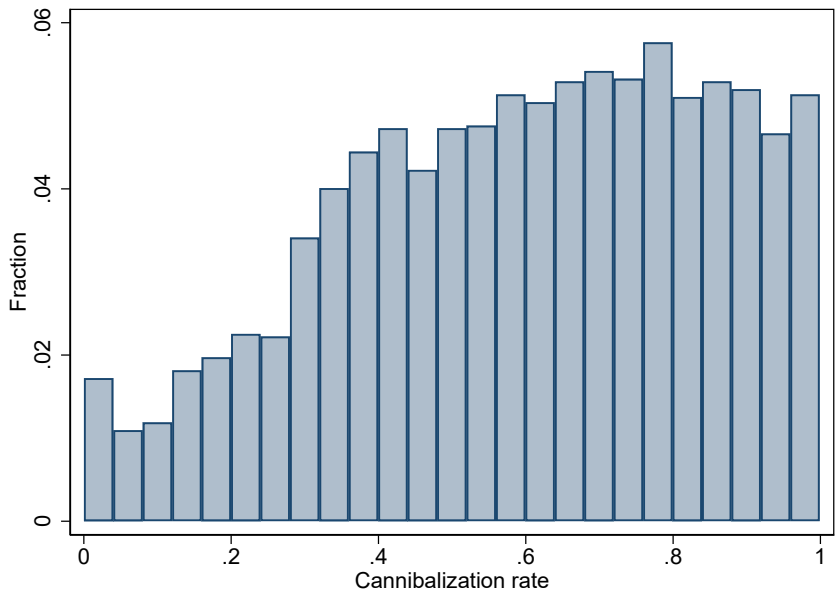


Figure 6: Distribution of cannibalization rate r_{mn}

Notes: $r_{mn} = 1 - S_{mn}^{\kappa}$ computed from observed prices and quantities. Excludes single-product brand–segments in the histogram; zeros can still occur when a model nearly monopolizes within brand-segment sales.

it were the only product its brand offered in that segment that are cannibalized by the brand's other products in market n. By construction, $r_{mn} = 0$ in single-product brand–segments because $S_{mn} = 1$ implies $\Delta_{mn}^{c} = 1$. We emphasize that r_{mn} is a partial-equilibrium measure: it holds fixed segment–country expenditure A_{sn}, demand shifters (ψ_{bsn}, ψ_{mn}), and prices at their observed/predicted levels, capturing within-brand substitution but not portfolio re-optimization, or general-equilibrium feedbacks across markets or plants. For illustration, a within-brand share of $S_{mn} = 0.20$ implies $r_{mn} \approx 1 - 0.2^{0.65} \approx 0.65$, i.e., the model sells about 35% of what it would sell if it were the only product in the brand-segment.

Figure 6 shows that cannibalization is sizable in multi-product brand–segments. The mean cannibalization rate is 0.55 (s.d. 0.29) and the median is 0.58. At the upper tail, 10% of model–markets have $r_{mn} > 0.91$, 5% exceed 0.96, and 1% exceed 0.99. The mean being below the median (skewness -0.40) is consistent with a thin low-cannibalization tail alongside a substantial mass at high cannibalization for fringe models. Economically, this pattern follows from within–brand substitutability being much stronger than across–brand substitutability ($\sigma^{M} \gg \sigma^{B}$). Higher σ^{M} (or lower σ^{B}) amplifies cannibalization and shifts mass to the right in the distribution of r_{mn}. With $\kappa \approx 0.65$, the upper tail corresponds to very small within–brand shares: $r_{mn} > 0.91 \Rightarrow S_{mn} < (1 - 0.91)^{1/0.65} \approx 0.025$, and $r_{mn} > 0.99 \Rightarrow S_{mn} \approx 0.00084$. Conversely, the median $r_{mn} \approx 0.58$ implies $S_{mn} \approx (1 - 0.58)^{1/0.65} \approx 0.26$, consistent with typical brand–segments offering about 3–4 models.

5.3 Fixed Cost Estimates

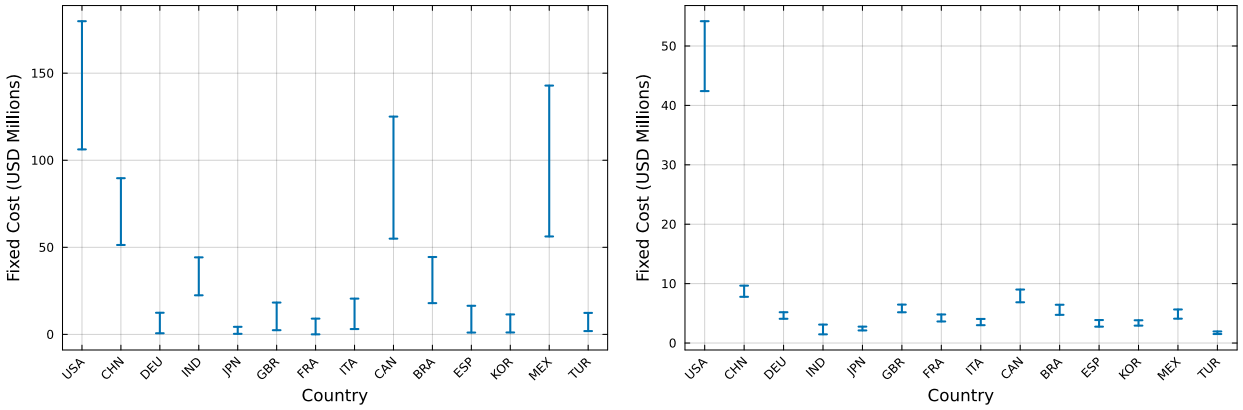
We implement our moment inequality procedure, discussed in Section 4.3, using a grid-search procedure in combination with Andrews and Soares (2010). Table 3 reports the limits of the 95% confidence sets for each of the estimated parameters.

Table 3: Fixed Costs Estimates

	95% Confidence Set Limits
Market Entry Fixed Cost	
γ_1 (Location)	[3.0, 5.25]
σ_e (Scale)	[4.75, 7.75]
Production Plant Fixed Costs	
γ_2 (Location)	[8.5, 12.5]
γ_3 (Free Plant)	[-2.75, -1.75]
σ_p (Std Shock)	[3.25, 5.5]
Observations	1111

Notes: Confidence sets computed using Andrews and Soares (2010). Units are in millions of USD. Variance matrix clustered at the brand-segment level.

Figure 7: Fixed Costs Conditional Producing/Selling for Brands with U.S. HQs



(a) Production Fixed Costs

(b) Selling Fixed Costs

Notes: These plots show bounds on the average fixed costs across 128 fixed cost draws and across products sold by brands with headquarters in the United States, conditional on production (Panel (a)) and selling (Panel (b)). The plots exclude zero fixed cost origins and destinations due to the free plant and free destination lottery. The bounds reflect both statistical uncertainty regarding the value of the fixed cost parameters as well as potential uncertainty about the solution due to computational infeasibility. We condition on (sure) production- locations for which the solution algorithm solves for a 1.

To interpret the fixed cost parameters, we solve the model using our algorithm under each point in our confidence set and under 128 draws of the fixed cost shocks, and report the mean fixed cost (across products and draws) conditional on producing or selling. Panel (a) of Figure 7 shows that mean production fixed costs (excluding free plants) do not exceed 200 million U.S. dollars per product-origin. Variation across countries in the mean fixed costs conditional on production

reflect selection. Brands with U.S. headquarters find it more profitable to sell in the United States, Canada, or Mexico, leading to production even at higher fixed costs. Instead, to produce in other countries, such as Japan or Turkey, American brands require low realizations of production fixed costs. Panel (b) illustrates a similar story for fixed selling costs, for which the conditional mean does not exceed 60 million U.S. dollars. In countries other than the United States, American brands do not incur more than 10 million dollars per product-destination. These estimates of realized fixed costs accord well with industry estimates of production and entry costs.

6 The Effect of U.S. Policy on Global Automobile Production

Accounting for 3% of global GDP, the automobile industry has historically been subject to important policy packages and intervention. In recent years, the automobile industry has especially captured policymakers' attention, with environmental policies seeking to incentivize the adoption of environmentally-friendly vehicles, and trade policies aiming to promote and protect domestic manufacturers. The goal of our analysis is to shed light on the effectiveness of these policies in achieving their intended goals, rather than to prescribe policies based on a normative analysis. For instance, a current topic in policy discussions is the extent to which tariffs on manufactured goods produced abroad can induce firms to relocate their production in the country imposing the tariff. Through the lens of our quantitative model, the answer to this question will depend crucially on the interaction between demand-side and supply-side factors that enter firms' plant location decision. If demand for such goods is large in the United States and the fixed cost of setting up a plant is relatively small, our model predicts that firms will profitably choose to re-locate production to the United States. If production fixed costs are very large and the product is not particularly demanded in the US market, this is less likely. A crucial factor affecting how profitable a product is in a country is cannibalization. Tariffs on foreignly-produced electric vehicles may fail to induce brands to re-locate the United States if such brands abstain from selling electric vehicles that cannibalize sales of their non-electric varieties.

We use our quantitative model to evaluate the effect of multiple policies enacted in the United States on global production in the car industry. First, we study the effects of policies on the production and sales decisions of a single brand, Audi. This exercise allows us showcase the non-trivial implications of firm and product-specific policies on global production and sales decisions of multi-product firms. Then, inspired by recent policies affecting electric vehicles, we study the effects of electric-vehicle-specific policies on global production and sales decisions of the average brand producing both electric and non-electric products.

6.1 Case Study: Audi in the Market for Wagons

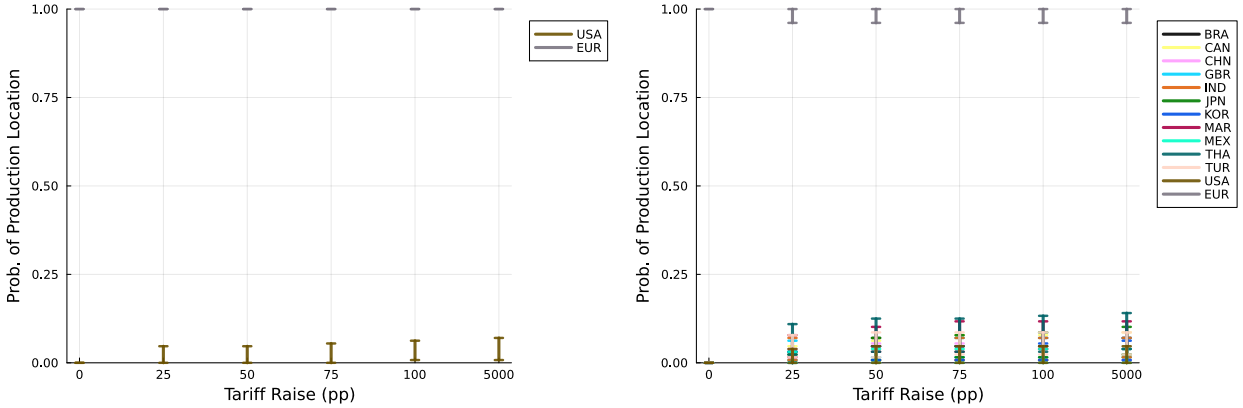
We study the market entry and plant location decisions of Audi within the wagon segment. In this segment, we observe Audi offering two products: the A4 and the A6. In the data, we observe that Audi produces both of these products in Germany only, and that only the A4 is sold in the

United States. Both products are sold both in the United Kingdom (UK) and in the European Union.

An advantage of focusing only on one brand-segment is that we can feasibly select fixed cost draws that perfectly match those features in the data, and compute counterfactual exercises conditional on such draws. We select 128 fixed cost draws that match production in Germany, the sale of both products in the UK and at least one large European country (Germany, France, Italy, or Spain), as well as the the availability and non-availability of the A4 and A6 in the United States, respectively. At these draws, we evaluate the effects of various policies on Audi's production and sales decisions for the A4.

6.1.1 Tariffs

Figure 8: Probability of Production Location - Tariffs on the A4



(a) U.S. Tariffs on ROW

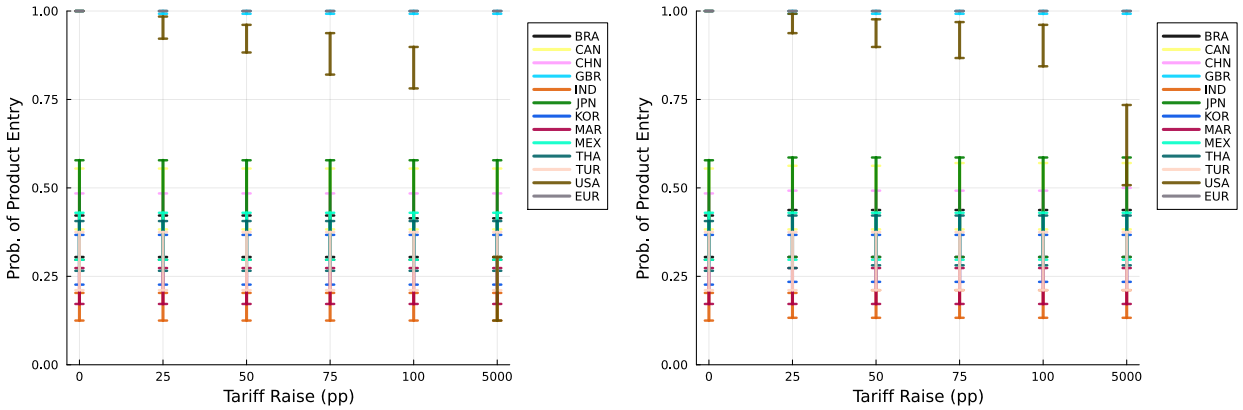
(b) U.S. Tariffs on the European Union

Notes: These plots show bounds on the probability that the Audi A4 is produced in different countries/regions at different U.S. tariffs that affect the A4 and not the A6. The bounds reflect both statistical uncertainty regarding the bounds on the fixed cost parameters and uncertainty about the solution arising from potential inability to solve the full brand problem. "EUR" denotes the European Union.

In this section, we explore the possibility of using tariffs to affect the global structure of production, and to attract manufacturing plants to the United States. To this end, we study the effect of increasing tariffs on the Audi A4, which is sold but not produced in the United States. A motivation for studying this policy is that Audi announced that its A4 product-line would go fully electric by 2026, and the United States announced tariffs on electric products (mostly produced in China).

Figure 8 shows how tariffs affecting only the A4 affect Audi's production location choices for this product. Panel (a) demonstrates that even a hypothetical 5,000 percentage-point increase in U.S. product-specific tariffs on the rest of the world leads to a small probability of production in the United States of $[0.8, 7]\%$. Intuitively, such a policy does not cause Audi to build A4 plants in any other countries. The reason for lack of production relocation to the United States is that fixed costs are high relative to the potential sales when sourcing the A4 locally within the United States.

Figure 9: Probability of Product Offering - Tariffs on the A4



(a) U.S. Tariffs on the ROW

(b) U.S. Tariffs on the European Union

Notes: These plots show the probability that the Audi A4 is sold in different countries/regions at different U.S. tariffs that affect the A4 and not the A6. The bounds reflect both statistical uncertainty regarding the bounds on the fixed cost parameters and uncertainty about the solution arising from potential inability to solve the full brand problem. "EUR" denotes the European Union.

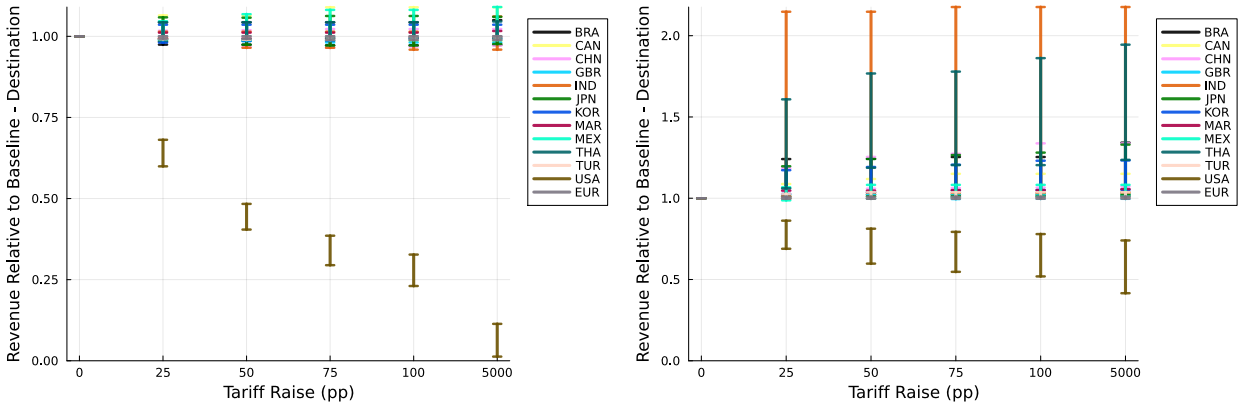
We find that at the $5,000$ percentage-point increase in tariffs, fixed costs would have to decline an average (across draws) of $[90.0, 97.3]\%$ to produce in the United States. A reason for this lack of relocation is that we are conditioning on not producing in the United States initially, which tends to select high fixed production costs for the United States.

In contrast, Panel (b) shows that tariffs only on the European Union lead to substantial production relocation, although not necessarily to the United States. Among the countries with highest probability of production relocation are Thailand and Morocco.[20] A $5,000$ percentage-point increase in tariffs to the European Union causes the probability of production in the United States to move from 0 (baseline) to $[0, 4.7]\%$. However, the model predicts a highly significant increase in the probability of opening of a new plant somewhere in the world. Indeed, the number of plants moves from 1 at baseline to $[1.375, 1.648]$ under a $5,000$ percentage point tariff increase.

Figures 9 demonstrates the effects of both types of tariffs on the probability that the Audi A4 is available in various countries. We find small changes in the bounds of the probability of entry in all countries except for the United States. There are two important takeaways from these plots. First, the probability that the A4 remains available in the United States declines monotonically with tariffs, with entry probabilities in the United States of [86.7, 96.9]% and [88.3, 96.1]% under a 50 percentage-point increase in ROW and European tariffs, respectively. At a 5,000 percentage-point increase, the bounds are instead [12.5, 30.5]% under ROW tariffs and [50.8, 73.4]% under European tariffs. Therefore, for "moderate" tariff raises, the effect on U.S. availability is similar under both types of tariffs, while the decline is much larger under ROW-tariffs. Second, the bounds on the number of destination countries in which the A4 is sold decreases from $[10.32, 11.70]$

[20] Thailand is the tenth largest automotive producer in the world and hosts Ford, BMW, Mercedes-Benz, Mazda, Mitsubishi, and Toyota, and Nissan plants, among others. Morocco has seen a rapid expansion of its car industry recently, with key investments by Renault-Nissan-Mitsubishi, PSA, and BYD.

Figure 10: Revenues Relative to Baseline at the Destination - Tariffs on the A4



(a) U.S. Tariffs on the ROW

(b) U.S. Tariffs on the European Union

Note: These plots show the revenues from the Audi A4 in different countries/regions at different U.S. tariffs that affect the A4 and not the A6, relative to the revenues earned at the baseline case with a 0 pp tariff raise. The bounds reflect both statistical uncertainty regarding the bounds on the fixed cost parameters and uncertainty about the solution arising from potential inability to solve the full brand problem. "EUR" denotes the European Union.

to $[9.617, 10.836]$ under ROW-tariffs, but only decline to $[10.05, 11.47]$ under EU tariffs. Thus, large tariffs on the European Union cause an increase in the availability of the A4 in some country other than the United States, which almost fully compensates the decline in U.S. availability. This shows that large tariffs on the European Union cause production and sales to relocate to other markets, not necessarily to serve the United States.

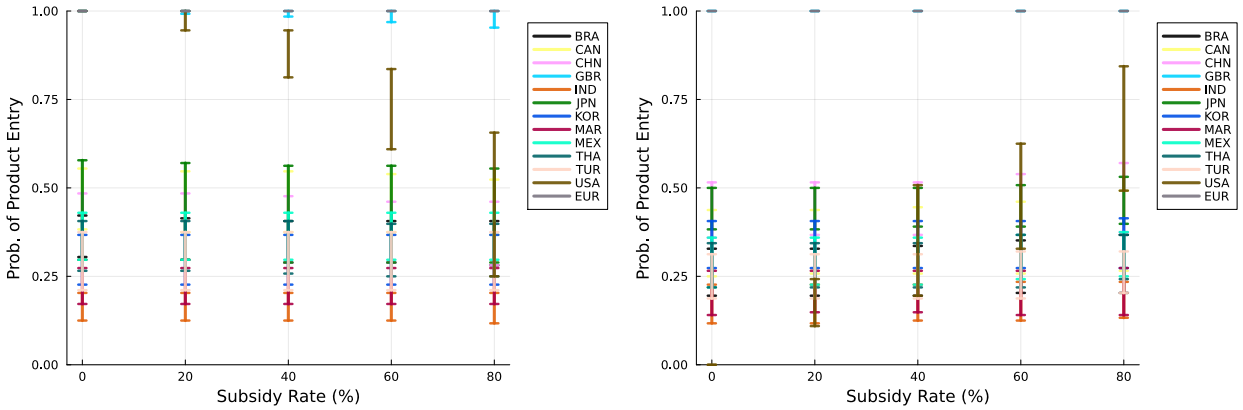
This point is made clear in Figure 10. Tariffs on the ROW cause a progressive decline in A4 revenues in the United States. This is also true under tariffs on the European Union, though even with a 5,000 percentage-point increase, revenues decline by not more than 60%, to $[41.6, 74.0]$% of baseline revenues. These plots clearly show the potential for large increases in A4 revenues in third countries in response to tariffs on the European Union, with for instance a rise of $[23.6, 94.5]$% in Thailand.

To summarize, tariffs on the A4 have potentially large consequences on Audi's production and market-entry strategy. However, the way the tariffs are implemented matters. Tariffs on the ROW lead to small relocation effects. Under most draws, Audi prefers to withdraw from the U.S. market than to establish a costly A4 plant in the United States. Instead, tariffs on the European Union, lead to a substantial reorganization of production, and due to interdependencies with product entry choices, on sales in different markets. This counterfactual exercise illustrates the importance of interdependencies between production and sales when analyzing the global effects of tariffs.

6.1.2 Demand-Side Subsidy on the A6

Policymakers have also proposed demand-side subsidies to incentivize the purchase and offerings of certain product categories, either due to environmental motivations (e.g., electric-vehicle

Figure 11: Consumer Subsidy on the A6



(a) A4: Product Entry Probability

(b) A6: Product Entry Probability

Notes: These plots show bounds on the probability that the Audi A4 and Audi A6 are sold in different countries/regions at different U.S. consumer subsidies on the A6. The bounds reflect both statistical uncertainty regarding the bounds on the fixed cost parameters and uncertainty about the solution arising from potential inability to solve the full brand problem. "EUR" denotes the European Union.

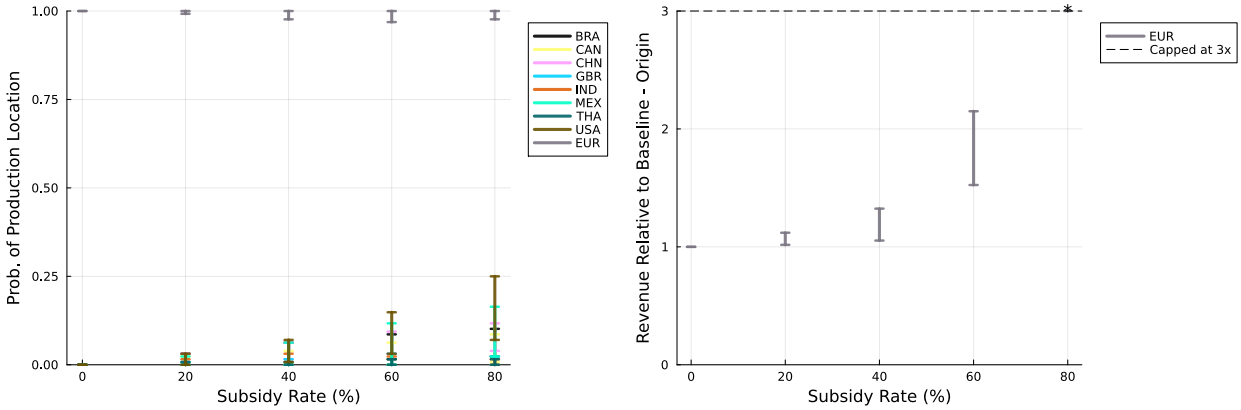
subsidies), or to spur domestic manufacturing (e.g., Buy American). Our model allows to speak to the effectiveness of demand-side interventions in attracting the production and offering of subsidized products to the United States. Given that the A6 is not offered in the United States in the data, we study the effect of a demand-side intervention that shifts the appeal ψ_{mn} of the Audi A6 in the United States to $\psi'_{mn} = \psi_{mn}/(1-s)$.[21]

Figure 11 shows the effects of the demand-side intervention on the availability of both the A4 and the A6 throughout the world. Panel (b) shows that already a 20% consumer subsidy on the A4 increases the probability that it is available in the United States from zero to $[10.9, 24.2]\%$. At an 80% subsidy, the bounds increase to $[49.2, 84.4]\%$. Panel (a) shows that cannibalization effects also matter. Subsidizing the A6 can cause the exit of the A4. When the A6 receives an 80% subsidy, the probability that the A4 is still offered shifts to $[25.0, 65.6]\%$. The bounds on product entry probabilities also shift in countries other than the United States, though less significantly. In Mexico, the probability that the A6 is offered shifts from a baseline interval of $[21.9, 35.9]\%$ to $[25.0, 37.5]\%$. Indeed, while the consumer subsidy incentivizes sales in the United States, it affects incentives to sell the product abroad through endogenous production location choices. If the change in production location does not lead to a substantial decrease in marginal costs of serving third countries, Audi's incentives to sell the A6 there are not substantially affected. Qualitatively, both cannibalization forces and economies of scale jointly determine the availability of products globally in response to a demand-side intervention.

Figure 12 shows that the subsidy increases the probability of A6 production in multiple countries, particularly in the United States, Mexico, and China. However, large consumer subsidies are needed to substantially increase the probability that a new plant is opened. At a 60% consumer

[21] In a single-country model, this policy is equivalent to a marginal-cost subsidy. With trade, this policy is distinct because ψ_{mn} varies at the destination-level, while marginal costs vary at the origin-level.

Figure 12: Consumer Subsidy on the A6



(a) A6: Production Location Probability

(b) A6: Revenues Relative to Baseline

Notes: Panel (a) shows bounds on the probability that the Audi A6 is produced in different countries/regions at different U.S. consumer subsidies on the A6. Panel (b) shows bounds on the A6 revenues coming from production in the European Union relative to baseline (pre-policy) revenues. The bounds reflect both statistical uncertainty regarding the bounds on the fixed cost parameters and uncertainty about the solution arising from potential inability to solve the full brand problem. "EUR" denotes the European Union.

subsidy, the probability of production is $[3.1, 14.8]\%$ in the United States, and $[0, 11.7]\%$ in Mexico. At an 80% subsidy, these probabilities are $[7, 25]\%$ and $[2.3, 16.4]\%$, respectively. While under some fixed cost draws a new plant is opened to serve the higher-demand U.S. market, we find it is likely that Audi prefers to abstain from building a new plant and to continue sourcing the Audi A6 to the United States from Europe. Inded, Panel (b) shows that revenues relative to baseline in the European Union increase by $[52.4, 215]\%$ under a 60% subsidy and more than triple under an 80% subsidy. Computing bounds on origin-level shares, we find that even at an 80% subsidy, the European Union accounts $[30.4, 55.9]\%$ of A6 production value, while the United States only for $[12.6, 31.9]\%$. Therefore, while production relocation is somewhat likely, it is also highly plausible for Europe to continue exporting the A6 to the U.S. market. Intuitively, if fixed production costs are high enough and the United States not productive or large enough to serve U.S. consumers, Audi will not find it profitable to build a U.S. plant.

6.2 Alternative Policies

In Appendix E, we discuss the effects of U.S. subsidies that reduce the unit cost of producing the A4 in the United States. This is an additional policy intervention that can be used to incentivize production relocation to the United States. We find that large subsidies of over 50% are needed for the United States to significantly increase the probability that the Audi A4 is produced in the United States. However, at large reductions in U.S. costs, the United States becomes a viable and profitable export platform for this product. Large marginal cost subsidies increase the probability that the A4 is introduced globally, reduce the probability that the A6 is sold in different countries due to cannibalization effects, and increase perceived U.S. origin-level revenues, as well as global

destination-level revenues for the A4. Due to production relocation on the intensive margin, Europe produces less of the A4 as the U.S. increases the marginal cost subsidy, but at large subsidies it also produces less of the A6 due to cannibalization forces leading to reduced demand for the A6.

6.3 Average Effects of U.S. Policies Favoring Electric Vehicles

Electric vehicle policies have become commonplace worldwide. We study the effects of such policies implemented in the United States on the global structure of automobile production. Through the lens of our partial equilibrium model, we report the impact of U.S. tariffs and consumer subsidies by averaging across all firms with at least one electric product in their portfolio.[22] We evaluate the effects of various levels of subsidies and tariffs on average firm-level outcomes, both on the extensive and the intensive margin. To solve for the optimal production and offering choices of firms under different policies, we implement our solution method at 128 draws of fixed costs and under each of the points in the confidence set of fixed cost parameters $\boldsymbol{\gamma}$. Let $h_{bs}(\boldsymbol{\gamma}, \tau, r)$ denote an outcome for brand b in segment s under parameter vector $\boldsymbol{\gamma}$, policy τ, and realization of fixed cost shocks r. Our solution method yields bounds on production location and product entry decisions $(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$, and as such, bounds on outcomes h that are weakly monotone in such choices. To obtain bounds on average brand-segment responses, we first average across the 128 realizations of shocks to obtain averaged lower and upper bounds, $\underline{h}^a_{bs}(\boldsymbol{\gamma}, \tau)$ and $\overline{h}^a_{bs}(\boldsymbol{\gamma}, \tau)$, respectively. Then, we average across all treated brand-segments to obtain bounds for the averages, $\underline{H}^a(\boldsymbol{\gamma}, \tau)$ and $\overline{H}^a(\boldsymbol{\gamma}, \tau)$. Finally, we obtain bounds on the counterfactual by minimizing and maximizing $\underline{H}^a(\boldsymbol{\gamma}, \tau)$ and $\overline{H}^a(\boldsymbol{\gamma}, \tau)$ over $\boldsymbol{\gamma}$ in the confidence set CS, respectively. This procedure yields bounds on levels $\underline{H}(\tau) = \min_{\boldsymbol{\gamma} \in CS} \underline{H}^a(\boldsymbol{\gamma}, \tau)$ and $\overline{H}(\tau) = \max_{\boldsymbol{\gamma} \in CS} \overline{H}^a(\boldsymbol{\gamma}, \tau)$.[23] The reported bounds reflect both the fact that it is computationally infeasible in some cases to find the full solution (the branching algorithm applied after the first-step algorithm does not converge after a maximal number of iterations), and statistical uncertainty regarding the value of $\boldsymbol{\gamma}$.

6.3.1 U.S. Tariffs on Electric Vehicles Produced Abroad

Figures 13-14 show the effects of different levels of U.S. tariffs on electric vehicles. In both figures, Panel (a) shows the consequences of tariffs on the rest of the world, while Panel (b) shows

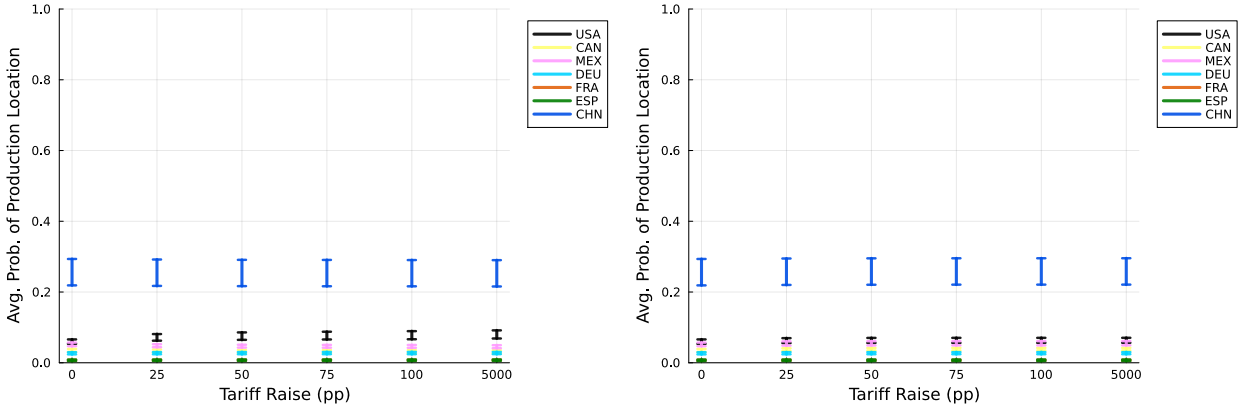
[22] We define a product as if electric trims of the product account for more than 50% of its global sales.

[23] A drawback of performing the counterfactual exercises in levels is that the parameter vector $\boldsymbol{\gamma}$ minimizing (maximizing) a given counterfactual outcome could vary across different policy levels, even though there is only one true parameter that generates the data. An alternative approach that addresses this concern is to report counterfactual changes, which compute

$$\begin{aligned}\Delta\overline{H}(\boldsymbol{\gamma}, \tau) := \max\{&\overline{H}^a(\boldsymbol{\gamma}, \tau) - \underline{H}^a(\boldsymbol{\gamma}, \tau_0), \overline{H}^a(\boldsymbol{\gamma}, \tau) - \overline{H}^a(\boldsymbol{\gamma}, \tau_0),\\ &\underline{H}^a(\boldsymbol{\gamma}, \tau) - \overline{H}^a(\boldsymbol{\gamma}, \tau_0), \underline{H}^a(\boldsymbol{\gamma}, \tau) - \underline{H}^a(\boldsymbol{\gamma}, \tau_0)\}\\ \Delta\underline{H}(\boldsymbol{\gamma}, \tau) := \min\{&\overline{H}^a(\boldsymbol{\gamma}, \tau) - \underline{H}^a(\boldsymbol{\gamma}, \tau_0), \overline{H}^a(\boldsymbol{\gamma}, \tau) - \overline{H}^a(\boldsymbol{\gamma}, \tau_0),\\ &\underline{H}^a(\boldsymbol{\gamma}, \tau) - \overline{H}^a(\boldsymbol{\gamma}, \tau_0), \underline{H}^a(\boldsymbol{\gamma}, \tau) - \underline{H}^a(\boldsymbol{\gamma}, \tau_0)\},\end{aligned}$$

where τ_0 denotes the case without policies in the United States (or anywhere). Then, one would report $\Delta\overline{H}(\tau) = \max_{\boldsymbol{\gamma} \in CS} \Delta\overline{H}(\boldsymbol{\gamma}, \tau)$ and $\Delta\underline{H}(\tau) = \min_{\boldsymbol{\gamma} \in CS} \Delta\underline{H}(\boldsymbol{\gamma}, \tau)$.

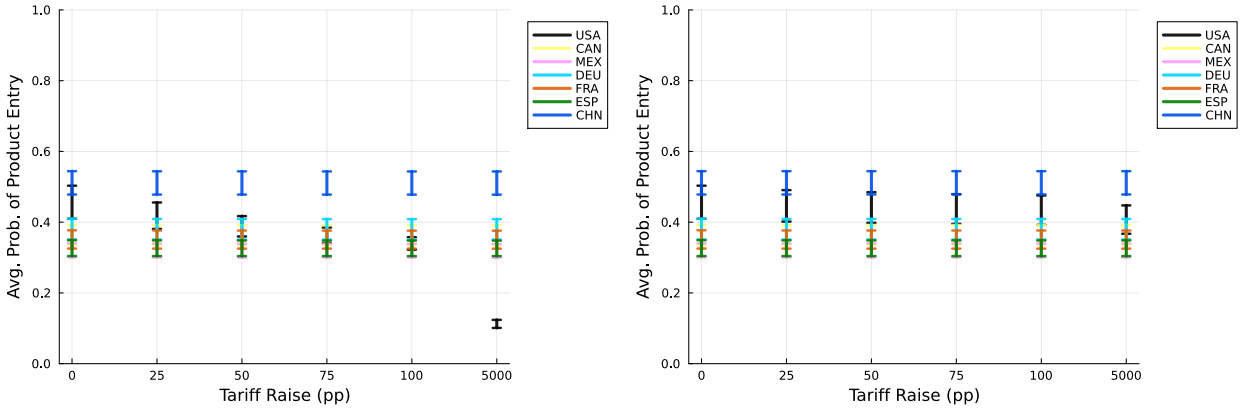
Figure 13: Probability of Production Location of Electric Vehicles - Tariffs



(a) Tariffs on the ROW (b) Tariffs on the European Union

Note: These plots display bounds on the average share of electric products that are produced in different countries by brand-segments with at least one electric product in their portfolio. Panel (a) shows the bounds at different levels of U.S. tariffs on the rest of the world on electric products. Panel (b) displays the effect of tariffs imposed only on electric products produced in the European Union. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

Figure 14: Probability of Product Offering of Electric Vehicles - Tariffs



(a) Tariffs on the ROW (b) Tariffs on the European Union

Note: These plots display bounds on the average share of electric products that are offered in different countries by brand-segments with at least one electric product in their portfolio. Panel (a) shows the bounds at different levels of U.S. tariffs on the rest of the world on electric products. Panel (b) displays the effect of tariffs imposed only on electric products produced in the European Union. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

the effects of tariffs imposed only on the European Union. Figure 13 shows that, consistent with our Audi case-study, production relocation to the United States is strongest under a tariff on the rest of the world. Panel (a) shows that the bounds on the average share of electric products that brands produce in the United States increase monotonically from $[5.4, 6.6]\%$ at baseline to $[6.9, 9.2]\%$ at a 5,000 percentage-point increase in ROW-tariffs. Thus, even at such an extreme level of tariffs,

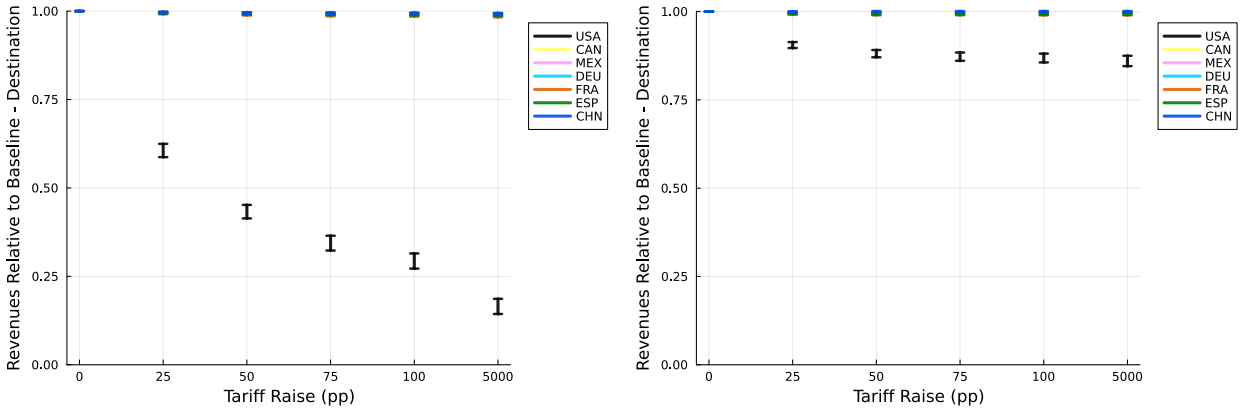
the average probability that a given brand produces an electric product in the United States is at most 9.2%. Production relocation to the United States is therefore not large in magnitude. Panel (b) demonstrates that if tariffs were imposed only on the European Union, the bounds move only up to $[5.6, 7.1]$%, a very small change compared to baseline.

As seen in Figure 14, tariffs have very different effects on product availability depending on their nature. When imposed on the rest of the world, electric-vehicle tariffs lead to a decrease in the probability that electric product is offered. This is consistent with little relocation of production to the United States. At the extreme with a 5,000 percentage-point tariff raise, the average probability that brand sells an electric vehicle in the United States approximately coincides with the probability that it produces it there. This is not the case with tariffs on the European Union only. In this case, even extreme tariffs lead to small effects on electric-product availability, with U.S. bounds shifting from $[41.0, 50.3]$% at baseline to $[36.7, 44.7]$% under the highest tariff raise. This small decline is explained by two forces. As with Audi, tariffs on the European Union cause some plants to open in other countries. The average number of electric-vehicle plants per product increases in this case from $[1.213, 1.277]$ at baseline to $[1.244, 1.310]$. However, this effect is smaller than in the Audi case, because unlike Audi wagons, many products are produced in multiple plants at baseline. Therefore, for many brands, tariffs on the European Union affect brands' sourcing decision, but do not lead to substantial effects on the extensive margin.

Figure 15 shows the consequences of U.S. tariffs on electric products for average revenues earned per electric product in different destination countries. As seen in Panel (a), ROW tariffs cause a sharp decline in electric vehicle sales in the United States, which fall to $[41.4, 45.2]$% and $[14.4, 18.6]$% of baseline values at a 50 and 5,000 percentage-point tariff raise, respectively. Note that that the limited production relocation to the United States does cause a decline in electric vehicle sales in Europe, with revenues falling by up to $[0.7, 1.4]$% in Germany or $[0.7, 1.8]$% in France. Global electric-vehicle sales fall by much less under European-Union tariffs.

In Figure 23 of Appendix E, we report the effects on origin-level electric-vehicle revenues and find the largest declines under ROW tariffs in Canada and Mexico. For example, in Canada, electric-vehicle plants shrink their production by $[23.3, 34.4]$% at a 25 percentage-point increase in tariffs and $[40.0, 55.0]$% under a $5,000$ percentage-point raise. These effects are instead mute under European Union tariffs. We also report the effects of tariffs on non-electric products. The main takeaways are that (i) non-electric product offerings increase slightly in the United States under ROW tariffs, though only from $[26.1, 30.5]$% to at most $[28.0, 34.0]$%, (ii) exports of non-electric vehicles tend increase globally due to these tariffs due to cannibalization effects, (iii) purchases of non-electric products rise primarily in the United States, with a much larger increase under ROW tariffs than under E.U. tariffs (e.g., $[37.6, 77.5]$% versus $[3.7, 6.1]$% at a 5,000 percentage-point tariff raise). These results suggest that while cannibalization forces have important repercussions for products not directly affected by tariffs, the effects are quantitatively much more significant on the intensive rather than the extensive margin.

Figure 15: Revenues Relative to Baseline at Destination - Tariffs



(a) Tariffs on the ROW

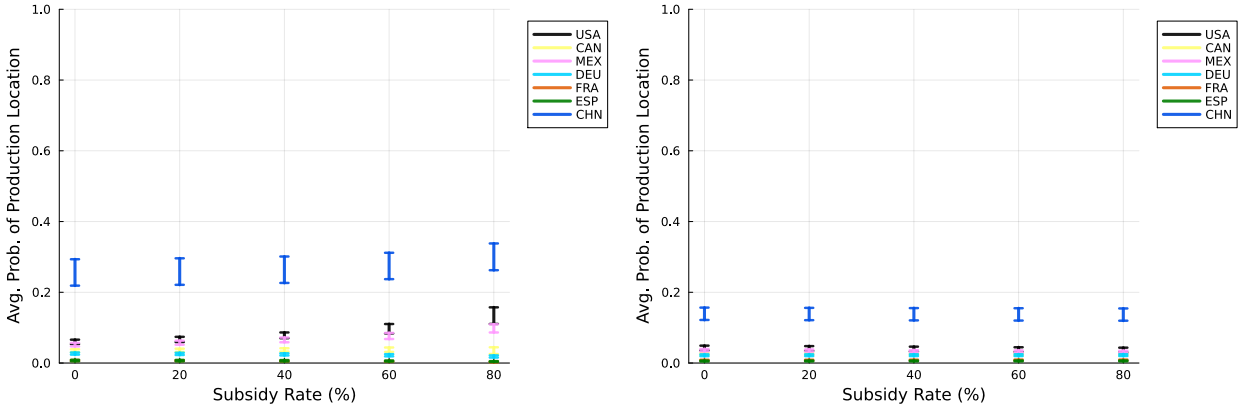
(b) Tariffs on the European Union

Note: These plots display bounds on the average revenue per electric product across brand-segments with at least one electric product in their portfolio relative to baseline in different destination countries. Panel (a) shows the bounds at different levels of U.S. tariffs on the rest of the world on electric products. Panel (b) displays the effect of tariffs imposed only on electric products produced in the European Union. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

6.3.2 Consumer Subsidies on Electric Vehicles in the United States

Figure 16 illustrates how consumer subsidies on electric products in the United States affects brands' production location choices. As seen in Panel (a), larger consumer subsidies cause brands to open additional electric-product in the United States, but also in other countries like Mexico, Canada, or China. Therefore, countries other than the United States are often better positioned

Figure 16: Probability of Production Location - Consumer Subsidy

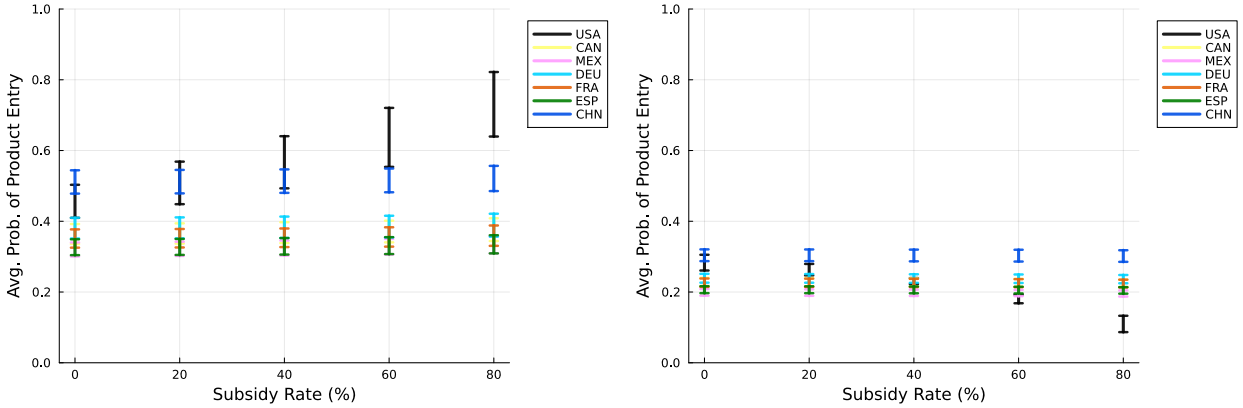


(a) Electric

(b) Non-Electric

Notes: These plots show the average lower and upper bound share of electric (panel (a)) and non-electric (panel (b)) products produced by brands in several countries/regions under different US electric-vehicle consumer subsidy rates. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

Figure 17: Probability of Product Offering - Consumer Subsidy



(a) Electric (b) Non-Electric

Notes: These plots show the average lower and upper bound share of electric (panel (a)) and non-electric (panel (b)) products offered by brands in several countries/regions under different US electric-vehicle consumer subsidy rates. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

to serve the U.S. market with certain products. As consumer subsidies rise, the average share of electric products produced in the United States shifts from $[5.4, 6.6]\%$ at baseline to $[11.1, 15.7]\%$ at an 80% consumer subsidy. In Mexico, the shift is from $[4.7, 5.7]\%$ to $[8.6, 11.0]\%$. The consumer subsidies cause production relocation. We find that while the probability rises of production rises very significantly in the United States or Mexico, the average number of plants per electric product rises only from $[1.213, 1.277]$ to $[1.240, 1.388]$. For non-electric products, we find that the number of plants tends to decrease globally, though to a lesser extent. In the United States, the bounds for non-electric products shift from $[3.6, 4.9]\%$ at baseline to $[3.0, 4.3]\%$ at an 80% consumer subsidy. In Figure 28 of Appendix E, we find the corresponding effects on plant sales. Thus, consumer subsidies can effectively attract production to origin countries that are well-positioned to serve the U.S. market, which for many products we find is not the United States.

The opening up of new plants a result of US demand-side interventions has implications for global consumers. American consumers see increasing availability of electric products as the consumer subsidy rises, with intervals $[41.0, 50.3]\%$ offered at baseline, and $[49.3, 64.0]\%$ and $[64.0, 82.2]\%$ offered at 40% and 60% subsidies, respectively. We also find small increases in availability in other countries. In Mexico, the bounds shift from $[30.1, 34.0]\%$ to $[30.9, 35.8]\%$ at an 80% subsidy in the United States. Overall, we find that firms sell the average electric product in $[14.4, 15.9]$ countries at baseline, and $[15.7, 17.6]$. Therefore, we find quantitatively large local and small cross-country effects on the extensive margin.

We find a similar and reverse story for non-electric products in Figure 17. Cannibalization forces lead to substantial exit of non-subsidized products, though mostly in the United States. We find that the non-electric probability bounds in the United States shift from $[26.1, 30.5]\%$ at baseline, to $[22.3, 23.9]\%$ and $[8.7, 13.3]\%$ under 40% and 80% consumer subsidies, respectively.

While we find small cross-country effects on the extensive margin, this does not mean that these effects are quantitatively small or economically unimportant. On the intensive margin, U.S. consumer subsidies on electric products have a direct impact on U.S. electric vehicle purchases, as seen in Panel (a) Figure 18. A 20% consumer subsidy raises U.S. purchases by more than twofold, and subsidies greater than 40% by more than threefold. In the United States, such subsidies have the effect of sharply reducing purchases of non-subsidized products. Already a 20% consumer subsidy decreases non-electric vehicle sales by $[20.7, 28.4]$%. While smaller quantitatively, we also find that brands' production relocation in response to the U.S. consumer subsidy affects purchases of non-electric products worldwide. In Mexico, electric purchases increase by $[1.5, 3.1]$% under a U.S. 20% subsidy, and by $[17.4, 27.6]$% under an 80% subsidy. At such subsidy rates, Mexican purchases of non-electric products decline by $[0.6, 1.5]$% and $[4.5, 7.2]$%, respectively. Thus, endogenous production relocation in response to a local electric consumer subsidy has quantitatively significant effects on electric purchases abroad through production relocation and on non-electric purchases through cannibalization forces.

6.4 Alternative Policies

In Appendix E, we also study the effect of U.S. marginal-cost subsidies favoring production of electric vehicles. We find that high marginal cost subsidies cause production of electric vehicles to relocate to the United States, increased sales of electric vehicles globally, and reduced sales of non-electric vehicles due to cannibalization effects.

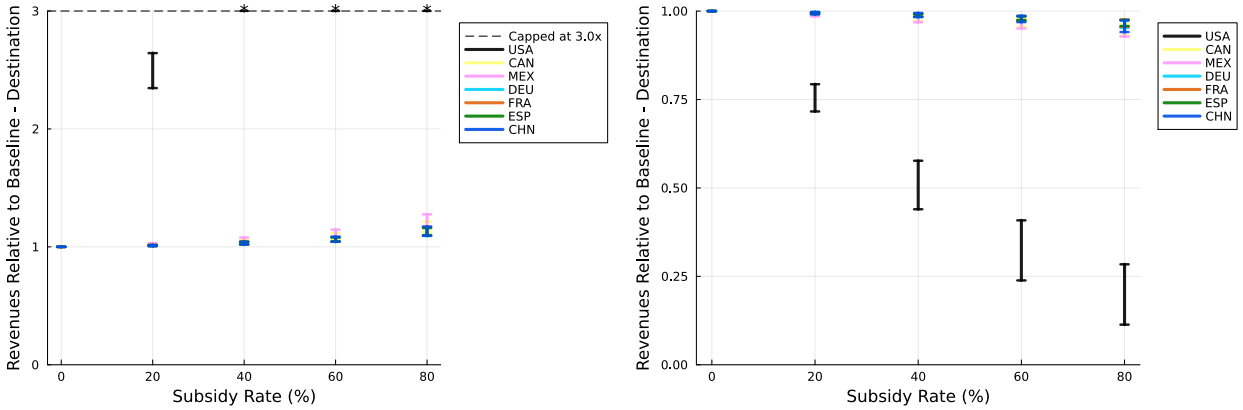
6.5 Why Is Production Relocation Limited in Response to U.S. Tariffs?

In sections 6.1.1 and 6.3.1, we showed that even large tariffs cause limited production relocation to the United States. To understand this result, we conduct additional counterfactual experiments to study the effect of tariffs in a world with alternative U.S. demand and cost fundamentals. We plot the results in Appendix G.

For production relocation to the United States to take place, by definition there must be no U.S. production before tariffs are implemented, and there must be production after tariffs are implemented. Fundamental changes to demand and costs in the United States affect both pre-treatment and post-treatment production probabilities.

Both for Audi and for the average firm, we find that production relocation to the United States would be higher if tariffs were implemented in a world with greater U.S. demand for the treated products. We also find that lower fixed costs of producing in the United States potentially cause higher rates of plant relocation. At some parameters in the confidence set, we find that non-monotonicities occur. For instance, Figure 39 in Appendix G shows that production relocation is higher under a 60% reduction in U.S. production fixed costs rather than an 80% reduction, due to higher baseline production probabilities in the latter case. We also find that even relatively modest U.S. marginal costs reductions increase baseline production probabilities, and therefore limit the scope for production relocation. All in all, we find quantitatively that the set of products for which

Figure 18: Revenues Relative to Baseline at Destination - Consumer Subsidy



(a) Electric (b) Non-Electric

Note: These plots display bounds on the average revenue per product across brand-segments with at least one electric product in their portfolio relative to baseline in different destination countries. Panel (a) shows the bounds at different levels of U.S. consumer subsidies on electric products. Panel (b) displays the bounds at different levels of U.S. consumer subsidies on non-electric products. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

tariffs induce relocation is small – products for which production in the U.S. is profitable are likely to already be produced at baseline, while products for which production is unprofitable at baseline are unlikely to be relocated due to insufficient demand relative to the costs of setting up a plant.

7 Alternative Models

In this section, we discuss alternative modelling assumptions, and whether or not our approach could be used to study such cases. First, we consider an alternative model with brand-country-level fixed production/selling costs. Second, we consider whether it is possible to study models with product-origin-destination-level fixed selling costs. Finally, we consider a version of the model where no fixed selling cost is incurred for a product if the model is produced in that country.

7.1 Economies of Scope

In our main specification, we specify fixed production and selling costs to vary at the product-country level. In reality, brands may first pay a fixed cost for entering a market, and subsequently pay additional fixed costs per product. For instance, fixed production costs for product m in country n might take the alternative form $\tilde{F}^p_{mn} = \gamma_4 \mathbb{1}\{\sum_{m' \neq m} D_{m'n} = 0\} + F^p_{mn}$, such that an additional fixed cost γ_4 is paid whenever m is the first product produced in country n. Under this specification, there is now a source of complementarities across production locations for a given product. Overall, it may be the case that production locations for the same product are substitutes, due to minimum-cost sourcing, or complements, due to economies of scope. In Appendix B.3.1, we show that our methodology can be extended to allow for payoff functions that are additively

separable in components, and where for each component the sign of the interdependence across any pair of choices is known. In this case, production locations for the same products are substitutes under variable profits (net of fixed selling costs), but complements under fixed production costs, which enter additively in the payoff function.

7.2 Fixed Selling Costs Varying at the Origin-Destination Level

Suppose that fixed selling costs varied across models, destination, and origin countries, i.e., F^e_{mno}. In this case, instead of sourcing a product to destination n from the minimum marginal cost location, brands would optimally source products from the maximal profit location. For single-product brands, this would pose no additional conceptual challenge, other than computing the optimal sourcing location according to $\max_o\{\pi_{no}(\boldsymbol{D}) - F^e_{no}\}$, instead of $\min_o\{c_{no}\}$. However, a complication arises in the case of multi-product brands. Without minimum-cost sourcing, there is now a potential complementarity between producing a product m, and selling a product m' in a given destination n. Indeed, if producing product m in a low fixed cost location now leads to sourcing model m from a higher marginal cost origin, this would *increase* the marginal value of selling product m' due to reduced cannibalization. In this case, because the interdependence between producing m and selling m' enter non-additively through the price index, it is not possible to leverage the extension from Appendix B.3.

7.3 Zero Fixed Selling Costs in Producing Locations

Consider the following specification for fixed selling costs: $\hat{F}^e_{mn} = (1 - D_{mn})F^e_{mn}$. That is, if product m is produced in country n, then selling product m in country n happens at a zero fixed cost, independently of whether m is sourced to n from n. Then, total profits for brand b in segment s can be written as,

$$\begin{aligned}\hat{\Pi}_{bs}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) &= \sum_{m,n} I_{mn}(\pi_{mn}(\boldsymbol{I}^{bs}_n, \boldsymbol{D}^{bs}) - F^e_{mn}(1 - D_{mn})) - \sum_{m,n} D_{mn}F^p_{mn} \\ &= \underbrace{\sum_{m,n} I_{mn}(\pi_{mn}(\boldsymbol{I}^{bs}_n, \boldsymbol{D}^{bs}) - F^e_{mn}) - \sum_{m,n} D_{mn}F^p_{mn}}_{\Pi_{bs}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})} + \sum_{m,n} I_{mn}D_{mn}F^e_{mn}.\end{aligned}$$

Therefore, due to constant-signed cross-partials under Π_{bs}, and because the term $\sum_{m,n} I_{mn}D_{mn}F^e_{mn}$ does not change the sign of the cross-partials, we can solve this alternative model.[24] Notice that overall, production locations remain substitutes under this payoff function, but there is an additional source of complementarity between producing and selling product m in the same location.

[24] While we could obtain bounds on the solution using our baseline mapping, it complicates using the extension from equation (24), as it introduces a potential source of (non-additive) complementarity between production locations.

8 Conclusion

Our analysis highlights how multinational, multiproduct firms make interconnected production and market entry decisions in response to policy changes. By modeling these choices within a unified framework, we show that subsidies and tariffs have distinct effects on global production. Demand-side EV subsidies in the U.S. favor production in or near the United States, and EV market entry, while marginal cost subsidies shift production toward the United States, both reducing the availability of non-electric alternatives worldwide. Tariffs, in contrast, have limited effects on global production structures due to high fixed costs, inducing selective production relocation but reducing product variety.

More broadly, our framework illustrates how national policies can reshape firms' global strategies, creating spillover effects across markets. The presence of fixed costs and economies of scale generates a proximity-concentration tradeoff, meaning that decisions in one country influence firms' entire global operations. These findings underscore the importance of considering cross-market effects when analyzing trade and industrial policies in industries dominated by multinational firms.

This paper provides a novel method that permits solving combinatorial discrete choice single-agent and static models featuring simultaneous complementarities and substitutabilities. We plan to extend our method to allow for strategic interactions across agents and for dynamic interdependent choices in future work.

References

Alfaro-Ureña, Alonso, Juanma Castro-Vincenzi, Sebastian Fanelli, and Eduardo Morales, "Firm Dynamics with Interdependent Markets," *mimeo*, January 2024.

Allcott, Hunt, Reigner Kane, Max Maydanchik, Joseph Shapiro, and Felix Tintelnot, "The Effects of "Buy American": Electric Vehicles and the Inflation Reduction Act," October 2024. Mimeograph.

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

Antràs, Pol, Teresa Fort, and Felix Tintelnot, "The Margins of Global Sourcing: Theory and Evidence from U.S. Firms," *American Economic Review*, 2017, *107* (9), 2514–2564.

Arkolakis, Costas, Fabian Eckert, and Rowan Shi, "Combinatorial Discrete Choice: A Quantitative Model of Multinational Location Decisions*," *mimeo*, November 2023.

_ , Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple, "Innovation and Production in the Global Economy," *American Economic Review*, August 2018, *108* (8), 2128–73.

_ , Sharat Ganapati, and Marc-Andreas Muendler, "The Extensive Margin of Exporting Products: A Firm-Level Analysis," *American Economic Journal: Macroeconomics*, October 2021, *13* (4), 182–245.

Barwick, Panle Jia, Hyuk-Soo Kwon, and Shanjun Li, "Attribute-based Subsidies and Market Power: an Application to Electric Vehicles," Working Paper 32264, National Bureau of Economic Research March 2024.

Bernard, Andrew B., J. Bradford Jensen, Stephen J. Redding, and Peter K. Schott, "Global Firms," *Journal of Economic Literature*, June 2018, *56* (2), 565–619.

Berry, Steve, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63*, 841–890.

Castro-Vincenzi, Juanma, "Climate Hazards and Resilience in the Global Car Industry," *mimeo*, February 2024.

Coşar, Kerem A., Paul L. E. Grieco, Shengyu Li, and Felix Tintelnot, "What Drives Home Market Advantage?," *Journal of International Economics*, 2016, *110*, 135–150.

Dhingra, Swati, "Trading Away Wide Brands for Cheap Brands," *American Economic Review*, October 2013, *103* (6), 2554–2584.

Eckel, Carsten and J. Peter Neary, "Multi-Product Firms and Flexible Manufacturing in the Global Economy," *The Review of Economic Studies*, 01 2010, *77* (1), 188–217.

Fan, Ying and Chenyu Yang, "Competition, Product Proliferation, and Welfare: A Study of the US Smartphone Market," *American Economic Journal: Microeconomics*, May 2020, *12* (2), 99–134.

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

Grieco, Paul L E, Charles Murry, and Ali Yurukoglu, "The Evolution of Market Power in the U.S. Automobile Industry*," *The Quarterly Journal of Economics*, 09 2023, *139* (2), 1201–1253.

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

_ , _ , Marc Melitz, and Chenying Yang, "Industrial policies for multi-stage production: The battle for battery-powered vehicles," Manuscript 2025.

Helpman, Elhanan and Benjamin Niswonger, "Dynamics of Large Multinationals," *Journal of the European Economic Association*, 01 2023, *21* (5), 1994–2042.

Holmes, Thomas J., "The Diffusion of Wal-Mart and Economies of Density," *Econometrica*, 2011, *79* (1), 253–302.

Hottman, Colin J., Stephen J. Redding, and David E. Weinstein, "Quantifying the Sources of Firm Heterogeneity," *The Quarterly Journal of Economics*, 2016, *131* (3), 1291–1364.

ITC, "Market Access Map," 2022.

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

Mayer, Thierry and Soledad Zignago, "Notes on CEPII's Distances Measures: the GeoDist Database," *CEPII Working Paper*, 2011, *2011-25.*

_ , Marc J. Melitz, and Gianmarco I. P. Ottaviano, "Market Size, Competition, and the Product Mix of Exporters," *American Economic Review*, February 2014, *104* (2), 495–536.

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

Ramondo, Natalia and Andrés Rodríguez-Clare, "Trade, Multinational Production, and the Gains from Openness," *Journal of Political Economy*, 2013, *121* (2), 273–322.

Sabal, Alejandro, "Product Entry in the Global Automobile Industry," *mimeo*, 2025.

Tarski, Alfred, "A lattice-theoretical fixpoint theorem and its applications.," *Pacific Journal of Mathematics*, 1955, *5* (2), 285 – 309.

Tintelnot, Felix, "Global Production with Export Platforms," *The Quarterly Journal of Economics*, 2017, *132* (1), 157–209.

Appendices

A Model Properties . . . 47

A.1 Interdependence Properties under Π_{bs} . . . 47

A.2 Interdependence Properties under $V_{b,s,m}$ . . . 50

B Extensions of the First-Step Algorithm . . . 54

B.1 Extension 1: Value-Function Mapping and Composition . . . 54

B.2 Generalization of Extension 1: Composition of Mappings . . . 55

B.3 Extension 2: Additive Payoff Function . . . 58

B.4 Applying Extension 1 to the Value Function . . . 60

C Method Proofs . . . 61

C.1 Proof of Theorems 2 and 3 . . . 61

C.2 Proof of Proposition 6 . . . 62

C.3 Properties of the Branching Algorithm . . . 62

C.4 Simple Example: Branching Algorithm . . . 63

D Practical Implementation of Algorithms 1 and 3 . . . 64

E Plots of Counterfactual Exercises . . . 67

E.1 Audi Wagons: U.S. Electic-Vehicle Marginal Cost Subsidy . . . 67

E.2 Average Counterfactual Exercises . . . 70

F Realized Fixed Costs . . . 75

G Tariffs under Counterfactual Worlds . . . 76

G.1 Audi - Confidence Intervals . . . 76

G.2 Audi - Single Parameter Vector . . . 78

G.3 Average Firm – Confidence Intervals . . . 79

A Model Properties

In this section, we derive the key model properties that guarantee that Assumption 1 is satisfied under different specifications of the payoff function. We begin by proving the interdependence structure across the binary choices that enter Π_{bs}. We then prove the interdependence properties of value functions $V_{b,s}$ and $V_{b,s,m}$.

A.1 Interdependence Properties under Π_{bs}

In this subsection, we prove the properties of Π_{bs} stated in section 3.1.1. Total profits for brand b in segment s are,

$$\Pi_{bs}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) = \sum_{n=1}^{N} \sum_{m=1}^{M_{bs}} \{I_{mn}(\pi_{mn}(\boldsymbol{I}_n^{bs}, \boldsymbol{D}^{bs}) - F_{mn}^e) - D_{mn}F_{mn}^p\} \tag{A1}$$

$$= \sum_{n=1}^{N} A_{bsn} P_{bsn}(\boldsymbol{I}_n^{bs}, \boldsymbol{D}^{bs})^{1-\sigma^B} - \sum_{m,n} (I_{mn}F_{mn}^e + D_{mn}F_{mn}^p), \tag{A2}$$

where

$$P_{bsn}(\boldsymbol{I}_n^{bs}, \boldsymbol{D}^{bs}) = \Big(\sum_{m=1}^{M_{bs}} I_{mn} \big(P_{mn}(\boldsymbol{D}_m)/\psi_{mn}\big)^{1-\sigma^M} \Big)^{\frac{1}{1-\sigma^M}}. \tag{A3}$$

A.1.1 Property 1: D_{mo} and I_{mn} are complements for any m and any $o, n \in \mathcal{N}$

To prove this property, consider the marginal value of I_{mn},

$$\Delta_{I_{mn}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) = A_{bsn}[P_{bsn}(\boldsymbol{I}_n^{bs,m\to 1}, \boldsymbol{D}^{bs})^{1-\sigma^B} - P_{bsn}(\boldsymbol{I}_n^{bs,m\to 0}, \boldsymbol{D}^{bs})^{1-\sigma^B}] - F_{mn}^e. \tag{A4}$$

To prove Property 1, we need to show that $\Delta_{I_{mn}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$ is increasing in $\boldsymbol{D}_m$. Consider $\hat{\boldsymbol{D}}_m \geq \boldsymbol{D}_m$. From the perspective of country n sales, there are two possibilities. More production location decisions for product m could either: (i) leave the minimum-cost origin unchanged, and therefore $P_{mn} = \hat{P}_{mn}$, or (ii) decrease the cost of supplying m in n, causing $\hat{P}_{mn} < P_{mn}$. Thus, the difference-in-differences yields,

$$\Delta_{I_{mn}}(\boldsymbol{I}^{bs}, \hat{\boldsymbol{D}}^{bs}) - \Delta_{I_{mn}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) = A_{bsn}[P_{bsn}(\boldsymbol{I}_n^{bs,m\to 1}, \hat{\boldsymbol{D}}^{bs})^{1-\sigma^B} - P_{bsn}(\boldsymbol{I}_n^{bs,m\to 0}, \hat{\boldsymbol{D}}^{bs})^{1-\sigma^B} \tag{A5}$$

$$- P_{bsn}(\boldsymbol{I}_n^{bs,m\to 1}, \boldsymbol{D}^{bs})^{1-\sigma^B} + P_{bsn}(\boldsymbol{I}_n^{bs,m\to 0}, \boldsymbol{D}^{bs})^{1-\sigma^B}] \tag{A6}$$

$$= A_{bsn}[P_{bsn}(\boldsymbol{I}_n^{bs,m\to 1}, \hat{\boldsymbol{D}}^{bs})^{1-\sigma^B} - P_{bsn}(\boldsymbol{I}_n^{bs,m\to 1}, \boldsymbol{D}^{bs})^{1-\sigma^B}] \tag{A7}$$

where the second equality follows from the fact that whenever $I_{mn} = 0$, the brand-segment price index is equal under $\hat{\boldsymbol{D}}^{bs}$ and under $\boldsymbol{D}^{bs}$. Now, under case (i) with $P_{mn} = \hat{P}_{mn}$, such a difference is clearly zero, while under case (ii) the difference is clearly positive. Thus,

$$\Delta_{I_{mn}}(\boldsymbol{I}^{bs}, \hat{\boldsymbol{D}}^{bs}) - \Delta_{I_{mn}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) \geq 0,$$

which proves complementarity.

We can also prove complementarity in the other direction. Consider the marginal value of D_{mo},

$$\Delta_{D_{mo}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) = \sum_n A_{bsn}[P_{bsn}(\boldsymbol{I}^{bs}_n, \boldsymbol{D}^{b,mo\to 1})^{1-\sigma^B} - P_{bsn}(\boldsymbol{I}^{bs}_n, \boldsymbol{D}^{b,mo\to 0})^{1-\sigma^B}] - F^p_{mo}. \tag{A8}$$

Consider the differences-in-differences,

$$\Delta_{D_{mo}}(\boldsymbol{I}^{bs,mk\to 1}, \boldsymbol{D}^{bs}) - \Delta_{D_{mo}}(\boldsymbol{I}^{bs,mk\to 0}, \boldsymbol{D}^{bs}). \tag{A9}$$

Clearly, for countries $n \neq k$, the terms in the differences-in-differences cancel out, since the marginal contribution of D_{mo} to country n profits is independent of the value of I_{mk}. Therefore,

$$\Delta_{D_{mo}}(\boldsymbol{I}^{bs,mk\to 1}, \boldsymbol{D}^{bs}) - \Delta_{D_{mo}}(\boldsymbol{I}^{bs,mk\to 0}, \boldsymbol{D}^{bs}) \tag{A10}$$

$$= A_{bsk}[P_{bsk}(\boldsymbol{I}^{bs,m\to 1}_k, \boldsymbol{D}^{bs,mo\to 1})^{1-\sigma^B} - P_{bsk}(\boldsymbol{I}^{bs,m\to 1}_k, \boldsymbol{D}^{bs,mo\to 0})^{1-\sigma^B}] \geq 0, \tag{A11}$$

where the equality follows from the fact that the terms with no sales in country n of model m cancel out and the inequality follows from potentially lower prices from $D_{mo} = 1$ and $\sigma^B > 1$.

A.1.2 Property 2: D_{mo} and I_{kn} are substitutes for $m \neq k$ and for any $o, n \in \mathcal{N}$

To prove Property 2, consider the marginal value of I_{kn},

$$\Delta_{I_{kn}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) = A_{bsn}[P_{bsn}(\boldsymbol{I}^{bs,k\to 1}_n, \boldsymbol{D}^{bs})^{1-\sigma^B} - P_{bsn}(\boldsymbol{I}^{bs,k\to 0}_n, \boldsymbol{D}^{bs})^{1-\sigma^B}] - F^e_{kn}. \tag{A12}$$

To prove Property 2, we need to show that $\Delta_{I_{kn}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$ is decreasing in $\boldsymbol{D}_m$. Consider $\hat{\boldsymbol{D}}_m \geq \boldsymbol{D}_m$. Clearly, if $I_{mn} = 0$, changes in production location choices for product m will not affect the brand-segment price index. Thus, consider the case in which $I_{mn} = 1$. In this case, more production locations for product m can only lead to declines in the price of product m in country n. Differentiating with respect to P_{mn}, we obtain,

$$\frac{\partial \Delta_{I_{kn}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})}{\partial P_{mn}} = \frac{A_{bsn}}{\psi_{mn}}\left(\frac{P_{mn}}{\psi_{mn}}\right)^{-\sigma^M}(1-\sigma^B)[P_{bsn}(\boldsymbol{I}^{bs,k\to 1}_n, \boldsymbol{D}^{bs})^{\sigma^M-\sigma^B} - P_{bsn}(\boldsymbol{I}^{bs,k\to 0}_n, \boldsymbol{D}^{bs})^{\sigma^M-\sigma^B}]$$

$$\geq 0, \tag{A13}$$

where the inequality follows from $\sigma^B > 1$ and the fact that the price index is decreasing in varieties offered and there is cannibalization, $\sigma^M > \sigma^B$. Therefore, increases in prices of model m in country n increase the marginal value of offering product k in country n. Or, increases in production locations for model m, which weakly decrease prices for product m in country n, decrease the marginal value of offering product k country n. It follows that Property 2 holds.

A.1.3 Property 3: I_{kn} and I_{mn} are substitutes in country n for $k \neq m$

The proof for this property mimics the proof for Property 2. Indeed, we proved that $\Delta_{I_{kn}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$ is increasing in P_{mn}. Offering product m in country n is equivalent to decreasing its price from ∞ to a

finite value. Therefore, $\Delta_{I_{kn}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$ is also decreasing in I_{mn}.

A.1.4 Property 4: I_{mn} and I_{kt} are independent for $n \neq t$ for all m, k with $m \neq k$

This property follows by inspection of equation (A4) defining $\Delta_{I_{mn}}$, which clearly does not depend on the sales decisions in markets other than n.

A.1.5 Property 5: D_{mo} and D_{mt} are substitutes for $o \neq t$

To prove Property 5, consider the marginal value of producing product m in country o,

$$\Delta_{D_{mo}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) = \sum_n A_{bsn}[P_{bsn}(\boldsymbol{I}^{bs}_n, \boldsymbol{D}^{bs,mo\to 1})^{1-\sigma^B} - P_{bsn}(\boldsymbol{I}^{bs}_n, \boldsymbol{D}^{bs,mo\to 0})^{1-\sigma^B}] - F^p_{mo}. \tag{A14}$$

We introduce notation and write,

$$\Delta_{D_{mo}}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) = \sum_n \Delta^o_{bsn}(\boldsymbol{D}_m) - F^p_{mo}. \tag{A15}$$

Note that $\Delta^o_{bsn}(\boldsymbol{D}_m) \geq 0$ for all n and all $\boldsymbol{D}_m$. We abuse notation and write Δ^o_{bsn} as a function of only $\boldsymbol{D}_m$ even though it also depends on production location of other products and of country-n offerings. We prove Property 5 by showing that $\Delta^o_{bsn}(\overline{\boldsymbol{D}}_m) \leq \Delta^o_{bsn}(\underline{\boldsymbol{D}}_m)$ for all n. This shows that the marginal value of producing m in o is smaller when m is produced in more locations. Consider two cases:

1. $\Delta^o_{bsn}(\overline{\boldsymbol{D}}_m) = 0$. In this case, producing product m in country o causes no change in variable profits in country n because m is not sourced to n from o. Thus, necessarily $0 = \Delta^o_{bsn}(\overline{\boldsymbol{D}}_m) \leq \Delta^o_{bsn}(\underline{\boldsymbol{D}}_m)$.

2. $\Delta^o_{bsn}(\overline{\boldsymbol{D}}_m) > 0$. If this is the case, then necessarily country o is the minimum cost location to source product m to country n, under $\overline{\boldsymbol{D}}_m$. The fact that $\underline{\boldsymbol{D}}_m \leq \overline{\boldsymbol{D}}_m$ implies that $\Delta^o_n(\underline{\boldsymbol{D}}_m) \geq 0$ – if o is the lowest cost production location under a superset of possible sourcing locations, it must also be the lowest cost location under a subset. However, the change in cost (and therefore in price) must be weakly smaller under $\overline{\boldsymbol{D}}_m$ than under $\underline{\boldsymbol{D}}_m$, which implies that $\Delta^o_n(\overline{\boldsymbol{D}}_m) \leq \Delta^o_n(\underline{\boldsymbol{D}}_m)$.

We have therefore proven Property 5.

A.1.6 Property 6: D_{mo} and D_{kt} are substitutes for $m \neq k$

To prove this property, it suffices to prove that $\frac{\partial}{\partial P_{kn}}\Delta^o_n(\boldsymbol{D}_m) \geq 0$. This follows from the fact that, as seen in equation (A14), the marginal value of producing m in country o consists of the change in variable profits across all countries arising from this decision minus the production fixed cost. If some other product k is produced is now produced in some country t, this could potentially reduce the price of product k in any country, affecting the marginal contribution of D_{mo} through the brand-segment price index.

We have,

$$\frac{\partial}{\partial P_{kn}}\Delta_n^o(\boldsymbol{D}_m) = \frac{A_{bsn}}{\psi_{kn}}\left(\frac{P_{kn}}{\psi_{kn}}\right)^{-\sigma^M}(1-\sigma^B)[P_{bsn}(\boldsymbol{I}_n^{bs},\boldsymbol{D}^{b,mo\to 1})^{\sigma^M-\sigma^B} - P_{bsn}(\boldsymbol{I}_n^{bs},\boldsymbol{D}^{b,mo\to 0})^{\sigma^M-\sigma^B}] \quad \text{(A16)}$$

$$\geq 0, \quad \text{(A17)}$$

since $\sigma^B > 1$, $\sigma^M > \sigma^B$ and the brand-segment price index is weakly decreasing in the set of production locations. We have therefore proven Property 6.

A.2 Interdependence Properties under $V_{b,s,m}$

We now prove the interdependence across production locations under the mapping,

$$V_{b,s,m}(\boldsymbol{D}_m;\boldsymbol{I}_{-m},\boldsymbol{D}_{-m}) := \max_{\boldsymbol{I}_m}\sum_{n=1}^{N} I_{mn}\Delta\pi_n^m(\boldsymbol{D}_m;\boldsymbol{I}_{-m},\boldsymbol{D}_{-m}) - \sum_{n=1}^{N} D_{mn}F_{mn}^p, \quad \text{(A18)}$$

introduced in equation (24). We first consider the single-product firm case and prove Proposition 5. We then prove general properties that hold in the general case with multiproduct firms.

A.2.1 Proof of Proposition 5

Consider the mapping,

$$V_{b,s}(\boldsymbol{D}_m) := \max_{\boldsymbol{I}_m}\sum_{n=1}^{N} I_{mn}\Delta\pi_n^m(\boldsymbol{D}_m) - \sum_{n=1}^{N} D_{mn}F_{mn}^p. \quad \text{(A19)}$$

The aim is to prove that D_{mo} and $D_{mo'}$ are substitutes under $V_{b,s}$. It suffices to prove that, holding fixed all production location decisions except for two, we have that,

$$V_{b,s}((1,1),\boldsymbol{D}_m) - V_{b,s}((0,1),\boldsymbol{D}_m) \leq V_{b,s}((1,0),\boldsymbol{D}_m) - V_{b,s}((0,0),\boldsymbol{D}_m). \quad \text{(A20)}$$

While b is single-product in segment s, we label the product as m (this helps with notation when we extend the result).

Introducing notation, let $I_{mn}^{(i,j)}$ denote the optimal product entry decision rule under the production location decision implied by $(i,j) \in \{(1,1),(1,0),(0,1),(0,0)\}$. We abuse notation and write $((i,j),\boldsymbol{D}_m)$ as the argument of the function, which should be interpreted as all production locations other than (i,j) are fixed at the values corresponding to $\boldsymbol{D}_m$.

The first step of the proof is to realize that,

$$I_{mn}^{(1,1)} = I_{mn}^{(1,0)} + I_{mn}^{(0,1)} - I_{mn}^{(1,0)}I_{mn}^{(0,1)}. \quad \text{(A21)}$$

Equation (A21) says that if product m is sold in country n under production vector $(1,1)$, it must be sold either under production vector $(1,0)$ or under production vector $(0,1)$ (or both). This property arises due

to single-origin sourcing – sourcing from two locations cannot be beneficial. In addition, we must have that,

$$I_{mn}^{(1,1)} \geq \max\{I_{mn}^{(1,0)}, I_{mn}^{(0,1)}\}, \tag{A22}$$

and

$$I_{mn}^{(0,0)} \leq \min\{I_{mn}^{(1,0)}, I_{mn}^{(0,1)}\} \tag{A23}$$

Inequality (A22) says that if m is sold in n under either $(1,0)$ or $(0,1)$, then it is necessarily sold under $(1,1)$. Inequality (A23) says that if m is sold in n under $(0,0)$, it must be sold either under $(1,0)$ or $(0,1)$. Combining inequalities (A21)-(A23), we obtain,

$$I_{mn}^{(1,1)} - I_{mn}^{(0,1)} = I_{mn}^{(1,0)} - I_{mn}^{(1,0)} I_{mn}^{(0,1)} \leq I_{mn}^{(1,0)} - I_{mn}^{(0,0)}. \tag{A24}$$

We are now ready to prove the proposition. Expanding inequality (A20), the goal is to prove that

$$\begin{aligned} &\sum_n \left(I_{mn}^{(1,1)} - I_{mn}^{(0,1)}\right) \Delta\pi_n^m(1,1) + \sum_n I_{mn}^{(0,1)} \left[\Delta\pi_n^m(1,1) - \Delta\pi_n^m(0,1)\right] \\ &\qquad\qquad \leq \\ &\sum_n \left(I_{mn}^{(1,0)} - I_{mn}^{(0,0)}\right) \Delta\pi_n^m(1,0) + \sum_n I_{mn}^{(0,0)} \left[\Delta\pi_n^m(1,0) - \Delta\pi_n^m(0,0)\right] \end{aligned}$$

or equivalently,

$$\begin{aligned} &\sum_n \left(I_{mn}^{(1,1)} - I_{mn}^{(0,1)}\right) \Delta\pi_n^m(1,1) + \sum_n (I_{mn}^{(0,1)} - I_{mn}^{(0,0)}) \left[\Delta\pi_n^m(1,1) - \Delta\pi_n^m(0,1)\right] \\ &\qquad\qquad + \sum_n I_{mn}^{(0,0)} [\Delta\pi_n^m(1,1) - \Delta\pi_n^m(0,1)] \\ &\qquad\qquad \leq \\ &\sum_n \left(I_{mn}^{(1,0)} - I_{mn}^{(0,0)}\right) \Delta\pi_n^m(1,0) + \sum_n I_{mn}^{(0,0)} \left[\Delta\pi_n^m(1,0) - \Delta\pi_n^m(0,0)\right]. \end{aligned} \tag{A25}$$

Due to Property 5 of the model (proven in subsection A.1.5),

$$\Delta\pi_n^m(1,1) - \Delta\pi_n^m(0,1) \leq \Delta\pi_n^m(1,0) - \Delta\pi_n^m(0,0), \tag{A26}$$

which follows from the fact that production locations are substitutes. Combining inequalities (A25)-(A26), it follows that a sufficient condition to prove is,

$$\sum_n (I_{mn}^{(1,1)} - I_{mn}^{(0,0)}) \Delta\pi_n^m(1,1) \leq \sum_n (I_{mn}^{(0,1)} - I_{mn}^{(0,0)}) \Delta\pi_n^m(0,1) + \sum_n (I_{mn}^{(1,0)} - I_{mn}^{(0,0)}) \Delta\pi_n^m(1,0). \tag{A27}$$

Note that the LHS of inequality (A27) sums across the variable profits in all countries for which the brand enters only when the product is produced in either of the two countries. The RHS sums profits over all destination countries for which sourcing occurs from the second country and over all countries for which

sourcing occurs from the first country. It must be that the LHS is equal to the RHS, which concludes the proof.

A.2.2 Corollary of Proposition 5: $V_{b,s,m}$ exhibits decreasing differences in $\boldsymbol{D}_m$

A corollary of Proposition 5 is that the function $V_{b,s,m}$, defined in equation (A18), exhibits decreasing differences in $\boldsymbol{D}_m$. The proof follows immediately from the arguments in section A.2.1. Note that all of the arguments apply even in the case with multiproduct firms, given that they do not depend on the values of $(\boldsymbol{I}_{-m}, \boldsymbol{D}_{-m})$.

A.2.3 Argmax of $V_{b,s,m}(\boldsymbol{D}_m; \boldsymbol{I}^*_{-m}, \boldsymbol{D}^*_{-m})$ is $\boldsymbol{D}^*_m$

In this subsection, we prove the following proposition.

Proposition 2 *The optimal production location decisions for product m, $\boldsymbol{D}^*_m$, that solve the program $\max_{\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}} \Pi_{bs}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$ are also the argmax of $V_{b,s,m}(\boldsymbol{D}_m; \boldsymbol{I}^*_{-m}, \boldsymbol{D}^*_{-m})$ is $\boldsymbol{D}^*_m$.*

Proof. The problem $\max_{\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}} \Pi_{bs}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$ has a unique solution almost surely. Note that this problem can trivially be written as,

$$\max_{\boldsymbol{D}_m} \max_{\boldsymbol{I}_m} \max_{\boldsymbol{I}_{-m}, \boldsymbol{D}_{-m}} \Pi_{bs}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs}) = \sum_{n=1}^{N} \sum_{m=1}^{M_{bs}} \{I_{mn}(\pi_{mn}(\boldsymbol{I}^{bs}_n, \boldsymbol{D}^{bs}) - F^e_{mn}) - D_{mn} F^p_{mn}\}. \tag{A28}$$

Let $(\boldsymbol{I}_{-m}(\boldsymbol{I}_m, \boldsymbol{D}_m), \boldsymbol{D}_{-m}(\boldsymbol{I}_m, \boldsymbol{D}_m))$ denote the argmax of the innermost maximization problem. We can then write the problem as,

$$\max_{\boldsymbol{D}_m} \max_{\boldsymbol{I}_m} \sum_{n=1}^{N} \{I_{mn}(\Delta\pi^m_n(\boldsymbol{D}_m; \boldsymbol{I}_{-m}(\boldsymbol{I}_m, \boldsymbol{D}_m), \boldsymbol{D}_{-m}(\boldsymbol{I}_m, \boldsymbol{D}_m))) - D_{mn} F^p_{mn}\}. \tag{A29}$$

This step follows because conditional on production location choices, market entry decisions are independent across countries. As such, given any $\boldsymbol{D}_m$, the inner loop of problem (A29) simply chooses optimal entry decisions in each country independently, accounting for how such entry decisions affect overall profits. Now, program (A29) has the form of $V_{b,s,m}$, which proves Proposition 2. ■

A.2.4 $V_{b,s,m}$ exhibits decreasing differences in $\boldsymbol{D}_{-m}$ and $\boldsymbol{I}_{-m}$

To prove this result, consider the difference

$$\Delta V^o_{b,s,m}(\boldsymbol{D}_m; \boldsymbol{I}_{-m}, \boldsymbol{D}_{-m}) = V_{b,s,m}(\boldsymbol{D}^{o\to 1}_m; \boldsymbol{I}_{-m}, \boldsymbol{D}_{-m}) - V_{b,s,m}(\boldsymbol{D}^{o\to 0}_m; \boldsymbol{I}_{-m}, \boldsymbol{D}_{-m}) - F^p_{mo}.$$

The goal is to prove that $\Delta V^o_{b,s,m}(\boldsymbol{D}_m; \boldsymbol{I}_{-m}, \boldsymbol{D}_{-m})$ is decreasing in $\boldsymbol{I}_{-m}$ and $\boldsymbol{D}_{-m}$. We are interested in establishing that $\Delta V^o_{b,s,m}(\boldsymbol{D}_m; \underline{\boldsymbol{I}}_{-m}, \boldsymbol{D}_{-m}) - \Delta V^o_{b,s,m}(\boldsymbol{D}_m; \overline{\boldsymbol{I}}_{-m}, \boldsymbol{D}_{-m}) \geq 0$ and $\Delta V^o_{b,s,m}(\boldsymbol{D}_m; \boldsymbol{I}_{-m}, \underline{\boldsymbol{D}}_{-m}) - \Delta V^o_{b,s,m}(\boldsymbol{D}_m; \boldsymbol{I}_{-m}, \overline{\boldsymbol{D}}_{-m}) \geq 0$ for $\overline{\boldsymbol{D}}_{-m} \geq \underline{\boldsymbol{D}}_{-m}$ and $\overline{\boldsymbol{I}}_{-m} \geq \underline{\boldsymbol{I}}_{-m}$. Using $\underline{V}_{b,s,m}$ to denote dependence on $\underline{\boldsymbol{I}}_{-m}$ or $\underline{\boldsymbol{D}}_{-m}$, and $\overline{V}_{b,s,m}$ to denote dependence on $\overline{\boldsymbol{I}}_{-m}$ or $\overline{\boldsymbol{D}}_{-m}$, the difference-in-differences can be

expanded as,

$$\begin{aligned}\Delta\underline{V}^{o}_{b,s,m}(\boldsymbol{D}_m) - \Delta\overline{V}^{o}_{b,s,m}(\boldsymbol{D}_m) = &\sum_n \underline{I}^1_{mn}[\Delta\underline{\pi}^m_n(\boldsymbol{D}^{o\to1}_m) - \Delta\overline{\pi}^m_n(\boldsymbol{D}^{o\to1}_m)] \\ &+ \sum_n (\underline{I}^1_{mn} - \overline{I}^1_{mn})\Delta\overline{\pi}^m_n(\boldsymbol{D}^{o\to1}_m) \\ &- \sum_n \underline{I}^0_{mn}[\Delta\underline{\pi}^m_n(\boldsymbol{D}^{o\to0}_m) - \Delta\overline{\pi}^m_n(\boldsymbol{D}^{o\to0}_m)] \\ &- \sum_n (\underline{I}^0_{mn} - \overline{I}^0_{mn})\Delta\overline{\pi}^m_n(\boldsymbol{D}^{o\to0}_m),\end{aligned} \tag{A30}$$

where under and over-lines are used to denote dependence on $\underline{\boldsymbol{D}}_{-m}(\underline{\boldsymbol{I}}_{-m})$ and $\overline{\boldsymbol{D}}_{-m}(\overline{\boldsymbol{I}}_{-m})$, respectively and superscripts on I_{mn} denote the optimal product-market entry decision under $D_{mo} = j$. Partition the set of countries as follows:

$$\begin{aligned}G_m &= \{n : \underline{I}^1_{mn} - \overline{I}^1_{mn} > \underline{I}^0_{mn} - \overline{I}^0_{mn}\} \\ E_m &= \{n : \underline{I}^1_{mn} - \overline{I}^1_{mn} = \underline{I}^0_{mn} - \overline{I}^0_{mn}\} \\ L_m &= \{n : \underline{I}^1_{mn} - \overline{I}^1_{mn} < \underline{I}^0_{mn} - \overline{I}^0_{mn}\}.\end{aligned}$$

We can partition equation (A30) and write,

$$\Delta\underline{V}^{o}_{b,s,m}(\boldsymbol{D}_m) - \Delta\overline{V}^{o}_{b,s,m}(\boldsymbol{D}_m) = \Sigma_{G_m} + \Sigma_{E_m} + \Sigma_{L_m} \tag{A31}$$

where Σ_{G_m} denotes the components of the sum in equation (A30) that sum over G_m, Σ_{E_m} denotes the components that sum over E_m, and Σ_{L_m} denotes the components that sum over L_m. It suffices to show that each of these components is weakly positive. Starting with G_m, note that if $n \in G_m$, it must be that $\underline{I}^1_{mn} = 1$ and $\overline{I}^1_{mn} = \underline{I}^0_{mn} = \overline{I}^0_{mn} = 0$, which implies,

$$\Sigma_{G_m} = \sum_{n\in G_m} \Delta\underline{\pi}^m_n(\boldsymbol{D}^{o\to1}_m) \geq 0, \tag{A32}$$

where the inequality follows from $n \in G_m \implies \underline{I}^1_{mn} = 1$. Focusing now on Σ_{E_m}, note that for all $n \in E_m$, either $\underline{I}^1_{mn} = \overline{I}^1_{mn} = \underline{I}^0_{mn} = \overline{I}^0_{mn} = 0$, in which case the difference-in-differences is $\Sigma_{E_m} = 0$, or $\underline{I}^1_{mn} = \overline{I}^1_{mn} = \underline{I}^0_{mn} = \overline{I}^0_{mn} = 1$, in which case,

$$\Sigma_{E_m} = \sum_{n\in E_m : \underline{I}^1_{mn}=\overline{I}^1_{mn}=\underline{I}^0_{mn}=\overline{I}^0_{mn}=1} [(\Delta\underline{\pi}^m_n(\boldsymbol{D}^{o\to1}_m) - \Delta\overline{\pi}^m_n(\boldsymbol{D}^{o\to1}_m)) - (\Delta\underline{\pi}^m_n(\boldsymbol{D}^{o\to0}_m) - \Delta\overline{\pi}^m_n(\boldsymbol{D}^{o\to0}_m))] \geq 0, \tag{A33}$$

where the inequality follows because the marginal value of producing m at country o is decreasing in the set of production and sales locations for products $-m$, as shown in Properties 2 and 6. Finally, consider the summation across $n \in L_m$. Note that if $n \in L_m$, then necessarily $\underline{I}^0_{mn} = \underline{I}^1_{mn} = \overline{I}^1_{mn} = 1$ and $\overline{I}^0_{mn} = 0$.

As such,

$$\Sigma_{L_m} = \sum_{n \in L_m} \{[(\Delta \underline{\pi}_n^m(\boldsymbol{D}_m^{o\to 1}) - \Delta \overline{\pi}_n^m(\boldsymbol{D}_m^{o\to 1})) - (\Delta \underline{\pi}_n^m(\boldsymbol{D}_m^{o\to 0}) - \Delta \overline{\pi}_n^m(\boldsymbol{D}_m^{o\to 0}))] - \Delta \overline{\pi}_n^m(\boldsymbol{D}_m^{o\to 0})\} \geq 0,$$

where the inequality follows from Properties 2 and 6 and from the fact that for $n \in L_m$, $\overline{I}_{mn}^0 = 0 \implies \Delta \overline{\pi}_n^m(\boldsymbol{D}_m^{o\to 0}) < 0$.

We have therefore proven that $V_{b,s,m}$ exhibits decreasing differences in $\boldsymbol{D}_{-m}$ and $\boldsymbol{I}_{-m}$.

B Extensions of the First-Step Algorithm

B.1 Extension 1: Value-Function Mapping and Composition

Depending on the structure of the objective payoff function π, it may be possible to solve for a subset of decisions easily given values of other choices. Indeed, suppose that the following properties/assumptions hold:

Assumption 2 *Given $C = (A, B)$ with $A \in \mathcal{B}^L$ and $B \in \mathcal{B}^{J-L}$, it is feasible to compute:*

$$\mu(A) := \max_{B \in \mathcal{B}^{J-L}} \pi(A, B), \qquad \text{for any } A \in \mathcal{B}^J.$$

Moreover, suppose that the following additional assumption holds.

Assumption 3 (Constant Sign of Cross-Partials of $\mu(A)$) *For all $A \in \mathcal{B}^L$, and all $i, j \in \{1, \ldots, L\}$, the sign of $\Delta_{ij}^{\mu}(A)$ is known and independent of A.*

Given Assumptions 2-3, we can define set-increasing mappings F_μ and F_π. For any $\boldsymbol{C} \in S(\mathcal{B}^J)$, define

$$\overline{\Omega}_\mu(\boldsymbol{C}) = \{i \in \{1, \ldots, L\} \colon \Delta_i^{\mu}(\sup_i^{\mu}(\boldsymbol{C})) < 0\},$$
$$\underline{\Omega}_\mu(\boldsymbol{C}) = \{i \in \{1, \ldots, L\} \colon \Delta_i^{\mu}(\inf_i^{\mu}(\boldsymbol{C})) \geq 0\}.$$

Both $\sup_i^{\mu}(\boldsymbol{C})$ and $\inf_i^{\mu}(\boldsymbol{C})$ are $L \times 1$ vectors with elements

$$[\sup_i^{\mu}(\boldsymbol{C})]_j := \sigma_{ij}^{\mu}[\sup(\boldsymbol{C})]_j + (1 - \sigma_{ij}^{\mu})[\inf(\boldsymbol{C})]_j,$$
$$[\inf_i^{\mu}(\boldsymbol{C})]_j := \sigma_{ij}^{\mu}[\inf(\boldsymbol{C})]_j + (1 - \sigma_{ij}^{\mu})[\sup(\boldsymbol{C})]_j.$$

Then, we can define mapping $F_\mu \colon S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ as

$$F_\mu(\boldsymbol{C}) = \{C \in \boldsymbol{C} \colon C_i = 0 \text{ for all } i \in \overline{\Omega}_\mu(\boldsymbol{C}) \text{ and } C_{i'} = 1 \text{ for all } i' \in \underline{\Omega}_\mu(\boldsymbol{C})\}.$$

Note that F_μ only updates coordinates $i \in \{1, \ldots, L\}$ according to the μ function, and uses Assumption 3 to construct decision-specific suprema and infima under μ.

Similarly, for any $\boldsymbol{C} \in S(\mathcal{B}^J)$, define

$$\overline{\Omega}_\pi(\boldsymbol{C}) = \{i \in \{L+1, \ldots, J\} \colon \Delta_i^{\pi}(\sup_i(\boldsymbol{C})) < 0\},$$

$$\underline{\Omega}_\pi(\boldsymbol{C}) = \{i \in \{L+1, \ldots, J\} \colon \Delta_i^\pi(\inf_i(\boldsymbol{C})) \geq 0\}.$$

Both $\sup_i(\boldsymbol{C})$ and $\inf_i(\boldsymbol{C})$ are $J \times 1$ vectors with elements

$$[\sup_i(\boldsymbol{C})]_j := \sigma_{ij}[\sup(\boldsymbol{C})]_j + (1-\sigma_{ij})[\inf(\boldsymbol{C})]_j,$$
$$[\inf_i(\boldsymbol{C})]_j := \sigma_{ij}[\inf(\boldsymbol{C})]_j + (1-\sigma_{ij})[\sup(\boldsymbol{C})]_j.$$

Let $F_\pi \colon S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ as

$$F_\pi(\boldsymbol{C}) = \{C \in \boldsymbol{C} \colon C_i = 0 \text{ for all } i \in \overline{\Omega}_\pi(\boldsymbol{C}) \text{ and } C_{i'} = 1 \text{ for all } i' \in \underline{\Omega}_\pi(\boldsymbol{C})\}.$$

Note that F_π is identical to F in the main text except for the fact that it only updates coordinates with indices larger than L. Moreover, both mappings have the property that $F_\pi(\boldsymbol{C}^*) = \boldsymbol{C}^*$ and $F_\mu(\boldsymbol{C}^*) = \boldsymbol{C}^*$, where $\boldsymbol{C}^* = \{C^*\}$, since at the global maximum, there is no one-coordinate deviation that improves the values of μ nor π. As such, $F_{\pi\circ\mu} : S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ defined as $F_{\pi\circ\mu} = F_\pi \circ F_\mu$ also has fixed point $\boldsymbol{C}^*$. Since a composition of set-increasing mappings is also increasing, it follows that one can apply $F_{\pi\circ\mu}$ iteratively starting at $\boldsymbol{C}_0 = \mathcal{B}^J$ to obtain the supremum of the set of fixed points (Tarski 1955).

One important example that satisfies Assumptions 2-3 worth mentioning in the context of our application is the case of the single-product firm. In this case, as shown in Appendix A, the mapping in equation (24),

$$V_{b,s}(\boldsymbol{D}) := \max_{\boldsymbol{I}} \sum_{n=1}^{N} I_n \Delta\pi_n^m(\boldsymbol{D}) - \sum_{n=1}^{N} D_n F_n^p$$

exhibits decreasing differences in $\boldsymbol{D}_m$, and thus satisfies Assumption 3. Moreover, given production location choices $\boldsymbol{D}$, sales decisions are independent across countries, which means only N profit evaluations have to be computed (rather than 2^N if the problem were combinatorial). Thus, this problem satisfies Assumption 2.

In this particular case, we only require knowledge of the interdependence across production locations to obtain bounds on the solution for such production locations. In this case, there are only substitutabilities, so our algorithm is equivalent to Arkolakis et al. (2023). If the researcher is also interested in bounding the profit-maximizing value of $\boldsymbol{I}$ given bounds on $\boldsymbol{D}$, then knowledge of interdependence between $\boldsymbol{I}$ and $\boldsymbol{D}$ is required. In this case, they are complements under $\pi = \Pi_{bs}(\boldsymbol{I}, \boldsymbol{D})$, as shown in Appendix A. The composition of two mappings – first bounding production location choices and then bounding offering choices – yields bounds on the overall solution of the problem. The first mapping exhibits only substitutabilities, while the second mapping exhibits only complementarities.

B.2 Generalization of Extension 1: Composition of Mappings

In this section, we provide a technical generalization of Extension 1. We start by defining the first-step algorithm that uses some mapping μ, which in general is not equal to π, to solve for bounds on the argmax of π. The function μ takes values in the L-dimension Boolean set, rather than the J-dimensional Boolean

set, where $L \leq J$. We define the marginal value and the cross-partial under μ.

Definition 3 (Marginal Value under μ) *The marginal value of discrete choice i under C given function $\mu : \mathcal{B}^L \to \mathbb{R}$, with $L \leq J$, is given by,*

$$\Delta_i^\mu(C) := \mu(C^{i\to 1}) - \mu(C^{i\to 0}) \tag{A34}$$

where $C^{i\to 1}$ is the vector C with the i^{th} coordinate set to 1 (analogously for $C^{i\to 0}$).

Thus, $\Delta_i^\mu(C)$ is the change in $\mu(C)$ when changing the i^{th} coordinate from zero to one holding all other coordinates constant. We also define the cross-partial of i and j under C.

Definition 4 (Cross-Partial under μ) *For any two choices $i, j \in \{1, 2, ..., J\}$ and a function $\mu : \mathcal{B}^L \to \mathbb{R}$, we define the cross-partial of i and j under C as,*

$$\Delta_{ij}^\mu(C) := \Delta_i^\mu(C^{j\to 1}) - \Delta_i^\mu(C^{j\to 0}). \tag{A35}$$

Thus, $\Delta_{ij}^\mu(C)$ is the difference in marginal value of coordinate i when j is set to 1 relative to when j is set to 0.

Assumption 4 *For all $C \in \mathcal{B}^L$, and any $i, j \in \{1, 2, ..., L\}$, the sign of $\Delta_{ij}^\mu(C)$ is known and independent of C.*

Note that Assumption 4 requires knowledge of the sign of the interdependence across any pair of choices within L-dimensional subset of decisions on which μ operates.

Given Assumption 4, we define decision-specific suprema and infima in Definition 5.

Definition 5 *Given any $\boldsymbol{C} \in S(\mathcal{B}^J)$, the supremum and infimum for decision $i \in \{1, 2, ..., L\}$ under function $\mu : \mathcal{B}^L \to \mathbb{R}$ are $L-$dimensional vectors, respectively, with j^{th} element equal to,*

$$[\sup_i^\mu(\boldsymbol{C})]_j := \mathbb{1}[\Delta_{ij}^\mu \geq 0][\sup(\boldsymbol{C})]_j + \mathbb{1}[\Delta_{ij}^\mu < 0][\inf(\boldsymbol{C})]_j \tag{A36}$$

$$[\inf_i^\mu(\boldsymbol{C})]_j := \mathbb{1}[\Delta_{ij}^\mu \geq 0][\inf(\boldsymbol{C})]_j + \mathbb{1}[\Delta_{ij}^\mu < 0][\sup(\boldsymbol{C})]_j. \tag{A37}$$

Notice that $\sup_i^\mu$ and $\inf_i^\mu$ map, for each $i \in \{1, 2, ..., L\}$, lattice $\boldsymbol{C} \in S(\mathcal{B}^J)$ to elements in $\mathcal{B}^L$.[25]

To construct a set-increasing mapping in $S(\mathcal{B}^J)$, define, for any $\boldsymbol{C} \in S(\mathcal{B}^J)$,

$$\overline{\Omega}_\mu(\boldsymbol{C}) = \{i \in \mathcal{U}(\mu) : \Delta_i^\mu(\sup_i^\mu(\boldsymbol{C})) < 0\} \tag{A38}$$

$$\underline{\Omega}_\mu(\boldsymbol{C}) = \{i \in \mathcal{U}(\mu) : \Delta_i^\mu(\inf_i^\mu(\boldsymbol{C})) \geq 0\}, \tag{A39}$$

[25] We assume without loss of generality that the indices under μ are a subset of the indices under π, and that the indices otherwise coincide, i.e., the decision corresponding to the j^{th} coordinate under μ is also the j^{th} coordinate under π. In practice, mapping μ could apply to a subset of decisions $\mathcal{A}_\mu \subseteq \{1, 2, ..., J\}$, with $\mathcal{A}_\mu \neq \{1, 2, ..., L\}$, in which case one could readily define a dictionary mapping the set $\{1, 2, .., L\}$ to $\mathcal{A}_\mu$.

where $\mathcal{U}(\mu) \subseteq \{1, 2, ..., L\}$ is the weak subset of coordinates which mapping μ updates. That is, mapping μ may update only a subset of the L-dimensional vector on which it operates, though it requires knowledge of the interdependence across each of these L choices to define the decision-specific infima and suprema.

We define mapping $F_\mu : S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ as,

$$F_\mu(\boldsymbol{C}) = \{C \in \boldsymbol{C} \colon C_i = 0 \text{ and } C_{i'} = 1 \, \forall i \in \overline{\Omega}_\mu(\boldsymbol{C}) \text{ and } \forall i' \in \underline{\Omega}_\mu(\boldsymbol{C})\}. \tag{A40}$$

Intuitively, by Assumption 1, it is always possible to define a decision-specific supremum and infimum. Since the sign of all pairwise interdependencies are known and independent of third choices, one can always define the vectors $\sup_i^\mu(\boldsymbol{C})$ and $\inf_i^\mu(\boldsymbol{C})$ that maximize and minimize the marginal value of i, respectively. If at the "worst-case scenario" for i, $\inf_i^\mu(\boldsymbol{C})$, the marginal value of i is positive, then $i \in \underline{\Omega}_\mu(\boldsymbol{C})$ and $F_\mu(\boldsymbol{C})$ returns only vectors C with $C_i = 1$. If at "best-case scenario" for i, $\sup_i^\mu(\boldsymbol{C})$, the marginal value of i is negative, then $i \in \underline{\Omega}_\mu(\boldsymbol{C})$ and $F_\mu(\boldsymbol{C})$ returns only vectors C with $C_i = 0$.

Importantly, note that while Assumption 4 requires knowledge of the sign of the pairwise interdependence across all elements in $\{1, 2, ..., L\}$, mapping μ may, in general, only update a subset of those decisions $\mathcal{U}(\mu)$.

We prove that the mapping F_μ is set-increasing.

Theorem 3 (Set-Increasing F_μ) *The mapping $F_\mu : S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ is set-increasing. That is, $\boldsymbol{C} \subseteq \boldsymbol{C}'$ implies that $F_\mu(\boldsymbol{C}) \subseteq F_\mu(\boldsymbol{C}')$.*

Proof. See Appendix C.1. ■

Theorem 2 results from the monotonicity of the infimum and supremum together with Assumption 1. Note that for $\boldsymbol{C} \subseteq \boldsymbol{C'}$, it must be that for all i, $\Delta_i^\mu(\pi(\sup_i(\boldsymbol{C}))) \leq \Delta_i^\mu(\pi(\sup_i(\boldsymbol{C'})))$ and $\Delta_i^\mu(\inf_i^\mu(\boldsymbol{C}))) \geq \Delta_i^\mu(\inf_i^\mu(\boldsymbol{C'}))$. Thus, $\overline{\Omega}_\mu(\boldsymbol{C'}) \subseteq \overline{\Omega}_\mu(\boldsymbol{C})$ and $\underline{\Omega}_\mu(\boldsymbol{C'}) \subseteq \underline{\Omega}_\mu(\boldsymbol{C})$, implying $F_\mu(\boldsymbol{C}) \subseteq F_\mu(\boldsymbol{C'})$. Intuitively, starting at worse "worst-case scenarios" and better "best-case scenarios" solves fewer coordinates, leading F to rule out fewer vectors.

We describe the first-step procedure in Algorithm 2 for future reference.

Algorithm 2 (First-Step Algorithm under F_μ) *We implement the algorithm as follows:*

1. *Start with $\boldsymbol{C}_0 = \mathcal{B}^J$.*
2. *Apply F_μ to obtain $\boldsymbol{C}_1 = F_\mu(\boldsymbol{C}_0)$; note that $\boldsymbol{C}_1 \in S(\mathcal{B}^J)$.*
3. *Obtain $\boldsymbol{C}_k = F_\mu(\boldsymbol{C}_{k-1})$ for $k = 1, 2, ...$ until convergence.*

To ensure that mapping F_μ converges to a set containing the argmax of π, an additional assumption is required relating function μ with the objective function, π. The required assumption is stated in Assumption 5.

Assumption 5 *Let $\boldsymbol{C}^* \in S(\mathcal{B}^J)$ be the set containing the maximizer of π, i.e., $\boldsymbol{C}^* = \{C^*\}$. The mapping $F_\mu : S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ is such that $F_\mu(\boldsymbol{C}^*) = \boldsymbol{C}^*$.*

Note that Assumption 5 is trivially satisfied, under Assumption 4, whenever $\mu = \pi$. Another possible function μ satisfying Assumption 5 is the max operator over a subset of coordinates, as discussed in Appendix B.1.

Proposition 3 states that Algorithm 1 converges monotonically to a set $\boldsymbol{C}^f$ such that $C^* \in \boldsymbol{C}^f$.

Proposition 3 *Under Assumptions 4-5, Algorithm 1 has the following properties:*

1. $\boldsymbol{C}_k = F_\mu(\boldsymbol{C}_{k-1}) \subseteq \boldsymbol{C}_{k-1}$ *for all* $k \in \mathbb{N}$.
2. *It converges to a fixed point* $\boldsymbol{C}^f$ *in a finite number of iterations* K *i.e.,* $F_\mu(\boldsymbol{C}_T) = \boldsymbol{C}^f$ *for all* $T \geq K$.
3. *The solution to the problem* C^* *is an element of* $\boldsymbol{C}^f$.

Proof. Due to Tarski (1955), we can start at $\boldsymbol{C}_0 = \mathcal{B}^J$ and iteratively apply F_μ until convergence to fixed point $\boldsymbol{C}^f$, which is the supremum of the set of fixed points. Convergence occurs monotonically because F is set-increasing, so that $\boldsymbol{C}_1 = F_\mu(\boldsymbol{C}_0) \subseteq \boldsymbol{C}_0$ implies that $F_\mu(\boldsymbol{C}_1) \subseteq F_\mu(\boldsymbol{C}_0)$. The algorithm converges in a finite number of iterations, because the number of iterations is bounded above by L, the number of binary decisions that are to be solved. The set $\{C^*\}$, the unique (almost surely) solution to the brand's problem is also a fixed point, since, by Assumption 5, $C_i^* = 1$ if and only if the marginal value of i is positive i.e., $\Delta_i^\mu(C^*) \geq 0$. Thus, $C^* \in \boldsymbol{C}^f$. ■

To bound the solution, the researcher may choose to apply Algorithm 1 with mapping $F_\mu = F_\pi$ only. However, as we show in our application, it may be more informative to apply compositions of set-increasing mappings, each satisfying Assumptions 4 and 5, where each of the components of the composition operates on a distinct subset of the J-dimensional coordinates that are to be solved. Proposition 4 shows that any composition of mappings satisfying Assumptions 4-5 can be used to bound the argmax of π.

Proposition 4 *Any finite composition of set-increasing mappings* $F = F_{\mu_1} \circ F_{\mu_2} \circ ... \circ F_{\mu_K}$*, each of which satisfies Assumptions 4-5, is set-increasing. Applying Algorithm 2 with mapping* F *therefore yields bounds on the argmax of* π.

Proof. Any composition of set-increasing functions is set-increasing. Therefore, $F : S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ is set-increasing. The set $\{C^*\}$ containing only the argmax of π is a fixed point of each of F_{μ_k}, for $k = 1, 2, ..., K$. As such, the set $\{C^*\}$ is also a fixed point of F. It follows that applying Algorithm 2 using F yields bounds on the argmax of π. ■

B.3 Extension 2: Additive Payoff Function

In this section, we relax Assumption 1 under additional structure on the payoff function, π. Suppose that π has the form,

$$\pi(C) = \sum_{k=1}^{K} \pi_k(C).$$

That is, π can be written as the sum of K additive components π_k, for $k = 1, 2, ..., K$. Under such a payoff function, we now only require that the sign of the interdependence across any pair of decisions i

and j is known *within* each component k. Thus, pairs of binary choices need not be always complements or substitutes regardless of the value of C according to the overall payoff function π. This assumption is stated below.

Assumption 6 *For all $C \in \mathcal{B}^J$, any $k \in \{1,2,...,K\}$, and any $i,j \in \{1,2,...,J\}$, the sign of $\Delta_{ij}^{\pi_k}(C)$ is known and independent of C.*

To construct a set-increasing mapping in $S(\mathcal{B}^J)$, define, for any $\boldsymbol{C} \in S(\mathcal{B}^J)$,

$$\overline{\Omega}_K(\boldsymbol{C}) = \{i \in \{1,2,...,J\} : \sum_{k=1}^{K} \Delta_i^{\pi_k}(\sup_i^k(\boldsymbol{C})) < 0\} \tag{A41}$$

$$\underline{\Omega}_K(\boldsymbol{C}) = \{i \in \{1,2,...,J\} : \sum_{k=1}^{K} \Delta_i^{\pi_k}(\inf_i^k(\boldsymbol{C})) \geq 0\} \tag{A42}$$

where, for each $i \in \{1,2,...,J\}$, $j \in \{1,2,...,J\}$, and $k \in \{1,2,...,K\}$,

$$[\sup_i^k(\boldsymbol{C})]_j := \mathbb{1}[\Delta_{ij}^{\pi_k} \geq 0][\sup(\boldsymbol{C})]_j + \mathbb{1}[\Delta_{ij}^{\pi_k} < 0][\inf(\boldsymbol{C})]_j \tag{A43}$$

$$[\inf_i^k(\boldsymbol{C})]_j := \mathbb{1}[\Delta_{ij}^{\pi_k} \geq 0][\inf(\boldsymbol{C})]_j + \mathbb{1}[\Delta_{ij}^{\pi_k} < 0][\sup(\boldsymbol{C})]_j. \tag{A44}$$

We define mapping $F_K : S(\mathcal{B}^J) \to S(\mathcal{B}^J)$ as,

$$F_K(\boldsymbol{C}) = \{C \in \boldsymbol{C} \colon C_i = 0 \text{ and } C_{i'} = 1 \, \forall i \in \overline{\Omega}_K(\boldsymbol{C}) \text{ and } \forall i' \in \underline{\Omega}_K(\boldsymbol{C})\}. \tag{A45}$$

Mapping F_K is a strict generalization of mapping F.

B.3.1 Example: Economies of Scope

A notable advantage of Extension 2 is that the same pair of choices may be complements under certain values of C and as substitutes under others. For instance, suppose that the fixed costs of production for a given model have the form,

$$\tilde{F}_{mn}^p = \gamma_4 \mathbb{1}\{\sum_{m' \neq m} D_{m'n} = 0\} + F_{mn}^p.$$

Under this specification, the brand pays a fixed cost γ_4 for entering a country, and then an additional fixed cost for each additional product it sells in that market. The market-level fixed cost γ_4 induces a complementarity across production locations stemming from fixed cost savings. In this case, brands' are profits are a sum of two components: (i) $\pi_{b,s}(\boldsymbol{I},\boldsymbol{D})$ and (ii) $(-F_{b,s}^p(\boldsymbol{D}))$, where the first component is the sum of market-level variable profits net of fixed selling costs and the second component is production location fixed costs. Under fixed costs exhibiting economies of scope, $F_{b,s}(\boldsymbol{D})$ now exhibits complementarities across production locations for the same product. Meanwhile, as in our baseline case, $\pi_{b,s}(\boldsymbol{I},\boldsymbol{D})$ exhibits substitutabilities across production locations for the same product arising from minimum cost sourcing.

B.4 Applying Extension 1 to the Value Function

While Algorithm 1 is guaranteed to converge to a set containing the solution, the set may still be large when directly applied to an arbitrary payoff function π. Equivalently, the solution for only a relatively small number of discrete choices might be found at convergence. In our model of multiproduct, multinational production, there is an economic reason why the algorithm may fail to provide informative bounds on the solution when applied directly to Π_{bs}. For simplicity, consider a single-product firm deciding whether to set a production plant for product m in country o, D_{mo}. Algorithm 1 solves for a 1 if the marginal value of setting $D_{mo} = 1$ relative to $D_{mo} = 0$ at the decision-specific infimum is positive. In other words, if in the most unfavorable conditions for D_{mo} given the interdependence structure, the marginal gains from setting this choice to 1 are positive. At the initialization of the algorithm, the most unfavorable conditions for producing m in country o consist of having plants in all countries $k \neq o$ open for product m, and selling product m nowhere. Crucially, *conditional* on selling product m nowhere, the marginal gains from producing product m in country o will be zero minus the fixed cost of production and, therefore, unlikely to be positive.[26]

A way around this practical problem is to apply Algorithm 2 (see Appendix B.2) to $V_{bs}(\boldsymbol{D}^{bs}) = \max_{\boldsymbol{I}^{bs}} \Pi_{bs}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$. To ensure that the algorithm can be applied, what is required is to show that the discrete choices in $\boldsymbol{D}^{bs}$ satisfy Assumption 4 under objective function V_{bs}. To see why this would help, note that in the previous example, we are now not comparing the gains from setting $D_{mo} = 1$ relative to $D_{mo} = 0$ conditional on selling m nowhere, but rather solving for product m offerings decisions under $D_{mo} = 1$ and separately under $D_{mo} = 0$ (conditional on all other production locations for product m being open). Such an evaluation is more likely to yield positive profits if the optimal offering choices conditional on $D_{mo} = 1$ yield higher profits than the optimal offering choices under $D_{mo} = 0$.

We show that for single-product brands, mapping $V_{bs}(\boldsymbol{D}^{bs})$ satisfies Assumption 4.

Proposition 5 *If $M_{bs} = 1$, the value function $V_{bs}(\boldsymbol{D}^{bs}) = \max_{\boldsymbol{I}^{bs}} \Pi_{bs}(\boldsymbol{I}^{bs}, \boldsymbol{D}^{bs})$ is such that D_{1o} and $D_{1o'}$ are substitutes for $o \neq o'$.*

Proof. See Appendix A.2.1. ■

When brands are multiproduct, we apply a composition of set-increasing mappings, as described in Appendix B.2, based on the value functions:

$$V_{b,s,m}(\boldsymbol{D}_m; \boldsymbol{I}_{-m}, \boldsymbol{D}_{-m}) := \max_{\boldsymbol{I}_m} \sum_{n=1}^{N} I_{mn} \Delta\pi_n^m(\boldsymbol{D}_m; \boldsymbol{I}_{-m}, \boldsymbol{D}_{-m}) - \sum_{n=1}^{N} D_{mn} F_{mn}^p, \tag{A46}$$

where $\Delta\pi_n^m$ is the marginal value of selling product m in country n. The value function in equation (A46) essentially conditions on all production location choices for all products and sales choices for products other than m, and maximizes the additional profits that the brand can earn by optimally choosing its sales decisions for model m across countries. In Appendix A.2.3, we show that conditional on the optimal

[26] Under our assumption of normal fixed costs, it is possible that the fixed cost of production is negative but with a small probability under a large mean and relatively small variance.

solution $(\boldsymbol{I}^*_{-m}, \boldsymbol{D}^*_{-m})$ for all decisions not concerning product m, the program in equation (A46) returns $\boldsymbol{I}^*_m$.

Of course, the objective is to find the optimal solution, so we cannot condition on $(\boldsymbol{I}^*_{-m}, \boldsymbol{D}^*_{-m})$. We prove that starting from any known bounds of the solution, we can iteratively apply Algorithm 2 product by product to obtain bounds of the solution. To apply Algorithm 2 under $V_{b,s,m}$, we prove in Appendix A.2 that mapping $V_{b,s,m}$ exhibits (i) decreasing differences in $\boldsymbol{D}_m$, and (ii) decreasing differences in $(\boldsymbol{I}_{-m}, \boldsymbol{D}_{-m})$. We describe the algorithm below.

Algorithm 3 (Production Location Solver) *1. Start at the known bounds of the solution. At the first iteration, initialize with* $(\overline{\boldsymbol{I}}^{bs}, \overline{\boldsymbol{D}}^{bs}) = (\mathbf{1}, \mathbf{1})$ *and* $(\underline{\boldsymbol{I}}^{bs}, \underline{\boldsymbol{D}}^{bs}) = (\mathbf{0}, \mathbf{0})$.

2. *Start at* $m = 1$.

3. *Apply Algorithm 1 to* $V_{b,s,1}$ *until convergence. Construct the decision-specific infima and suprema using knowledge of substitutability across all discrete choices.*

4. *Repeat Step 3 for all products* $m \in \{2, ..., M_{bs}\}$.

5. *Repeat Steps 2-3 until convergence.*

Intuitively, the algorithm starts at any known bounds of the solution and exploits substitutability across all discrete choices according to the $V_{b,s,m}$ mapping. Consider the problem of solving for $D_{mo} = 1$. Due to the substitutabilty properties proven in Appendix A.2, the algorithm evaluates the marginal change in $V_{b,s,m}$ by setting $D_{mo'} = 1$ for all $o' \neq o$, and $D_{m'n} = I_{m'n} = 1$ for all $m' \neq m$ and for all $n \in \mathcal{N}$. If this marginal change is *positive*, the algorithm says that $\underline{D}_{mo} = 1 = D^*_{mo}$. Consider the problem of solving for $D_{mo} = 0$. In this case, the algorithm evaluates the marginal change in $V_{b,s,m}$ by assuming that $D_{mo'} = 0$ for all $o' \neq o$, and $D_{m'n} = I_{m'n} = 0$ for all $m' \neq m$ and for all $n \in \mathcal{N}$. If this marginal change is *negative*, the algorithm says that $\overline{D}_{mo} = 0 = D^*_{mo}$.

Proposition 6 *Algorithm 3 converges to bounds on the solution* $\boldsymbol{D}^{bs,*}$ *of brand* b *in segment* s.

Proof. See Appendix C.2. ■

Algorithm 3 consists of applying Algorithm 1 iteratively to different products under the $V_{b,s,m}$ function. It is, therefore, a composition of set-increasing mappings. A composition of increasing functions is itself increasing. As such, we can apply the Algorithms 2 and 3 iteratively and in any order provided we start at any known bounds of the solution. In practice, we start with Algorithm 3 to solve for production location decisions, and then apply Algorithm 1 to Π_{bs} to solve for both sales and production location decisions. We describe the implementation in detail in Appendix D.

C Method Proofs

C.1 Proof of Theorems 2 and 3

We prove Theorem 3. For $\boldsymbol{C} \subseteq \boldsymbol{C'}$, it must be that for all i, $\Delta^{\mu}_i(\pi(\sup_i(\boldsymbol{C}))) \leq \Delta^{\mu}_i(\pi(\sup_i(\boldsymbol{C'})))$ and $\Delta^{\mu}_i(\inf^{\mu}_i(\boldsymbol{C}))) \geq \Delta^{\mu}_i(\inf^{\mu}_i(\boldsymbol{C'}))$. Thus, $\overline{\Omega}_{\mu}(\boldsymbol{C'}) \subseteq \overline{\Omega}_{\mu}(\boldsymbol{C})$ and $\underline{\Omega}_{\mu}(\boldsymbol{C'}) \subseteq \underline{\Omega}_{\mu}(\boldsymbol{C})$, implying $F_{\mu}(\boldsymbol{C}) \subseteq F_{\mu}(\boldsymbol{C'})$.

It follows that F_μ is a set-increasing mapping. The proof of Theorem 2 follows by setting $\mu = \pi$ and $\mathcal{U}(\mu) = \{1, 2, ..., J\}$.

C.2 Proof of Proposition 6

Given the proof for Proposition 3, the proof of Proposition 6 requires three steps. First, note that conditional on m and on $(\boldsymbol{I}_{-m}, \boldsymbol{D}_{-m})$, we can apply Algorithm 1 by applying $\mu = V_{b,s,m}$. Function $V_{b,s,m}$ maps $(\boldsymbol{D}_m, \boldsymbol{I}_{-m}, \boldsymbol{D}_{-m})$ to a real number, and exhibits substitutabilities across all discrete choices, as established in Corollary A.2.2. This implies that $V_{b,s,m}$ satisfies Assumption 1. Thus, applying algorithm 1 using $V_{b,s,m}$ yields a set-increasing mapping. Second, to ensure that the mapping yields bounds of the solution, we must show that $V_{b,s,m}$ satisfies Assumption 5. This follows immediately from Proposition 2, proven in section A.2.3. Finally, we show that the composition of mappings $F_{V_{b,s,M_{bs}}} \circ ... \circ F_{V_{b,s,1}}$ is set-increasing and satisfies Assumption 5. This follows from the proof of Proposition 4.[27]

C.3 Properties of the Branching Algorithm

We start by defining the second-step (branching) algorithm formally.

Let $\boldsymbol{C} \in S(\mathcal{B}^J)$ be a sublattice. Let $Z(\boldsymbol{C}) = \{i \in \{1, 2, ..., J\} : [\sup(\boldsymbol{C})]_i = 0\}$ be the set of coordinates that are solved to be 0 at $\boldsymbol{C}$ and $O(\boldsymbol{C}) = \{i \in \{1, 2, ..., J\} : [\inf(\boldsymbol{C})]_i = 1\}$ be the set of coordinates that are solved to be 1 at $\boldsymbol{C}$. Let $U(\boldsymbol{C}) = Z(\boldsymbol{C})^c \cap O(\boldsymbol{C})^c$. Since there are a finite number of elements in $U(\boldsymbol{C})^c$, we can define an ordering of coordinates in $U(\boldsymbol{C})^c$ as a function of $\boldsymbol{C}$. Let $W : S(\mathcal{B}^J) \to \{1, 2, ..., J\}$ be a function mapping a sublattice $\boldsymbol{C}$ to a coordinate $i \in U(\boldsymbol{C})$. We are now ready to define the branching mapping.

Definition 6 (Branching Mapping) *The branching mapping is defined by function* $B : \mathcal{P}(S(\mathcal{B}^J)) \to \mathcal{P}(S(\mathcal{B}^J))$ *as,*

$$B(\boldsymbol{P}) = \bigcup_{\boldsymbol{C} \in \boldsymbol{P}} \{F_*(\boldsymbol{C}_1), F_*(\boldsymbol{C}_0)\},$$

where $\boldsymbol{C}_1 = \{C \in \boldsymbol{C} : C_i = 1 \text{ for } i = W(\boldsymbol{C})\}$, $\boldsymbol{C}_0 = \{C \in \boldsymbol{C} : C_i = 0 \text{ for } i = W(\boldsymbol{C})\}$, *and* F_* *denotes application of the first-step set-increasing mapping* F *until convergence. Clearly,* $\boldsymbol{C}_1 \cup \boldsymbol{C}_0 = \boldsymbol{C}$ *and* $\boldsymbol{C}_1 \cap \boldsymbol{C}_0 = \emptyset$. *That is* $B(\boldsymbol{P})$ *maps a set of sublattices of* $S(\mathcal{B}^J)$ *to another set of sublattices by "branching" each of the sublattices and applying* F_* *in each of these branches.*

We can now define the branching algorithm.

Definition 7 (Branching Algorithm) *The branching algorithm is defined as follows:*

1. *Start at* $\boldsymbol{C} \in S(\mathcal{B}^J)$ *such that* $C^* \in \boldsymbol{C}$. *Obtain* $\boldsymbol{P}_1 = B(\{\boldsymbol{C}\})$.
2. *Compute* $\boldsymbol{P}_2 = B(\boldsymbol{P}_1)$

[27] In Appendix D, we discuss our practical implementation.

3. *Iterate until convergence to fixed point* $\boldsymbol{P}^*$*. By construction,* $\boldsymbol{P}^*$ *is a collection of sets with no more than one element.*

4. *Evaluate* π *at all elements of the singleton sets in* $\boldsymbol{P}^*$ *to obtain the argmax of* π*,* C^**.*

Note that the algorithm does not require to start at the supremum of $\mathcal{P}(S(\mathcal{B}^J))$, which is $S(\mathcal{B}^J)$. Though one could start at the supremum, we show that this is computationally inefficient compared to starting at $\{\boldsymbol{C}\}$, with $C^* \in \boldsymbol{C}$. Before proving that the Branching Algorithm yields the global maximizer of π, we prove the following proposition.

Proposition 7 *For any* $\boldsymbol{C}_0 \in S(\mathcal{B}^J)$*, applying* F *repeatedly converges to a fixed point of* F*,* $\underline{\boldsymbol{C}}^f$*, in a finite number of iterations. Moreover, if* $C^* \in \boldsymbol{C}_0$*, then* $C^* \in \underline{\boldsymbol{C}}^f$*.*

Proof. Mapping F is set-increasing, and by construction, $F(\boldsymbol{C}_0) \subseteq \boldsymbol{C}_0$. As such, and because any non-increasing sequence of sets is convergent, it must be that mapping F converges to $\underline{\boldsymbol{C}}^f$ in a finite number of iterations. Moreover, the set $\{C^*\}$ is a fixed point such that $\{C^*\} \subseteq \boldsymbol{C}_0$. Applying F yields $\{C^*\} \subseteq F(\boldsymbol{C}_0)$. Applying F any finite number of times K yields $\{C^*\} \subseteq F^K(\boldsymbol{C}_0)$, with $F^K(\boldsymbol{C}_0) \subseteq F^{K-1}(\boldsymbol{C}_0)$. Thus, for any finite K, $F^K(\boldsymbol{C}_0)$ is non-empty and contains the solution. Due to monotonicity, either $F^K(\boldsymbol{C}_0) = \{C^*\}$ for some finite K, or $F^K(\boldsymbol{C}_0) = F^{K-1}(\boldsymbol{C}_0) = \underline{\boldsymbol{C}}^f$. ■

Proposition 8 *The branching algorithm yields the global maximizer of* π*,* C^**.*

Proof. First, note that $(\mathcal{P}(S(\mathcal{B}^J)), \subseteq)$ is a complete lattice, since it a finite lattice. Indeed, the meet is the intersection of the two elements (which may be the empty set), and the join is the union. Second, we show that mapping B is set-increasing. Take $\boldsymbol{P}_1 \subseteq \boldsymbol{P}_2$. Then,

$$B(\boldsymbol{P}_1) = \bigcup_{\boldsymbol{C} \in \boldsymbol{P}_1} \{F_*(\boldsymbol{C}_1), F_*(\boldsymbol{C}_0)\} \subseteq \bigcup_{\boldsymbol{C} \in \boldsymbol{P}_2} \{F_*(\boldsymbol{C}_1), F_*(\boldsymbol{C}_0)\} = B(\boldsymbol{P}_2).$$

By Tarski (1955), it follows that the set of fixed points of B is a complete lattice. Third, if $\boldsymbol{P}^*$ is a fixed point of B, then necessarily all elements $\boldsymbol{C} \in \boldsymbol{P}^*$ are singletons or the empty set. Suppose this were not the case and there was an element $\hat{\boldsymbol{C}}$ with more than one element. Then, applying $\{F_*(\hat{\boldsymbol{C}}_1), F_*(\hat{\boldsymbol{C}}_0)\}$ further partitions $\hat{\boldsymbol{C}}$, implying that $\boldsymbol{P}^*$ was not a fixed point. Due to finiteness, convergence to the fixed point occurs in a finite number of iterations. Finally, we prove that $\{C^*\} \in \boldsymbol{P}^*$. This must hold due to Proposition 7. Indeed, since one starts at $\boldsymbol{C} \in S(\mathcal{B}^J)$ with $C^* \in \boldsymbol{C}$, necessarily one of the branches always contains C^*, and application of B (and F) at this branch always retains the solution C^*. ■

C.4 Simple Example: Branching Algorithm

In this subsection, we provide a simple example for the branching algorithm. For illustration purposes, suppose there are 3 binary choices to solve. Recall that $S(\mathcal{B}^3)$ includes the empty set and all sublattices of $\mathcal{B}^3$ that include all elements that lie between their suprema and infima. For instance, an element of $S(\mathcal{B}^3)$ is $\{(1,0,0),(1,0,1),(1,1,0),(1,0,0)\}$, but not $\{(1,1,1),(0,0,0)\}$. To describe branching, consider

now the power set of $S(\mathcal{B}^3)$. An example of an element of such a power set is

$$\big\{\{(1,0,0),(1,0,1),(1,1,0),(1,0,0)\},\{(0,0,1)\}\big\},$$

since this is a subset of $S(\mathcal{B}^3)$. Another example would be $S(\mathcal{B}^3)$ itself.

The branching algorithm starts at $\{\mathcal{B}^3\}$, since we know that $C^* \in \mathcal{B}^3$. Then, it follows the following steps:

1. Obtain $\boldsymbol{P}_1 = B(\{\mathcal{B}^3\}) = \{F_*(\mathcal{B}_1^3), F_*(\mathcal{B}_0^3)\}$, where,

$$\mathcal{B}_1^3 = \{(1,0,0),(1,0,1),(1,1,0),(1,1,1)\},$$
$$\mathcal{B}_0^3 = \{(0,0,0),(0,0,1),(0,1,0),(0,1,1)\}.$$

   Notice that $\mathcal{B}_1^3 \cup \mathcal{B}_0^3 = \mathcal{B}^3$. The solution C^* must therefore be an element of one of these sets. Suppose that,

$$F_*(\mathcal{B}_1^3) = \{(1,0,0),(1,0,1)\}$$
$$F_*(\mathcal{B}_0^3) = \{(0,1,0)\}.$$

   Conditional on entry in location 1, then either $(1,0,0)$ or $(1,0,1)$ are the solution. Conditional on non-entry in location 0, the solution is $(0,1,0)$.

2. We now have a set of sublattices,

$$\big\{\{(1,0,0),(1,0,1)\},\{(0,1,0)\}\big\},$$

   to which we apply mapping B again. We have,

$$B(\big\{\{(1,0,0),(1,0,1)\},\{(0,1,0)\}\big\}) = \{F_*(\{(1,0,0)\}), F_*(\{(1,0,1)\}, F_*(\{(0,1,0)\}\}.$$

   We are now left with a set of sublattices that are either empty or singletons. Note that at least one of them must be a singleton, given that one of the branches converges to $\{C^*\}$, as implied by Proposition 7.

3. The last step is to compute π at each of the singleton sets and to return the maximizer.

D Practical Implementation of Algorithms 1 and 3

Table 4 provides a pseudocode description of our first-step algorithm implementation, which combines Algorithms 1 and 3. The code is a straightforward for-loop over products (indexed by m) and potential production locations (indexed by l) for brand b in segment s. We describe each of the steps of the code in words:

- Step 1: In this step, we initialize a while loop by setting an exit condition equal to `true`. We also start at known bounds of the solution. The brand has $M_{bs} \times (N + L)$ binary decisions to make,

Table 4: Pseudocode description of the `CMMS_prod` function.

Step	Pseudocode Description
1	**Initialize:** Set `exit_condition = true`. Copy `fixed` matrix of binary decisions to `output`.
2	**While** `exit_condition is true`, **do**: Store previous output: `output_prev = output`.
3	**For each product** m: **For each potential production location** l: **If** `output_prev[m, N+l] == 0.5`, then: Set `index = (m, N+l)`. Upper bound: `upper = MVmax_sup(output, index, profit_func)` **If** `upper < 0.0`, then set `output[index] = 0.0`. **Else:** Lower bound: `lower = MVmax_inf(output, index, profit_func)` **If** `lower` $\geq$ `0.0`, then: Set `output[index] = 1.0`
4	**Update offering decisions for product** m**:** Call `CMMS_sales(output, profit_func, m)` to adjust sales.
5	**Check convergence:** **If** `output_prev == output`, set `exit_condition = false`.
6	**Return:** Output the final production allocation: `return output`.

where N denotes the potential destination countries and L the potential production countries. The matrix `fixed` is an $M_{bs} \times (N + L)$ matrix with values 0.5 at each coordinate, denoting that no binary decision has been solved for at the initialization step.

- Step 2: The second step stores in `output_prev` the matrix `output` computed in the previous iteration.
- Step 3: For-loop over products $m \in \{1, 2, ..., M_{bs}\}$ and potential production locations $l \in \{1, 2, ..., L\}$. If the matrix `output_prev` is 0.5 at coordinate $(m, N + l)$ (at production location l), the algorithm first computes the upper bound using the `MVmax_sup` function, which computes the largest possible marginal value of (m, l) under function $V_{b,s,m}$ (see equation (A46)) – that is, setting $(\boldsymbol{I}_{-m}, \boldsymbol{D}_{-m})$ at lower bounds, as per the substitutability. If such an upper bound is strictly negative, we update the solution matrix `output` and set coordinate $(m, N + l)$ to 0. Otherwise, we compute the lower bound by applying `MVmax_inf`, which computes the smallest possible marginal value of (m, l) under function $V_{b,s,m}$ – that is, setting $(\boldsymbol{I}_{-m}, \boldsymbol{D}_{-m})$ at upper bounds, as per the substitutability. If such a lower bound is weakly positive, we update the solution matrix `output` and set coordinate $(m, N+l)$ to 1.

- Step 4: Given the solved 0's or 1's for the production locations of product m, we apply function `CMMS_sales` to product m offering decisions. The function `CMMS_sales` effectively applies Algorithm 1 to all offering choices of product m. We provide the pseudocode description of `CMMS_sales` in Table 5.

- Step 5: After looping over all products, the algorithm checks whether binary decisions were solved in Steps 3 and 4. If none were solved, the algorithm sets `exit_condition = false` and exits the `while` loop. Otherwise, the algorithm repeats Steps 3-4.

- If `exit_condition = false`, the algorithm returns `output`, which is a matrix containing 0's, 1's, and 0.5's.

Our implementation of the algorithm composes Algorithms 1 and 3, which we proved were both set-increasing. Thus, starting at known bounds of the solution, such a composition converges to bounds of the solution of the brand's problem.

Table 5: Pseudocode description of the `CMMS_sales` function.

Step	Pseudocode Description
1	**Initialize:** Copy `fixed` matrix of binary decisions to `output`.
2	**Iterate over sales locations:** **For each sales destination** $n = 1$ to N: **If** `output[m, n] == 0.5`, then: Set `index = (m, n)`
3	**Compute bounds and update** `output`**:** Upper bound: `upper = MVsup(output, index, profit_func)` **If** `upper < 0.0`, then set `output[index] = 0.0`. **Else:** Lower bound: `lower = MVinf(output, index, profit_func)` **If** `lower` $\geq$ `0.0`, then set `output[index] = 1.0`
5	**Return:** Output the updated sales allocation: `return output`.

Formally, in our practical implementation, we apply set-increasing mapping $F^{bs} = (F^D_{M_{bs}} \circ F^I_{M_{bs}}) ... \circ (F^D_m \circ F^I_m) \circ ... \circ (F^D_2 \circ F^I_2) \circ (F^D_1 \circ F^I_1)$, which is a composition of set-increasing mappings, where, for $\boldsymbol{C} \in S(\mathcal{B}^{M_{bs} \times (N+L)})$,

$$F^I_m(\boldsymbol{C}) = \{C \in \boldsymbol{C} : C_i = 0 \text{ and } C_{i'} = 1 \ \forall i \in \overline{\Omega}^I_m(\boldsymbol{C}) \text{ and } \forall i' \in \underline{\Omega}^I_m(\boldsymbol{C})\}$$
$$F^D_m(\boldsymbol{C}) = \{C \in \boldsymbol{C} : C_i = 0 \text{ and } C_{i'} = 1 \ \forall i \in \overline{\Omega}^D_m(\boldsymbol{C}) \text{ and } \forall i' \in \underline{\Omega}^D_m(\boldsymbol{C})\},$$

and

$$\overline{\Omega}^I_m(\boldsymbol{C}) = \{n \in \{1, 2, ..., N\} : \Delta_{(m,n)}(\Pi_{bs}(\sup_{(m,n)}(\boldsymbol{C})) < 0\}$$

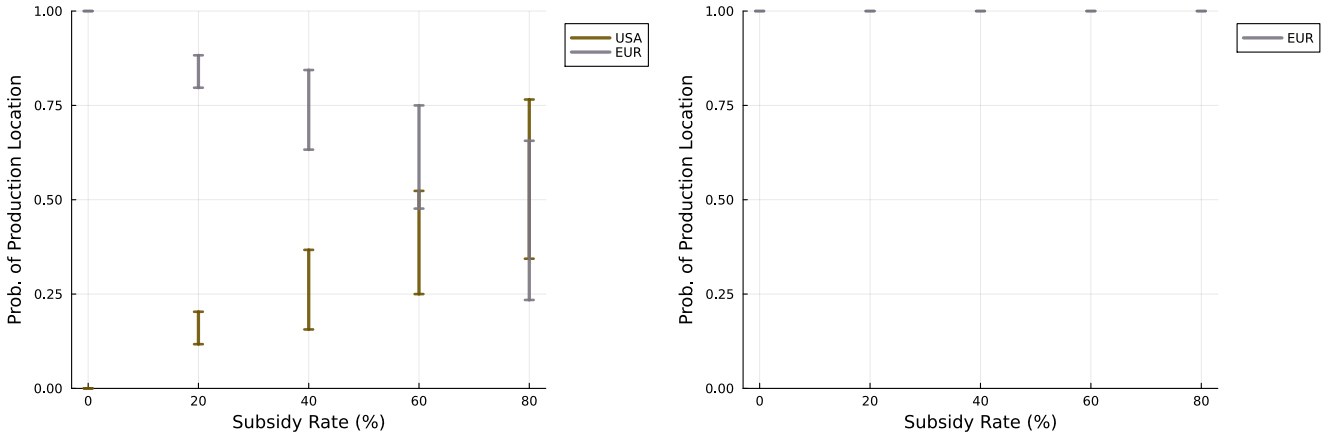
$$\underline{\Omega}_m^I(\boldsymbol{C}) = \{n \in \{1,2,...,N\} : \Delta_{(m,n)}(\Pi_{bs}(\inf_{(m,n)}(\boldsymbol{C})) \geq 0\}$$
$$\overline{\Omega}_m^D(\boldsymbol{C}) = \{l \in \{1,2,...,L\} : \Delta_{(m,N+l)}(V_{b,s,m}(\sup_{(m,N+l)}(\boldsymbol{C})) < 0\}$$
$$\overline{\Omega}_m^D(\boldsymbol{C}) = \{l \in \{1,2,...,L\} : \Delta_{(m,N+l)}(V_{b,s,m}(\inf_{(m,N+l)}(\boldsymbol{C})) \geq 0\}.$$

Due to Propositions 3-6, the algorithm we implement in practice is a composition of set-increasing mappings, which is itself set-increasing, and therefore converges to bounds of the solution of the problem under the interdependence properties proven under Π_{bs} and $V_{b,s,m}$.

E Plots of Counterfactual Exercises

E.1 Audi Wagons: U.S. Electic-Vehicle Marginal Cost Subsidy

Figure 19: Production Location - Marginal Cost Subsidy on the A6

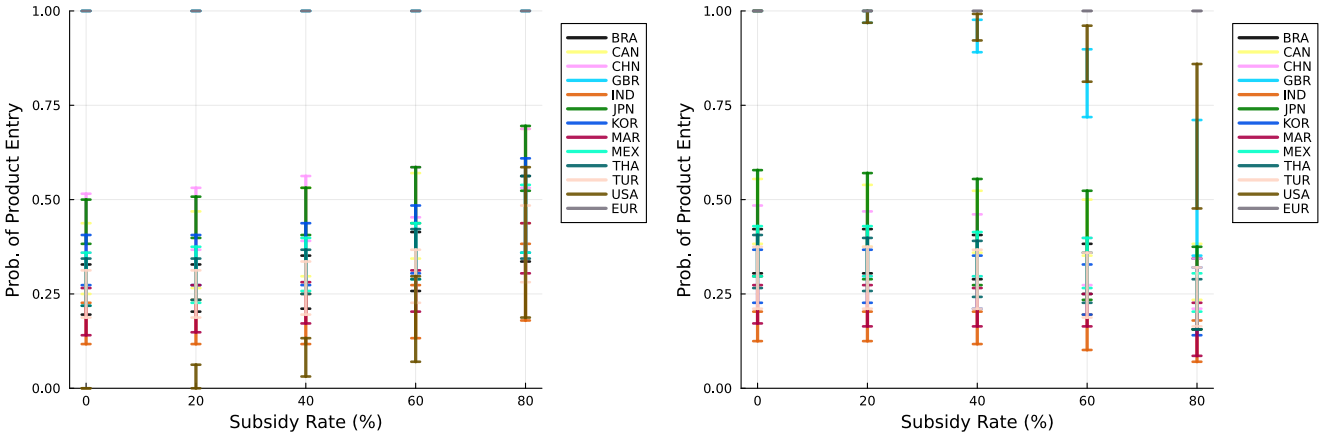


(a) A4: Production Location Probability

(b) A6: Production Location Probability

Note: These plots show the production location probability for the Audi A4 in different countries/regions at different U.S. marginal cost subsidies that affect the A4 and not the A6. The bounds reflect both statistical uncertainty regarding the bounds on the fixed cost parameters and uncertainty about the solution arising from potential inability to solve the full brand problem. "EUR" denotes the European Union.

Figure 20: Product Entry - Marginal Cost Subsidy on the A4

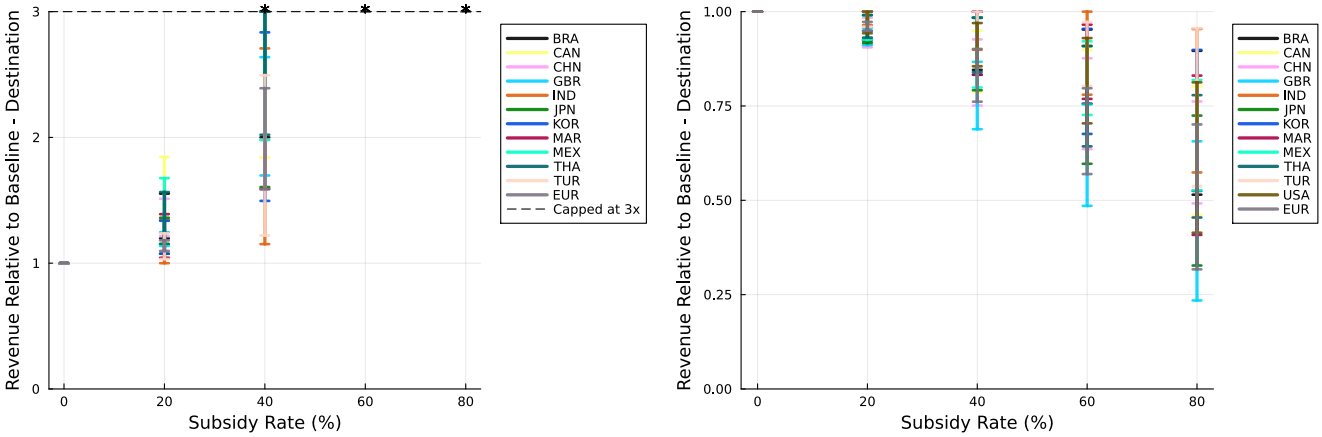


(a) A6: Product Entry Probability

(b) A4: Product Entry Probability

Note: These plots show the product offering probability for the Audi wagon products in different countries/regions at different U.S. marginal cost subsidies that affect the A4 and not the A6. The bounds reflect both statistical uncertainty regarding the bounds on the fixed cost parameters and uncertainty about the solution arising from potential inability to solve the full brand problem. "EUR" denotes the European Union.

Figure 21: Revenues Relative to Baseline at the Destination - Marginal Cost Subsidy on the A4

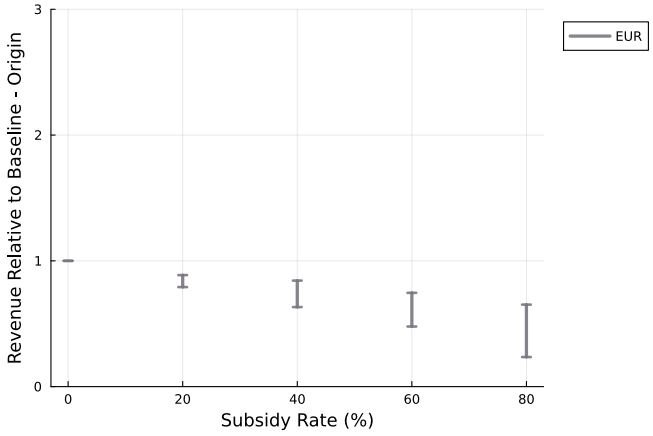


(a) A4: Revenues Relative to Baseline

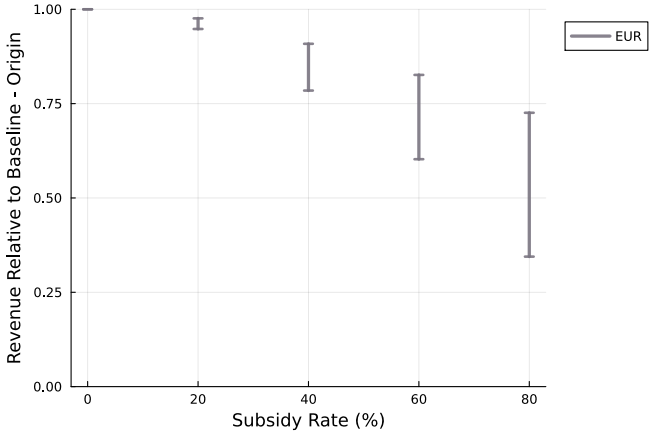
(b) A6: Revenues Relative to Baseline

Note: These plots show the revenues from the Audi wagon products in different countries/regions at different U.S. marginal cost subsidies that affect the A4 and not the A6, relative to the revenues earned at the baseline case. The bounds reflect both statistical uncertainty regarding the bounds on the fixed cost parameters and uncertainty about the solution arising from potential inability to solve the full brand problem. "EUR" denotes the European Union.

Figure 22: Revenues Relative to Baseline at the Origin - Marginal Cost Subsidy on the A4



(a) A4: Revenues Relative to Baseline



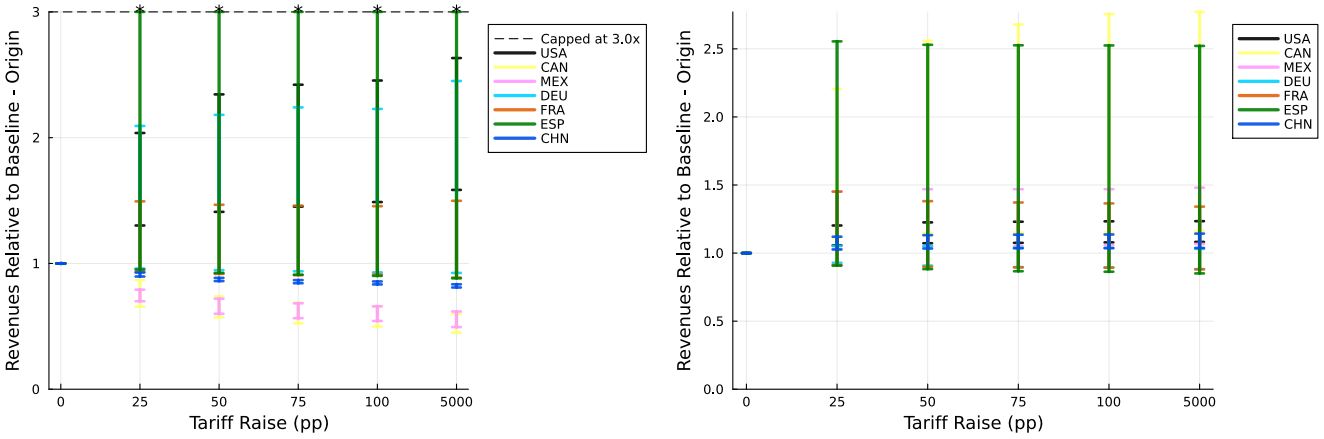
(b) A6: Revenues Relative to Baseline

Note: These plots show the revenues from the Audi wagon products in different countries/regions at different U.S. marginal cost subsidies that affect the A4 and not the A6, relative to the revenues earned at the baseline case. The bounds reflect both statistical uncertainty regarding the bounds on the fixed cost parameters and uncertainty about the solution arising from potential inability to solve the full brand problem. "EUR" denotes the European Union.

E.2 Average Counterfactual Exercises

E.2.1 Additional Plots: Tariff Counterfactual Exercises

Figure 23: Electric Revenues Relative to Baseline at Origin - Tariffs

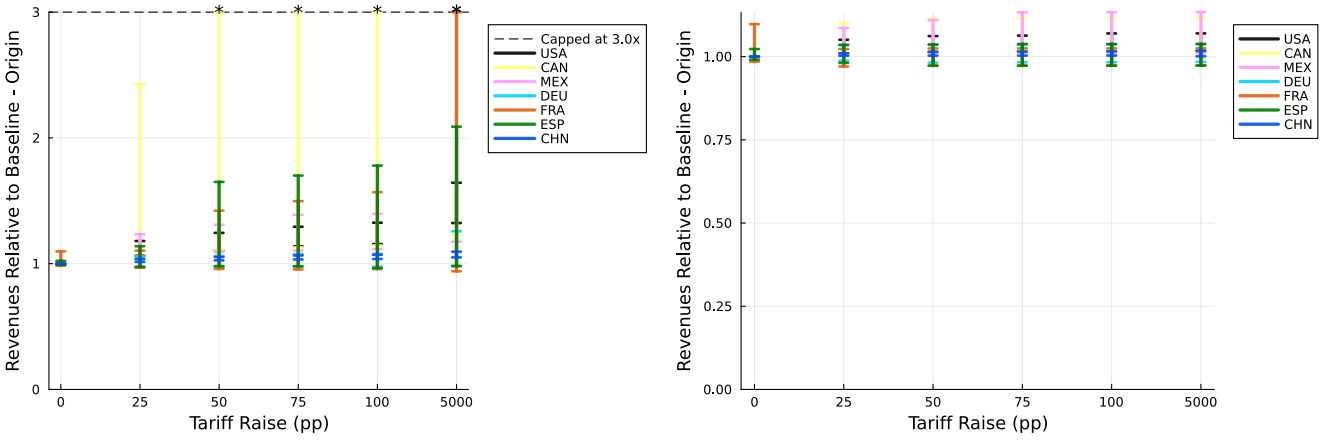


(a) Tariffs on the ROW

(b) Tariffs on the European Union

Note: These plots display bounds on the average revenue per electric product across brand-segments with at least one electric product in their portfolio relative to baseline in different origin countries. Panel (a) shows the bounds at different levels of U.S. tariffs on the rest of the world on electric products. Panel (b) displays the effect of tariffs imposed only on electric products produced in the European Union. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

Figure 24: Non-Electric Revenues Relative to Baseline at Origin - Tariffs

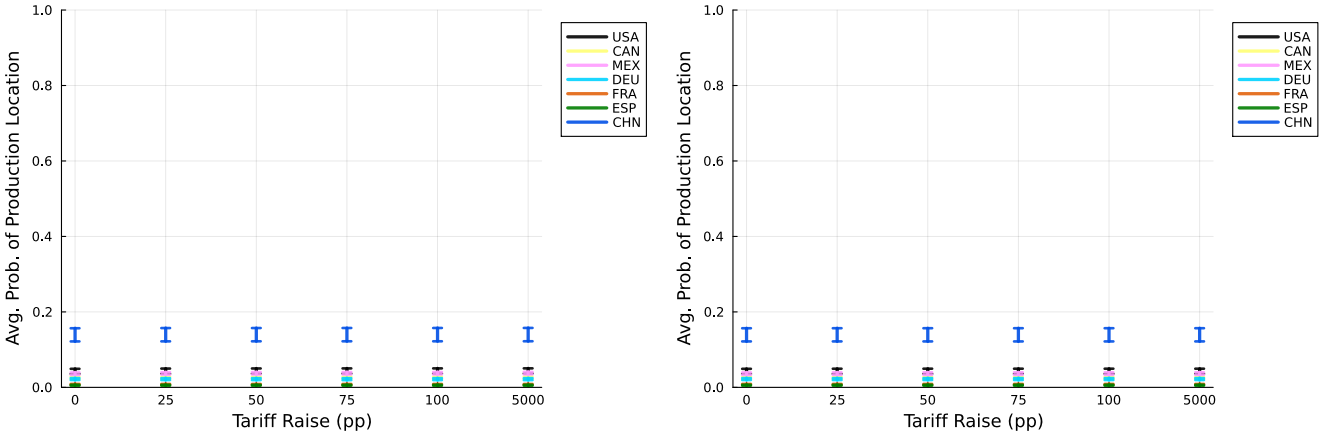


(a) Tariffs on the ROW

(b) Tariffs on the European Union

Note: These plots display bounds on the average revenue per non-electric product across brand-segments with at least one electric product in their portfolio relative to baseline in different origin countries. Panel (a) shows the bounds at different levels of U.S. tariffs on the rest of the world on electric products. Panel (b) displays the effect of tariffs imposed only on electric products produced in the European Union. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

Figure 25: Probability of Production Location of Non-Electric Products - Tariffs

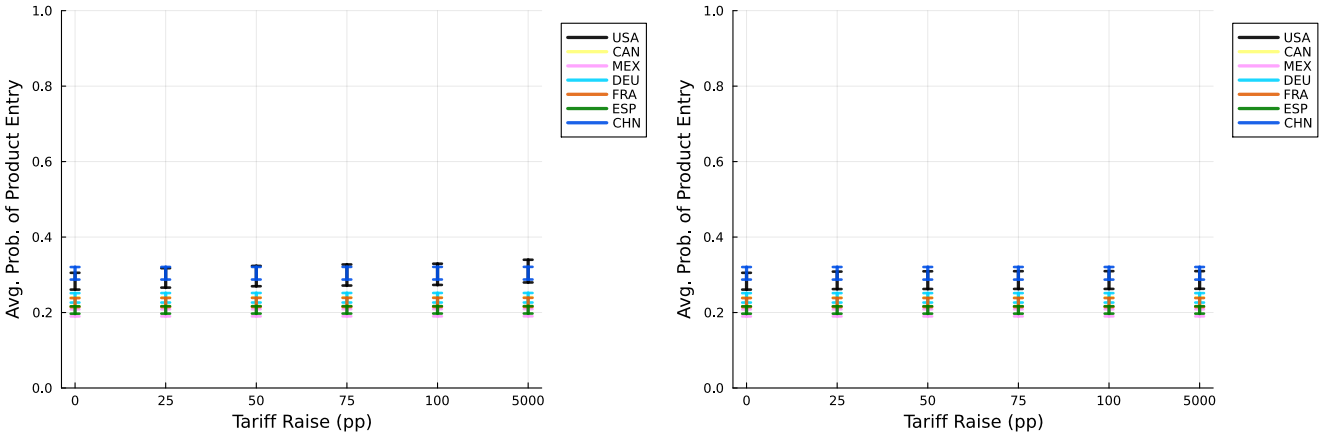


(a) Tariffs on the ROW

(b) Tariffs on the European Union

Note: These plots display bounds on the average share of non-electric products that are produced in different countries by brand-segments with at least one electric product in their portfolio. Panel (a) shows the bounds at different levels of U.S. tariffs on the rest of the world on electric products. Panel (b) displays the effect of tariffs imposed only on electric products produced in the European Union. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

Figure 26: Probability of Product Offering of Non-Electric Products - Tariffs

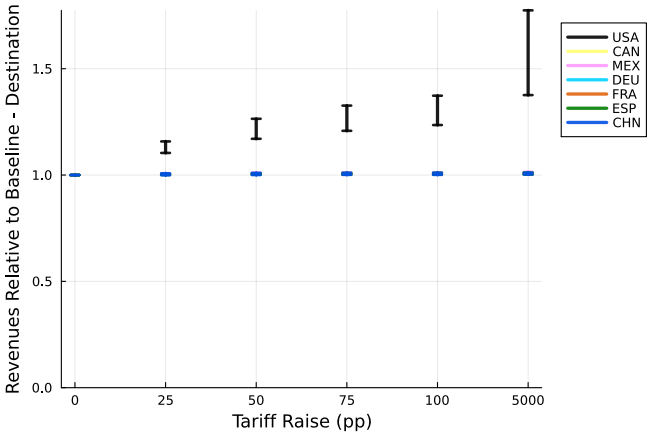


(a) Tariffs on the ROW

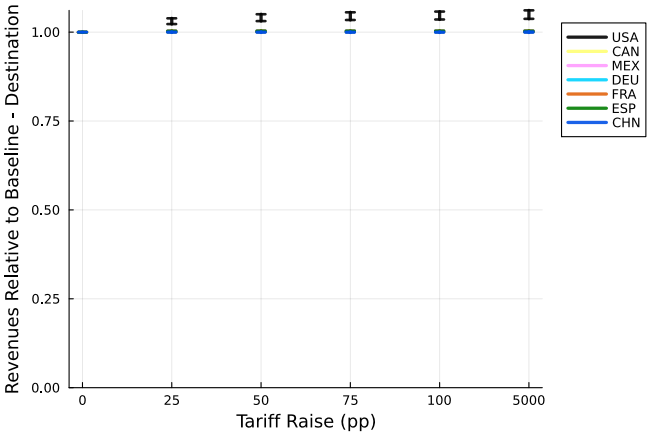
(b) Tariffs on the European Union

Note: These plots display bounds on the average share of non-electric products that are offered in different countries by brand-segments with at least one electric product in their portfolio. Panel (a) shows the bounds at different levels of U.S. tariffs on the rest of the world on electric products. Panel (b) displays the effect of tariffs imposed only on electric products produced in the European Union. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

Figure 27: Non-Electric Revenues Relative to Baseline at Destination - Tariffs



(a) Tariffs on the ROW

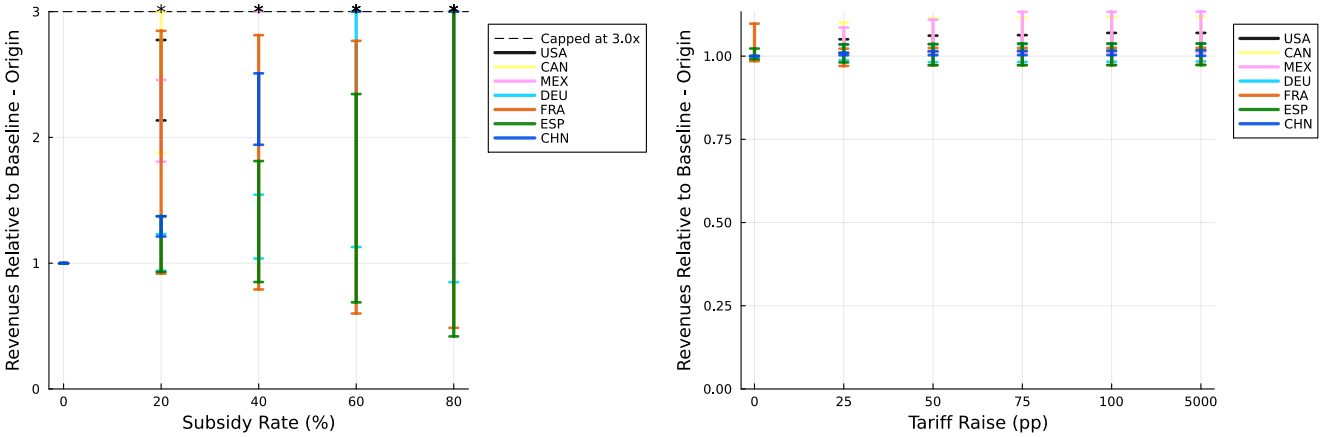


(b) Tariffs on the European Union

Note: These plots display bounds on the average revenue per non-electric product across brand-segments with at least one electric product in their portfolio relative to baseline in different destination countries. Panel (a) shows the bounds at different levels of U.S. tariffs on the rest of the world on electric products. Panel (b) displays the effect of tariffs imposed only on electric products produced in the European Union. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

E.2.2 Additional Plots: U.S. Electric-Vehicle Consumer Subsidy

Figure 28: Revenues Relative to Baseline at Origin - Consumer Subsidy

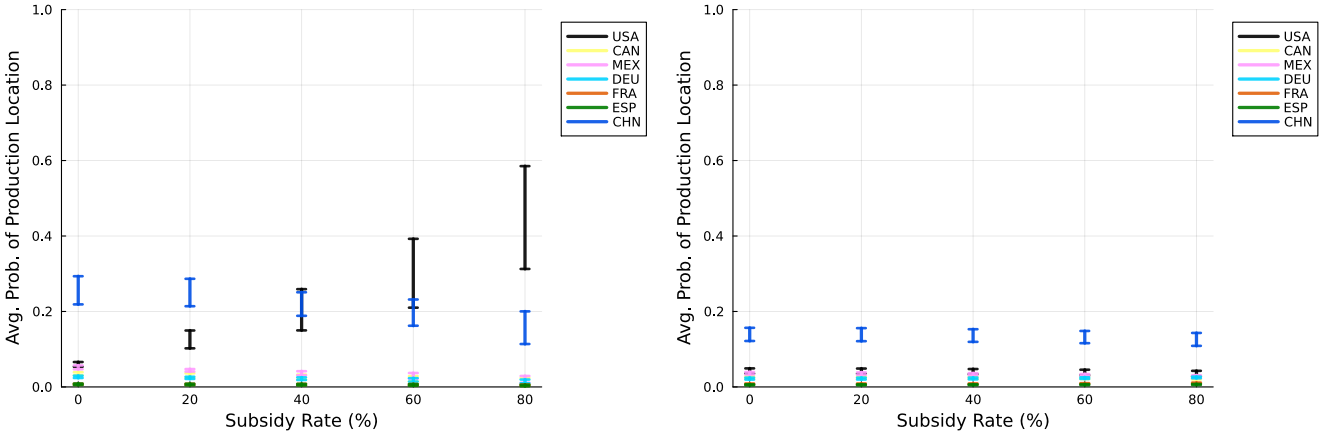


(a) Electric (b) Non-Electric

Note: These plots display bounds on the average revenue per product across brand-segments with at least one electric product in their portfolio relative to baseline in different destination countries. Panel (a) shows the bounds for electric products at different levels of U.S. consumer subsidies on electric products. Panel (b) displays the effect on non-electric products. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

E.2.3 Additional Plots: U.S. Electric-Vehicle Marginal Cost Subsidy

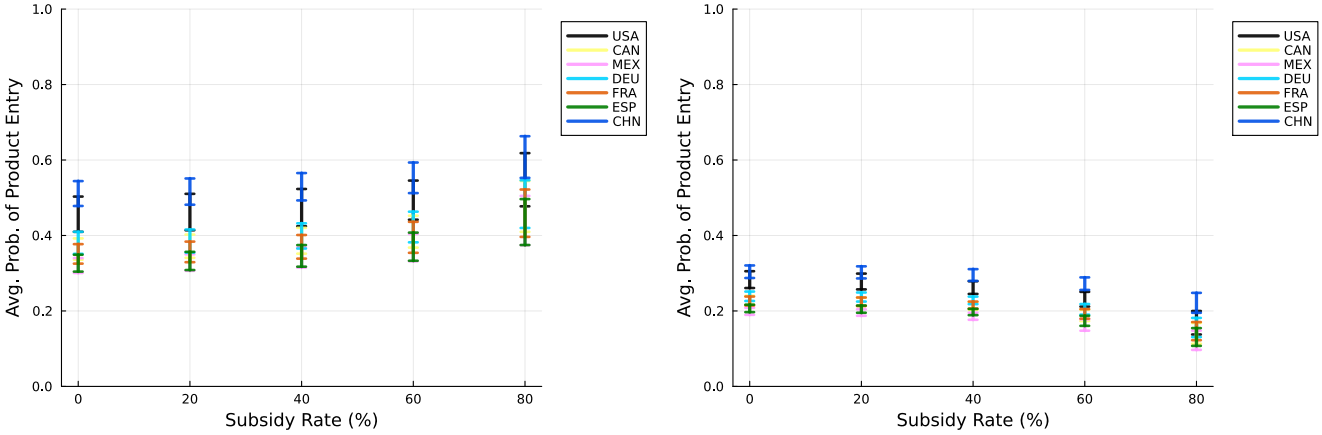
Figure 29: Probability of Production Location - Marginal Cost Subsidy



(a) Electric (b) Non-Electric

Notes: These plots show the average lower and upper bound share of electric (panel (a)) and non-electric (panel (b)) products produced by brands in several countries/regions under different US electric-vehicle marginal cost subsidy rates. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

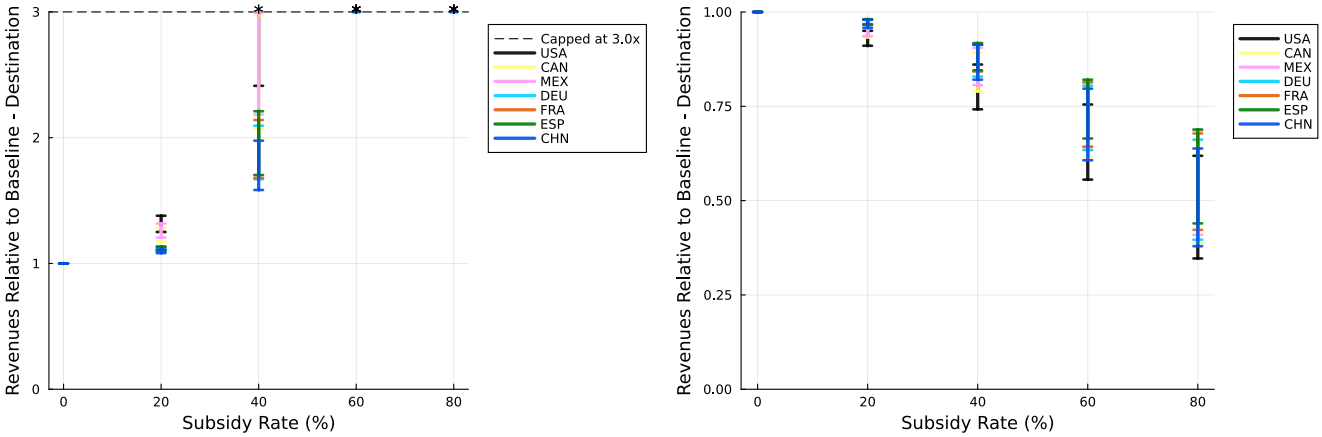
Figure 30: Probability of Product Offering - Marginal Cost Subsidy



(a) Electric (b) Non-Electric

Notes: These plots show the average lower and upper bound share of electric (panel (a)) and non-electric (panel (b)) products offered by brands in several countries/regions under different US electric-vehicle marginal cost subsidy rates. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

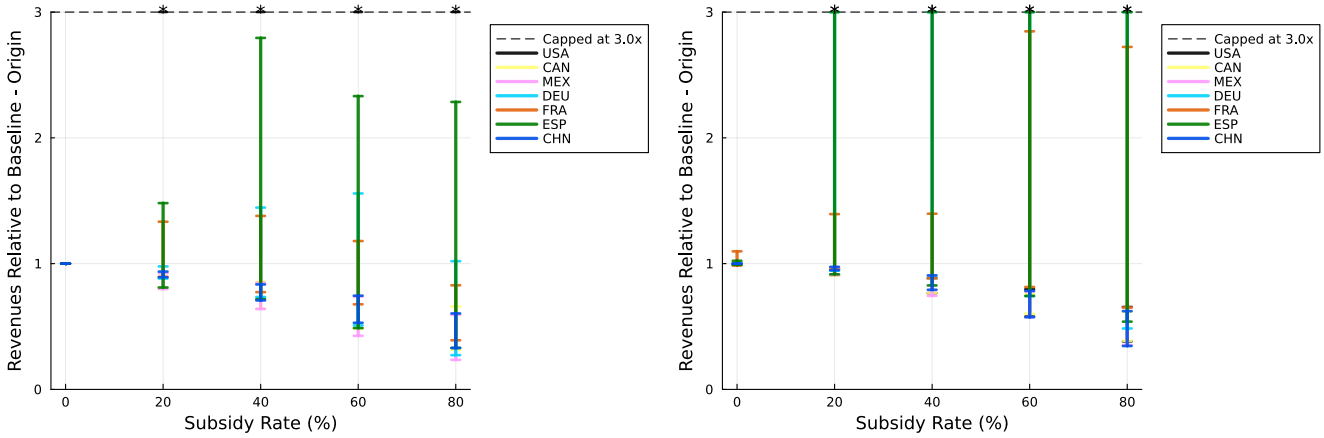
Figure 31: Revenues Relative to Baseline at Destination - Marginal Cost Subsidy



(a) Electric (b) Non-Electric

Note: These plots display bounds on the average revenue per product across brand-segments with at least one electric product in their portfolio relative to baseline in different destination countries. Panel (a) shows the bounds at different levels of U.S. marginal cost subsidies on electric products. Panel (b) displays the bounds at different levels of U.S. consumer subsidies on non-electric products. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

Figure 32: Revenues Relative to Baseline at Origin - Marginal Cost Subsidy

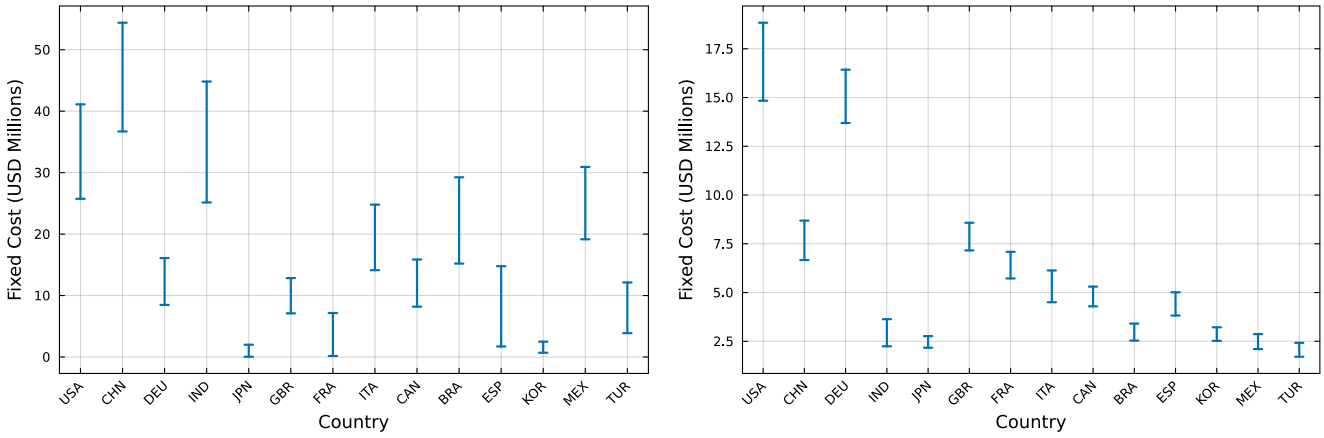


(a) Electric (b) Non-Electric

Note: These plots display bounds on the average revenue per product across brand-segments with at least one electric product in their portfolio relative to baseline in different destination countries. Panel (a) shows the bounds for electric products at different levels of U.S. marginal cost subsidies on electric products. Panel (b) displays the effect on non-electric products. The bounds reflect both uncertainty arising from unsolved binary decisions due to computational infeasibility, as well as from statistical uncertainty regarding the fixed cost parameter vector.

F Realized Fixed Costs

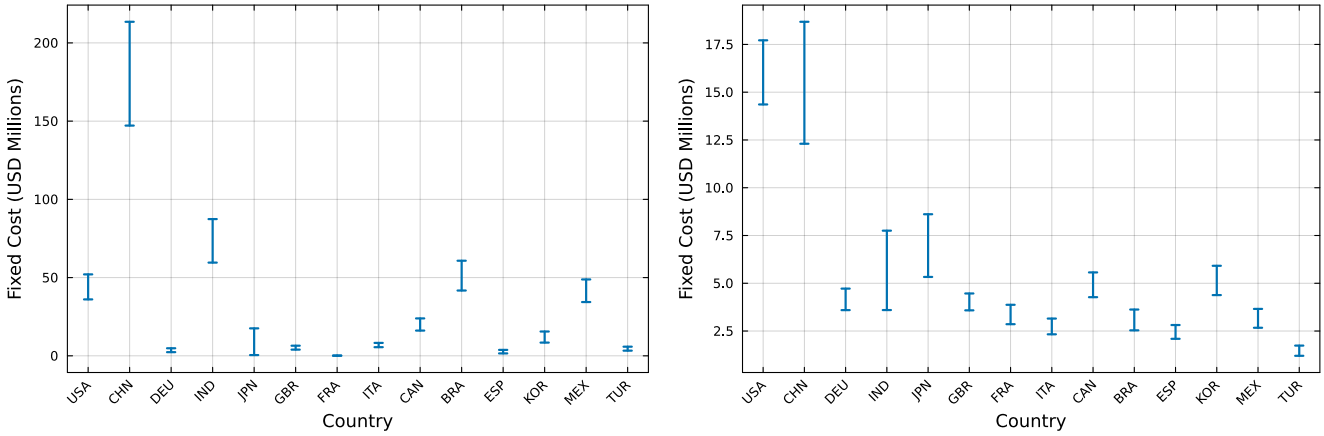
Figure 33: Fixed Costs Conditional Producing/Selling for Brands with German HQs



(a) Production Fixed Costs (b) Selling Fixed Costs

Notes: These plots show bounds on the average fixed costs across 128 fixed cost draws and across products sold by brands with headquarters in the United States, conditional on production (Panel (a)) and selling (Panel (b)). The plots exclude zero fixed cost origins and destinations due to the free plant and free destination lottery. The bounds reflect both statistical uncertainty regarding the value of the fixed cost parameters as well as potential uncertainty about the solution due to computational infeasibility. We condition on (sure) production- locations for which the solution algorithm solves for a 1.

Figure 34: Fixed Costs Conditional Producing/Selling for Brands with Japanese HQs



(a) Production Fixed Costs

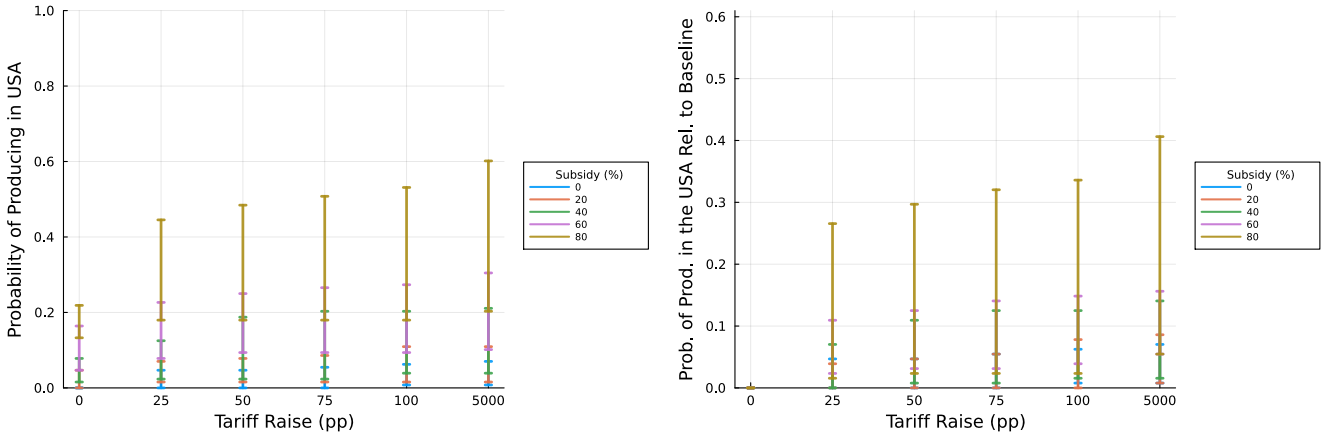
(b) Selling Fixed Costs

Notes: These plots show bounds on the average fixed costs across 128 fixed cost draws and across products sold by brands with headquarters in the United States, conditional on production (Panel (a)) and selling (Panel (b)). The plots exclude zero fixed cost origins and destinations due to the free plant and free destination lottery. The bounds reflect both statistical uncertainty regarding the value of the fixed cost parameters as well as potential uncertainty about the solution due to computational infeasibility. We condition on (sure) production- locations for which the solution algorithm solves for a 1.

G Tariffs under Counterfactual Worlds

G.1 Audi - Confidence Intervals

Figure 35: U.S. Tariffs on A4 at Different Demand Subsidy Rates

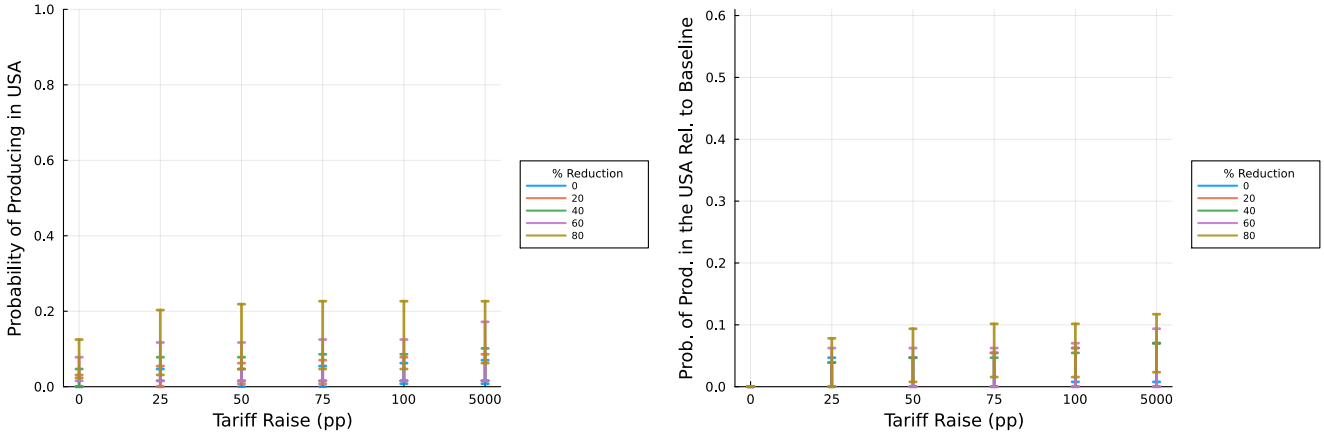


(a) Probability Level

(b) Probability Relative to Baseline

Notes: These plots show bounds on the probability of producing the A4 in the United States at different demand subsidy rates. The bounds reflect statistical uncertainty about the fixed-cost parameter values as well as uncertainty regarding the solution to the brand's problem. Panel (a) shows the bounds on the probability of producing in the United States, while Panel (b) displays bounds on the difference in the probability of producing in the United States relative to baseline (zero tariff raise).

Figure 36: U.S. Tariffs on A4 at Different Fixed U.S. Production Subsidy Rates

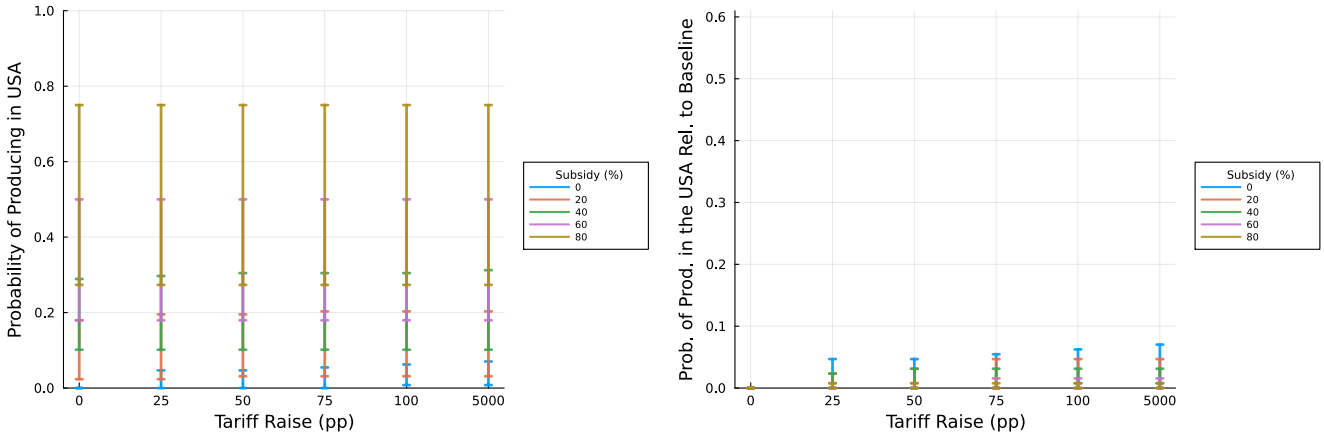


(a) Probability Level

(b) Probability Relative to Baseline

Notes: These plots show bounds on the probability of producing the A4 in the United States at different U.S. fixed production cost subsidy rates. The bounds reflect statistical uncertainty about the fixed-cost parameter values as well as uncertainty regarding the solution to the brand's problem. Panel (a) shows the bounds on the probability of producing in the United States, while Panel (b) displays bounds on the difference in the probability of producing in the United States relative to baseline (zero tariff raise).

Figure 37: U.S. Tariffs on A4 at Different Marginal Cost Subsidy Rates



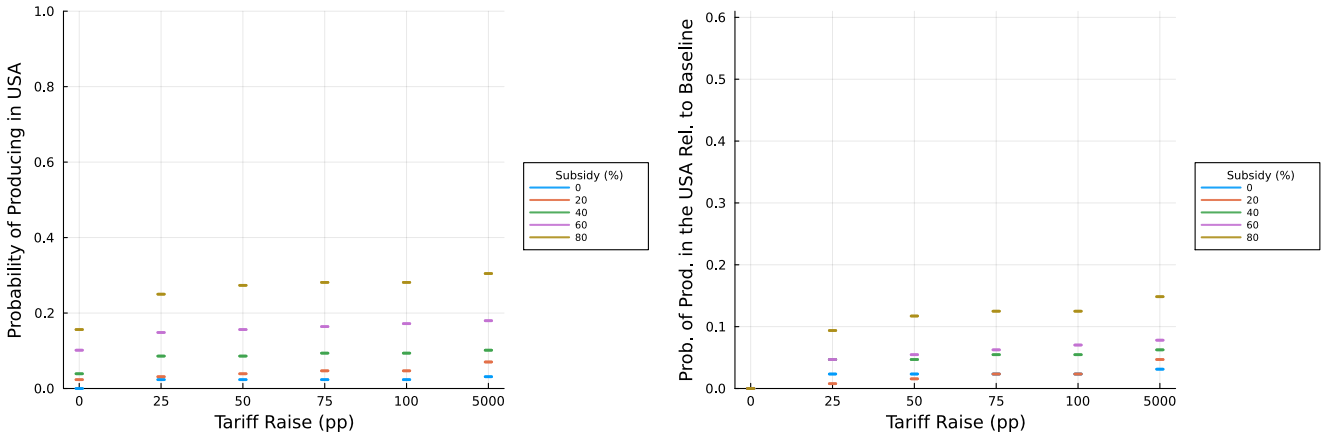
(a) Probability Level

(b) Probability Relative to Baseline

Notes: These plots show bounds on the probability of producing the A4 in the United States at different marginal cost subsidy rates. The bounds reflect statistical uncertainty about the fixed-cost parameter values as well as uncertainty regarding the solution to the brand's problem. Panel (a) shows the bounds on the probability of producing in the United States, while Panel (b) displays bounds on the difference in the probability of producing in the United States relative to baseline (zero tariff raise).

G.2 Audi - Single Parameter Vector

Figure 38: U.S. Tariffs on A4 at Different Demand Subsidy Rates

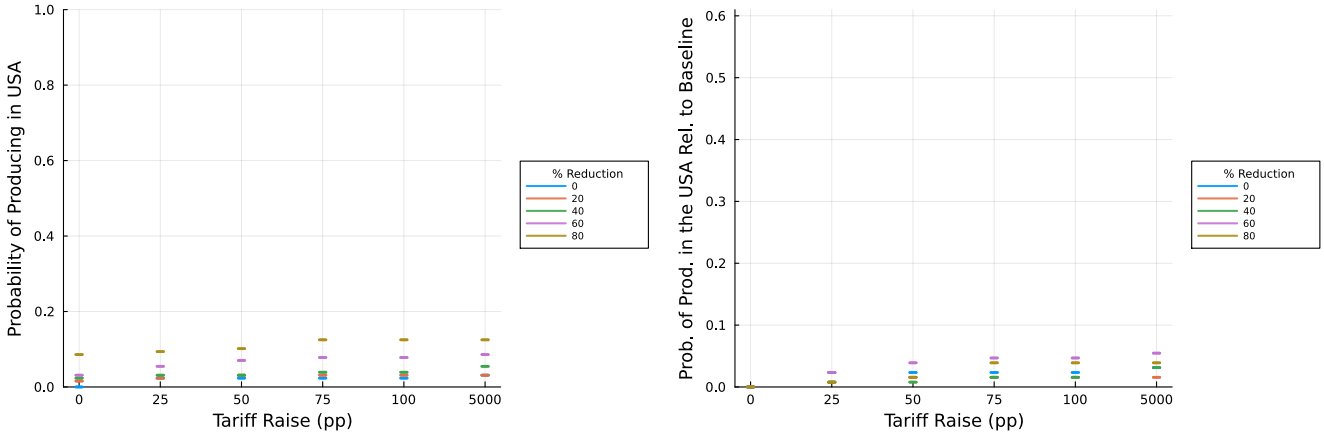


(a) Probability Level

(b) Probability Relative to Baseline

Notes: These plots show bounds on the probability of producing the A4 in the United States at different demand subsidy rates. The bounds reflect uncertainty regarding the solution to the brand's problem. Panel (a) shows the bounds on the probability of producing in the United States, while Panel (b) displays bounds on the difference in the probability of producing in the United States relative to baseline (zero tariff raise).

Figure 39: U.S. Tariffs on A4 at Different Fixed U.S. Production Subsidy Rates

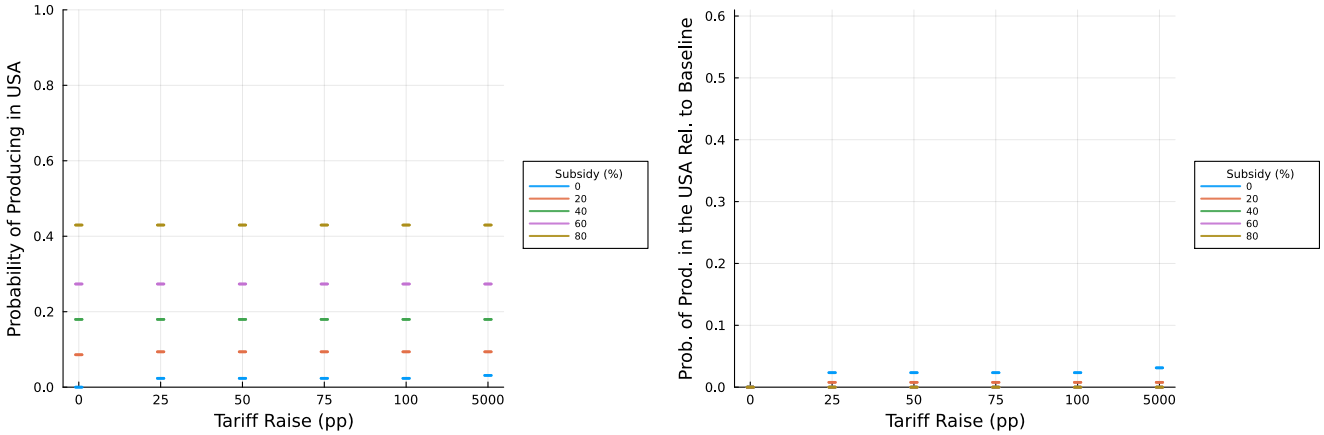


(a) Probability Level

(b) Probability Relative to Baseline

Notes: These plots show bounds on the probability of producing the A4 in the United States at different U.S. fixed production cost subsidy rates. The bounds reflect uncertainty regarding the solution to the brand's problem. Panel (a) shows the bounds on the probability of producing in the United States, while Panel (b) displays bounds on the difference in the probability of producing in the United States relative to baseline (zero tariff raise).

Figure 40: U.S. Tariffs on A4 at Different Marginal Cost Subsidy Rates



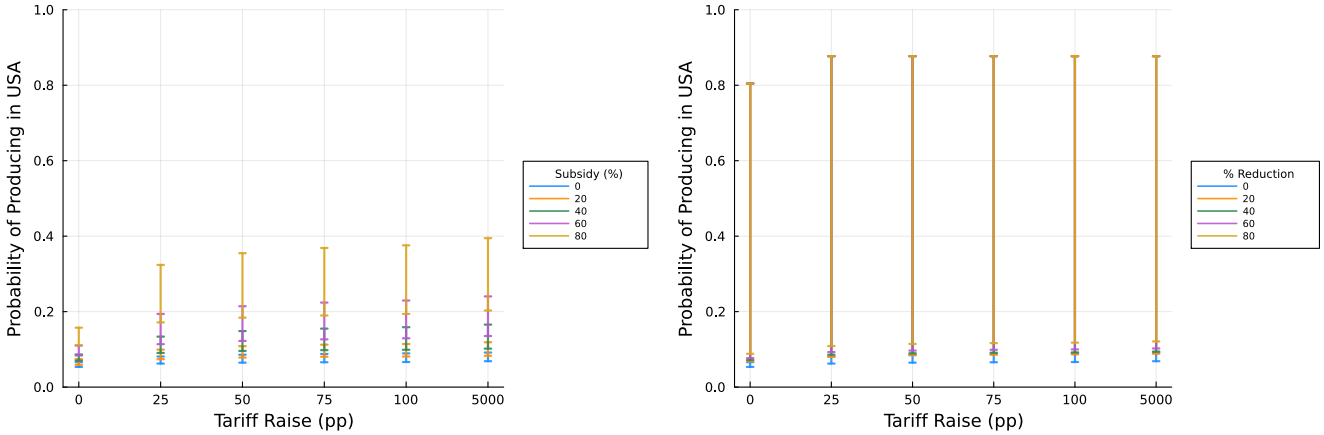
(a) Probability Level

(b) Probability Relative to Baseline

Notes: These plots show bounds on the probability of producing the A4 in the United States at different marginal cost subsidy rates. The bounds reflect uncertainty regarding the solution to the brand's problem. Panel (a) shows the bounds on the probability of producing in the United States, while Panel (b) displays bounds on the difference in the probability of producing in the United States relative to baseline (zero tariff raise).

G.3 Average Firm – Confidence Intervals

Figure 41: U.S. Electric-Vehicle Tariffs on the ROW



(a) Under E.V. Consumer Subsidies

(b) Under E.V. Fixed Production Cost Subsidies

Notes: These plots show bounds on the probability of producing electric vehicles in the United States for the average firm. The bounds reflect statistical uncertainty regarding the fixed cost parameter values, as well as uncertainty regarding the solution to the brand's problem. Panel (a) shows the bounds on the probability of producing in the United States under various levels of E.V. consumer subsidies, while Panel (b) shows the bounds under various levels of E.V. fixed production cost subsidies.