Tariffs and Production Relocation in a Multi-Product Oligopoly

Juanma Castro-Vincenzi
University of Chicago

Jin Liu
University of Pennsylvania

Eduardo Morales
Princeton University

Alejandro Sabal[*]
Yale University

June 1, 2026

PRELIMINARY AND INCOMPLETE.

Abstract

We build an oligopolistic model of production location and market entry decisions for multi-product and multinational firms. We develop a novel solution algorithm for entry games that are neither supermodular nor submodular. Our algorithm only requires that any pair of binary actions within or across players be either complements or substitutes, independently of third choices. Using data from the global SUV market, we estimate heteregeneous consumer preferences, trade costs, marginal costs of production, and fixed exporting and production costs. We quantify the role of endogenous production reallocation in shaping the equilibrium welfare effects of U.S. tariffs on SUVs, and find substantial global spillover effects due to the reorganization of production and offerings.

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

---

[*] We thank Elhanan Helpman, Michal Kolesár, Robin Lee, Marc Melitz, Ariel Pakes, and Katja Seim for insightful conversations, and seminar participants at Harvard, and Yale, for helpful comments. All errors are our own. Emails: Juanma Castro-Vincenzi (castrovincenzi@uchicago.edu), Jin Liu (jinliu52@upenn.edu), Eduardo Morales (ecmorale@princeton.edu), Alejandro Sabal (alejandro.sabal@yale.edu).

1 Introduction

Multinational firms often have substantial market shares, and produce and sell many products in multiple countries. Despite their global nature, such firms typically produce and sell their products in a relatively small subset of all potential countries. A standard explanation for this empirical fact is that firms must incur fixed costs both to produce and to distribute products worldwide. To predict how an oligopolistic industry responds to tariffs and other forms of policy intervention, it is therefore natural to account for the presence of fixed costs, as well as the potential interdependence of production and sales decisions across locations, products, and firms.

How do multinationals respond to a tariff levied by the United States on goods produced overseas? The answer is non-trivial. Even in the absence of any adjustment on production or sales decisions, a given firm may decide to reoptimize the quantities it supplies and the prices it charges in the U.S. market. The tax may also induce the exit of some of the least popular imported goods, as variable profits may now be insufficient to compensate fixed exporting costs. If a product was produced in multiple countries, the firm could also choose to source it from a country not subject to the new tariff. Or, potentially, the firm may decide to open a new plant for a product in a country, such as the United States, that is not subject to the tariff. The tariff may affect offerings and prices in other countries if the firm simultaneously decides to close a pre-existing plant for that product. Moreover, in the presence of substantial market shares, a given firm's (product's) reorganization of production and sales decisions has a direct impact on other firms' (product's) payoffs, and thus their decisions. A quantitative equilibrium model is needed to study how offerings and production decisions are determined in multinational oligopolies.

We develop a model and study this question in the context of the global SUV industry, which has recently been the target of ambitious trade and industrial policies. The model features a finite set of firms, each endowed with a finite portfolio of products, competing in a game with two stages. In the first stage, firms make production location and market offering decisions for each product, subject to fixed costs of producing and selling products in each country. In the second stage, firms compete à la Cournot and choose quantities for each offered product in each destination market. In the model, firms optimally source a product to a destination market from the available minimum-cost origin, accounting for both marginal production costs and trade costs. Firms also internalize that consumers substitute across available products. Minimum-cost sourcing and consumer demand naturally render production and sales decisions interdependent.

The first stage of our model features many interdependent binary decisions. In the model, minimum-cost sourcing implies that the production of a given product in any country weakly increases the returns of selling such a product anywhere. Minimum-cost sourcing also implies that producing a product in a country weakly decreases the value of producing it in any other country. Consumer substitution across products implies that selling and producing different products, both within and across firms, are always substitute decisions. Such an entry game is therefore neither supermodular nor submodular.

The key methodological contribution of our paper is to provide a solution procedure for many-binary-action complete-information games that simultaneously feature complementary and substitutable choices. Previous approaches to solving these games have focused on finding pure strategy Nash equilibria (PSNE). Jia (2008) exploits supermodularity to bound the PSNE. Fan and Yang (2020) deal instead with a submodular game and rely on the "greedy" algorithm to find (or approximate) one equilibrium. Two key challenges emerge in settings without supermodularity nor submodularity: (i) potential PSNE non-existence or multiplicity, and (ii) computational intractability in finding all the PSNE in settings with many binary actions.[1] Our approach is to depart from PSNE and focus instead on action profiles that survive iterated elimination of (pure) dominated strategies (IEDS), which contain all PSNE actions whenever they exist.

We develop a computationally feasible iterative solution algorithm that bounds the IEDS strategies in many-player and many-binary-action games. Our method allows for within-firm choice pairs that are complements or substitutes, as well as across-firm choice pairs that are strategic complements or strategic substitutes. The algorithm operates by iteratively discarding binary actions (zeros or ones) that are dominated, relaxing restrictions from previous methods in the literature (e.g. Jia 2008), which impose either super- or submodularity. Relative to Jia (2008), or best-response iteration in supermodular games more generally, our bounds are not sharp. However, we find highly informative bounds in practice.

We incorporate our solution method in a moment inequality procedure to estimate fixed-cost parameters. Our moment inequalities sharpen the bounds in Fan and Yang (2025), which only impose level-1 rationality, by iteratively discarding dominated actions, i.e., imposing level-∞ rationality. Given fixed-cost parameters, exogenous variables, and simulated fixed-cost shocks, our solution algorithm yields bounds on sales and production location decisions. We follow Castro-Vincenzi et al. (2023) and show that we can use moments that are monotone in such choices to estimate fixed-cost parameters in a simulated moment inequality procedure.

We estimate our model using information from *IHS Markit* (S&P Global) on the global SUV market in 2019. We combine data on new global SUV registrations, with information on quantities and production origins for products offered in 56 destination countries, with data on product prices and characteristics for a representative sample of 12 countries. We supplement our sample with information on gravity variables from CEPII, MacMap data on car tariffs, and World Bank data on country demographics (e.g., GDP per capita and population).

Using these data, we estimate our model in two steps. First, given observed prices and quantity shares, we estimate a multinomial logit demand model using two-stage least squares, instrumenting price with cost shifters (geographic distance and tariffs). Given our Cournot conduct assumption, we solve for markups and obtain marginal costs, which we project on fixed effects and observable product characteristics. In the second step, we use our moment inequality procedure and our second-stage model and estimates to recover fixed-cost parameters. We simulate our moment inequalities by using our demand, marginal cost, and the Cournot conduct assumption to predict

---

[1] Already with 30 binary actions in total, there are over 1 billion possible entry configurations, rendering brute-force procedures infeasible.

variable profits for products in countries in which they are not sold in our sample.

Our demand estimates reveal heterogeneity in preferences for horsepower and size across the country income distribution, as well as a substantial bias for brands with local headquarters, as in Coşar et al. (2018). Our marginal-cost estimates reflect gravity and heterogeneity in marginal costs of production across countries. SUV brands find it cheaper to sell products that are produced locally, and costlier to export products to geographically distant markets. A 1% increase in distance leads to a 0.037% increase in trade costs. Moreover, brands find it costlier to produce in countries that are either distant or impose larger part-tariffs on their headquarters country. Finally, we estimate higher costs of producing more powerful, heavier, and larger vehicles, all else equal.

We assume that both fixed costs of production and selling are log-normally distributed. Our estimates of fixed production and selling costs align with the estimates Castro-Vincenzi et al. (2023): fixed production costs are substantially larger than fixed selling costs. The median of the distribution of fixed selling costs per product-country lies at $20 million, and the median of the fixed production cost distribution per product-country is at least $1,096 billion. We estimate that the fixed cost distributions feature substantial dispersion, and that brands are more likely to have production plants in their headquarters country.

We use the estimated model to quantify the welfare effects of a counterfactual U.S. policy that imposes an additional 100 percentage-point ad-valorem tariff on all foreign-produced SUVs. When firms are allowed to fully reoptimize production and sales decisions, the tariff induces substantial reshoring: the share of U.S.-sold products sourced from U.S. plants rises by 13.3–20.2 percentage points. Although U.S. consumer surplus falls by $302–$569 per household as prices rise and product variety shrinks, the resulting gains in variable profits attributable to U.S. plants and in labor income more than offset these losses. Consequently, the net effect on U.S. households is unambiguously positive under the assumption that U.S. households receive all profits generated by U.S. plants.

To isolate the role of production relocation, we also consider a scenario in which firms are contrained to keep production locations fixed. In this case, the reshoring response shrinks to 3.8–17.8 percentage points, and the gains in profits and labor income decline sharply. Although tariff revenue rises substantially, from $38–$65 to $198–$344 per household, the overall welfare effect becomes ambiguous.

We then consider a second counterfactual in which the same tariff increase applies only to E.U. produced SUVs. The direct consumer impact is much smaller, because E.U. produced vehicles account for only about one fifth of models sold in the U.S. market and firms can relocate production to non-tariffed countries. Consumer prices rise by only $200–$600 per vehicle, compared with $3,800–$6,800 under the global tariff. At the same time, much of the affected production relocates to the U.S., causing tariff revenue to fall by $13–$45 per household.

We also find that countries not directly targeted by the tariff could experience meaningful spillovers. In particular, Canada potentially benefits as a third-country destination, with Canadian plants absorbing part of the demand displaced from E.U. producers. Comparing the two counterfactual experiments, a universal tariff generates larger reshoring responses and more robust gains for U.S. households, whereas a discriminatory tariff provides an escape valve that mitigates consumer harm but also weakens reshoring incentives.

Our quantitative results are subject to a few caveats worth mentioning. First, our model is static and our tariff counterfactuals do not incorporate policy uncertainty. In this sense, our results likely favor production relocation to the United States in response to tariffs, which in reality might be perceived as temporary or subject to future changes. Second, our baseline counterfactual exercises consist of unilateral tariffs, which also likely overstate the welfare benefits for U.S. households relative to a scenario with retaliation. Finally, our current specifications of fixed costs do not allow for economies of scope (i.e., fixed cost reductions when selling or producing multiple products in a market). By ignoring such fixed-cost complementarities, we likely underpredict production relocation to the United States.

Our paper relates to several literatures. First, it contributes to quantitative work on multinational production and export platforms, including Ramondo and Rodríguez-Clare (2013), Tintelnot (2017), Arkolakis et al. (2023), and Head and Mayer (2019). Within this literature, it is closest to papers that use the automobile industry to study global production, market access, and counterfactual exercises (e.g. Coşar et al. 2018, Head and Mayer 2019, Castro-Vincenzi 2024, Head et al. 2025). The paper is closest to Castro-Vincenzi et al. (2023), whose model also features joint production-location and product-offering decisions by multi-product firms. We build on their lattice-based algorithm for high-dimensional binary choices with both complements and substitutes, and extend it to entry games in which the sign of each pairwise interdependence is known and invariant to other own and rival choices.

Second, the paper contributes to the econometric literature on entry games and incomplete models. This literature uses observed discrete choices to infer payoff primitives and competitive effects while allowing for equilibrium multiplicity (Bresnahan and Reiss, 1991; Ciliberto and Tamer, 2009; de Paula, 2013). Recent papers extend these ideas to settings with many firms or many product decisions (Fan and Yang, 2020, 2025; Sabal, 2025). We differ from such papers by combining product-market entry with production-location choices, and by departing from Nash equilibrium and instead relying only on rationalizability in a complete-information setting.

Third, the paper is related to computational methods for interdependent discrete choices. Prior work develops methods for large entry or combinatorial-choice problems in which economic structure restricts the pattern of interactions across binary decisions (Jia, 2008; Arkolakis et al., 2023; Castro-Vincenzi et al., 2023; Head et al., 2025). Our extension is to a strategic setting with many multi-product firms, where pairwise interactions may be complementary or substitutable within and across firms. Because pure-strategy Nash equilibria need not be unique and are generally infeasible to enumerate, we impose common knowledge of rationality, following Bernheim (1984) and Pearce (1984), and compute bounds on action profiles that survive iterative deletion of actions that are never best responses.

The tariff counterfactuals relate to strategic-trade-policy models of oligopolistic industries (as surveyed by Brander 1995). Although our paper does not consider optimal tariff setting, it does explore how tariffs affect rent allocation when firms internalize rivals' output or pricing responses. In particular, our paper explores how tariffs can change which plants are active and which products are offered in each destination, and how these interact with the firms' strategic behavior. We therefore compare tariff responses under endogenous production relocation with responses obtained when production locations are held fixed.

2 Data

Our analysis uses information on global SUV sales, production, prices, product characteristics, trade costs, and destination-country demographics. The core sales and production data correspond to 2019 and are compiled by *IHS Markit* (S&P Global) from new-vehicle registration records. These data identify, for each car product, the country in which the vehicle is assembled and the country in which it is sold. We define a product as a parent–brand–nameplate combination, and in this paper we restrict attention to SUVs. Throughout the empirical analysis, the firm is the parent company. For example, the Hyundai parent company owns the Hyundai and Kia brands, and the Tucson and Sportage are separate SUV products in Hyundai's portfolio.

Our dataset contains information on differentiated SUVs that are offered in many, but not all, countries and that are often produced in more than one location. Table 1 reports summary statistics on product portfolios, sales markets, production locations, and international sourcing. The average parent firm owns 9.3 SUV products, while the quantity-weighted average is 20.6 products, reflecting the fact that high-volume firms have broader portfolios. The average SUV product is sold in 12.3 countries and produced in 1.4 countries. There is substantial international sourcing: across model-origin pairs, the unweighted share exported is 32.1 percent and the quantity-weighted share exported is 44.7 percent. At the same time, home-market production and sales remain important. The average parent produces 75.3 percent of its output in its headquarters country, although the quantity-weighted home-production share is lower, at 49.0 percent. The analogous shares for sales in the headquarters country are 67.4 percent and 34.2 percent.

Table 1: Summary Statistics

Metric	Mean (unw.)	Mean (w.)	p25	p50	p75	p90	Max
Models (parent)	9.3	20.6	2	7	13	24	34
Sales countries (model)	12.3	28.4	1	2	13	48	56
Production countries (model)	1.4	2.2	1	1	1	2	6
Models sold (market)	102.3	185.5	88	102	115	124	335
Sales countries (parent)	19.4	46.0	1	5	51	56	56
Prod. countries (parent)	2.9	7.2	1	1	4	8	12
Share exported (model-origin)	32.1%	44.7%	0.0%	1.5%	79.7%	100.0%	100.0%
HQ production share (parent)	75.3%	49.0%	50.4%	100.0%	100.0%	100.0%	100.0%
HQ sales share (parent)	67.4%	34.2%	23.7%	95.9%	100.0%	100.0%	100.0%

Notes: Weighted moments use quantities as weights. Market refers to countries where the good is sold. Model-origin refers to countries where the good is produced. Parent refers to our definition of firm. Home market production is defined using the headquarters' country of the parent.

Table 2 describes market concentration in sales markets. The average destination market has 102 SUV products, 17 parent firms, and 34 brands. Nevertheless, market shares are highly concen-

Table 2: Market Competition Among SUV Sellers

Metric	Mean (unw.)	Mean (w.)	p25	p50	p75	p90	Max
Models sold (market)	102.3	185.5	88	102	115	124	335
Parent firms (market)	17.4	26.5	15	16	18	22	48
Brands (market)	33.8	53.0	32	33	36	39	95
Top parent share (market)	27.8%	21.0%	22.1%	25.1%	32.4%	38.3%	71.3%
Top 3 parent share (market)	58.8%	45.6%	52.1%	59.4%	66.2%	72.8%	84.9%
Top model share (market)	11.5%	6.2%	6.9%	8.8%	13.6%	19.8%	38.5%
Top 5 model share (market)	34.8%	22.8%	25.1%	31.8%	44.2%	55.0%	69.9%

Notes: The table reports destination-market summary statistics for the 2019 SUV sample. Weighted moments use market-level quantities as weights.

trated. The largest parent firm in a market accounts for 27.8 percent of SUV sales on average, and the top three parent firms account for 58.8 percent. Concentration is also visible at the product level: the top model has an average market share of 11.5 percent, and the top five models account for 34.8 percent. These facts motivate the oligopolistic model in Section 3, where firms internalize the impact of their own quantities on prices and where each firm's actions affect rival firms' payoffs.

In addition to global production and sales, we use product-level price and characteristics data for Australia, Brazil, China, France, Germany, India, Italy, Japan, Mexico, Spain, the United Kingdom, and the United States. The price variable is the manufacturer-suggested retail price (MSRP). The characteristics used in estimation include horsepower, gross vehicle weight, and size, defined as width times length. We aggregate price and characteristics data to quantity-weighted annual product–market–origin observations and use the global production data to identify the assembly location serving each destination.

We measure trade costs using bilateral applied final-good tariffs and geographic distance between origin and destination countries. In the marginal-cost projection we also use bilateral auto-parts tariffs between the brand headquarters country and the production origin. Geographic distances come from CEPII. We use World Bank data on destination-country demographics, including GDP per capita and population, to construct demand shifters and market sizes. Following Grieco et al. (2024), we define the market size in destination n as proportional to the number of households and the average number of SUVs per household. Let H_n denote the number of households, v_n vehicles per household, and ω_n^{SUV} the SUV share among 2019 vehicle quantities destined for n. The SUV market size is

$$M_n = 0.2 H_n v_n \, \omega_n^{SUV}.$$

The factor 0.2 converts the household vehicle stock into an annual market under a five-year ownership cycle, while ω_n^{SUV} scales this market to SUVs.[2] The resulting outside option captures choices outside the modeled new-SUV choice set and pins down the logit shares used in demand estimation.

---

[2] When household-level vehicle ownership is unavailable, we compute v_n from vehicles per capita and average household size.

3 Strategic Model of Plant Location and Market Entry

We develop a game-theoretic model of plant location and market entry decisions for multi-product firms. Firm f, endowed with a portfolio of differentiated SUV products, Ω_f, decides in which potential countries, $\mathcal{N} = \{1, 2, ..., n, ..., N\}$, to produce and sell each of its products. A finite set of firms $\mathcal{F}$ make plant and market offering decisions for each of their products simultaneously under common knowledge of rationality, i.e., best-response behavior.

Firms make decisions in two stages. In Stage 1, each firm f draws for each product $m \in \Omega^f$ a country in which it can produce and sell at zero fixed costs, a set of plant location fixed cost shocks $\{\nu^p_{mo}\}_{m\in\Omega^f,o}$, and set of market entry fixed cost shocks $\{\nu^e_{mn}\}_{m\in\Omega_f,n}$.[3] Upon observing the realizations of these fixed cost shocks (unobserved by the econometrician), firms simultaneously choose in which countries to produce and in which countries to sell each of their products. In Stage 2, firms realize a vector of demand shocks $\{\xi_{mn}\}_{m\in\Omega_n,n}$ for all products offered in country n, Ω_n, and simultaneously choose the quantity supplied of each product in each destination country. We solve the game by backward induction and describe the two stages in detail.

3.1 Stage 2: Demand, Marginal Costs, and Quantity-Setting

We begin by introducing the features of the model conditional on a vector of first-stage production and market-entry decisions. Let

$$\Omega_{fn} \equiv \{m \in \Omega_f : I_{mn} = 1\} \tag{1}$$

denote the set of products that firm f offers in destination country n, and let $\Omega_n = \cup_{f\in\mathcal{F}}\Omega_{fn}$ denote the set of all products offered in n. The outside option has mean utility normalized to zero. For each inside product $m \in \Omega_{fn}$, consumers in destination n have mean utility,

$$u_{mn} = \delta_{mn} + \xi_{mn} - \alpha_n p_{mn}, \tag{2}$$

where δ_{mn} is an observed component of demand, ξ_{mn} is a demand shock realized after the production and entry choices have been made, α_n is the destination-specific price sensitivity, and p_{mn} is the consumer price. Given market size M_n, demand is multinomial logit:

$$\frac{q_{mn}}{M_n} = \frac{\exp(\delta_{mn} + \xi_{mn} - \alpha_n p_{mn})}{1 + \sum_{f'\in\mathcal{F}} \sum_{k\in\Omega_{f'n}} \exp(\delta_{kn} + \xi_{kn} - \alpha_n p_{kn})}. \tag{3}$$

The demand shifter δ_{mn} is flexible enough to capture product and destination heterogeneity. In the empirical specification, it includes product fixed effects, interactions between destination income and product characteristics, an indicator for whether destination n is the headquarters country of the brand producing product m, and the distance from destination n to the brand's headquarters country. The price coefficient varies with income in destination n, so consumers in richer and poorer

---

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

countries may differ in their sensitivity to prices.

Supplying product m to destination n from production origin o requires a constant delivered marginal cost. We write this cost as the product of an ad-valorem final-good tariff factor, a non-tariff transport cost, and a marginal production cost:

$$c_{m,o,n} = t_{on} \tau_{on} c_{mo}. \tag{4}$$

Here t_{on} is one plus the tariff rate applied by destination n to vehicles produced in origin o, τ_{on} is a non-tariff origin-destination trade cost, and c_{mo} is the marginal production cost of product m in origin o. The trade-cost component depends on the distance between the origin and the destination and on whether the vehicle is sold in its production country. The marginal production cost component depends on a brand and an origin fixed effect, product characteristics, the distance between the production origin and the brand headquarters country, whether production takes place in the headquarters country, and auto-parts tariffs between the headquarters and origin countries.

Given the first-stage production vector $D_m = \{D_{mo}\}_{o \in \mathcal{N}}$, firms source each sold product from the active plant with the lowest delivered marginal cost:

$$c_{mn}(D_m) = \min_{o: D_{mo}=1} \{t_{on} \tau_{on} c_{mo}\}. \tag{5}$$

If no plant is active for product m, we interpret $c_{mn}(D_m) = +\infty$, so the product cannot be profitably supplied.[4]

Conditional on $\{\xi_{mn}\}_{m,n}$, the second stage is a complete-information Cournot game. Firms simultaneously choose quantities $\{q_{mn}\}_{m \in \Omega_{fn}}$ in every destination market in which their products are offered. For notational ease, fix a destination n and let Q_{fn} denote firm f's total quantity sold in that destination across all of its offered products, and let Q_n denote aggregate quantity sold in destination n across all firms:

$$Q_{fn} = \sum_{m \in \Omega_{fn}} q_{mn}, \qquad Q_n = \sum_{f \in \mathcal{F}} Q_{fn}. \tag{6}$$

Inverting equation (3) yields the inverse demand system,

$$p_{mn} = \frac{1}{\alpha_n} \left[\delta_{mn} + \xi_{mn} - \log q_{mn} + \log(M_n - Q_n) \right]. \tag{7}$$

Firm f's variable profits in destination n are

$$\pi_{fn} = \sum_{m \in \Omega_{fn}} \left(p_{mn} - c_{mn}(D_m) \right) q_{mn}. \tag{8}$$

---

[4] This minimum-cost sourcing rule is optimal because, conditional on the product being offered, production origin does not enter demand for a given product-market except through delivered marginal costs; second-stage profits are therefore decreasing in the chosen origin's delivered marginal cost, holding demand and rival quantities fixed.

The first-order condition with respect to q_{mn}, for any $m \in \Omega_{fn}$, implies

$$p_{mn} - c_{mn}(D_m) = \frac{1}{\alpha_n}\left(1 + \frac{Q_{fn}}{M_n - Q_n}\right). \tag{9}$$

Thus, all products offered by the same firm in the same destination have the same additive markup. The markup is increasing in the firm's own quantity share and in the penetration of inside goods through the term $Q_{fn}/(M_n - Q_n)$.

To solve the Cournot equilibrium, define utility of product m net of marginal costs as,

$$v_{mn} \equiv \delta_{mn} + \xi_{mn} - \alpha_n c_{mn}(D_m), \tag{10}$$

and define firm f's destination-specific quality-cost index as,

$$K_{fn} \equiv \sum_{m\in\Omega_{fn}} \exp(v_{mn} - 1) = \sum_{m\in\Omega_{fn}} \exp\left(\delta_{mn} + \xi_{mn} - \alpha_n c_{mn}(D_m) - 1\right). \tag{11}$$

Combining equations (7) and (9), summing across the products of firm f, and using the definition of the Lambert W function gives,

$$\frac{Q_{fn}}{M_n - Q_n}\exp\left(\frac{Q_{fn}}{M_n - Q_n}\right) = K_{fn}, \qquad \frac{Q_{fn}}{M_n - Q_n} = W(K_{fn}). \tag{12}$$

Summing the latter expression across firms yields

$$Q_{fn} = \frac{W(K_{fn})}{1 + \sum_{g\in\mathcal{F}} W(K_{gn})} M_n. \tag{13}$$

The within-firm allocation of output across products is proportional to each product's contribution to K_{fn}:

$$q_{mn} = \exp(v_{mn} - 1)\frac{W(K_{fn})/K_{fn}}{1 + \sum_{g\in\mathcal{F}} W(K_{gn})} M_n. \tag{14}$$

Equilibrium prices then follow from the markup equation,

$$p_{mn} = c_{mn}(D_m) + \frac{1 + W(K_{fn})}{\alpha_n}. \tag{15}$$

Finally, substituting the equilibrium markup and quantity into variable profits yields,

$$\pi_{fn} = \frac{M_n}{\alpha_n}\frac{W(K_{fn})\left(1 + W(K_{fn})\right)}{1 + \sum_{g\in\mathcal{F}} W(K_{gn})}. \tag{16}$$

The index K_{fn} is the sufficient statistic through which firm f's own offered products and production locations affect its second-stage variable profits in destination n. Adding a product to destination n raises K_{fn} by adding a positive term to the sum in equation (11). Adding a production location weakly raises K_{fn} because it weakly lowers the delivered marginal cost $c_{mn}(D_m)$ in every destination served with that product. By contrast, rival firms' offerings and production locations affect firm f only through the denominator in equation (16). This representation ties the empirical demand and cost estimates directly to the interdependence properties used by the solution algorithm: own product and plant choices raise own variable profits, while rival choices that increase rivals' K_{gn} lower firm f's variable profits in the same destination market.

3.2 Stage 1: Plant Location and Market Entry

At the first stage of the game, firms simultaneously make plant location and market entry decisions for each of their products. Firms must incur fixed costs to produce and sell products.

We assume that each firm f draws, for each product $m \in \Omega_f$, a location $o^*_m \in \mathcal{N}$ at which it can both produce and sell a product at zero fixed costs.[5] The probability that a location o is the free location follows a multinomial logit,

$$P(o^*_m = o|f) = \frac{\exp\bigl(\gamma_3 \cdot \mathbf{1}\{o \neq h(m)\}\bigr)}{\sum_{k \in \mathcal{N}} \exp\bigl(\gamma_3 \cdot \mathbf{1}\{k \neq h(m)\}\bigr)}, \tag{17}$$

where $h(m)$ denotes the headquarters country of product m's brand. When $\gamma_3 < 0$, all non-headquarters locations share the same free-plant probability, which is lower than that of the headquarters country. We define dummy variable $d_{mo} = 1$ if and only if $o^*_m = o$. The fixed cost of producing product $m \in \Omega_f$ in country $o \in \mathcal{N}$ is, therefore,

$$F^p_{mo} = (1 - d_{mo}) \exp(\gamma_1 + \sigma^p \nu^p_{mo}). \tag{18}$$

The fixed cost of selling product m in country n is,

$$F^e_{mn} = (1 - d_{mn}) \exp(\gamma_2 + \sigma^e \nu^e_{mn}). \tag{19}$$

We assume that $\{\nu^e_{mn}\}_{m,n}$ and $\{\nu^p_{mo}\}_{m,o}$ are normally distributed. As such, fixed costs in non-free locations are log-normally distributed.

Introducing notation, we let $D_{mo} = 1$ if and only if firm f (endowed with product m) chooses to produce product m in origin country o. Analogously, we let $I_{mn} = 1$ if and only if firm f chooses to sell (introduce) product m in destination country n. The total profits for firm f are, therefore,

$$\Pi_f(I_f, D_f, I_{-f}, D_{-f}) = \sum_{n=1}^{N} \Bigl\{ \hat{\pi}_{fn}(I_{fn}, D_f, I_{-fn}, D_{-f}) - \sum_{m \in \Omega_f} I_{mn} F^e_{mn} \Bigr\} - \sum_{o=1}^{N} \sum_{m \in \Omega_f} D_{mo} F^p_{mo}, \tag{20}$$

---

[5] In the data, we only observe products that are produced and sold in at least one country. This assumption guarantees that independently of the realization of fixed cost draws, we match this feature of the data, overcoming a possible sample selection problem.

where I_f and D_f are the $M_f \times N$ matrices of offering and plant location decisions for firm f, respectively, and $\hat{\pi}_{fn} = \mathbb{E}_\xi[\pi_{fn}]$ are the expected profits given any configuration of production and sales decisions for all firms.[6] We denote by I_{-f} and D_{-f} the offering and production location decisions for all firms except for firm f.

Firms make market offering and plant location decisions in a setting of complete information. At Stage 1, firms' information set is,

$$\mathcal{J}_1 \equiv (\{\delta_{mn}\}_{m,n}, \{\alpha_n\}_n, \{(t_{on}, \tau_{on})\}_{o,n}, \{c_{mo}\}_{m,o}, \{F^p_{mo}\}_{m,o}, \{F^e_{mn}\}_{m,n}).$$

Set $\mathcal{J}_1$ includes information on all own and rival payoff-relevant variables except for demand shocks $\{\xi_{mn}\}_{m,n}$, which are realized in Stage 2. In addition to the elements in $\mathcal{J}_1$, we assume that firms in Stage 1 know the distribution of $\{\xi_{mn}\}_{m,n}$.

Additionally, we depart from Nash equilibrium, the solution concept used in previous papers studying complete-information games with multiple interdependent binary decisions (e.g., Jia 2008, Fan and Yang 2025). Under Nash, firms have correct equilibrium beliefs about the strategy profile that rival firms play. Instead, we only impose that firms have common knowledge of rationality i.e., best-response behavior. Our concept is therefore equivalent to point rationalizability in pure strategies, discussed in detail in Bernheim (1984) and Pearce (1984). It is also equivalent, as these papers demonstrate, to Iterated Elimination of Dominated (Pure) Strategies (IEDS). Henceforth, we use the acronym IEDS to refer to the solution concept we impose, which is weaker than the commonly assumed pure strategy Nash equilibria (PSNE).

We depart from Nash equilbria for three reasons. First, neither uniqueness nor existence of PSNE are guaranteed in non-supermodular games with multiple interdependent binary decisions. While Jia (2008) shows how to overcome these issues in settings with two players, the methods developed in that paper do not apply in oligopolistic settings with more than two firms. Our focus on finding strategy profiles that survive IEDS allows us to jointly overcome non-existence and multiplicity. Indeed, the set of IEDS strategies is always non-empty, and any existing PSNE is contained in the set.

Second, we show in Section 4 how to learn about the IEDS set in a computationally feasible manner. A main challenge in dealing with games (or even single-agent problems) with multiple binary decisions is that enumerating all alternatives is impossible. Indeed, the number of candidate equilibria is 2^J, where J is the number of total binary decisions in the game. Finding the full set of Nash equilibria in these settings is therefore an NP-hard problem that is not feasible. Procedures such as the greedy algorithm (i.e., Fan and Yang 2020, Fan and Yang 2025) have been used in previous work to approximate the equilibria, but in general, even computing the best response for a single player making multiple binary decisions is infeasible. We show that focusing instead on IEDS allows us learn about firms' counterfactual behavior without relying on heuristic approaches.

Finally, by imposing weaker behavioral assumptions, IEDS provides more robust estimates and

---

[6] We assume that firms know the distribution of ξ, which is assumed to coincide with the empirical distribution recovered at the estimation of demand.

predictions. We do not impose that firms are best-responding to each other; we only impose that firms iteratively eliminate action profiles that are dominated under common knowledge of rationality. As such, even if we could feasibly solve for all the Nash equilibria, our approach has the advantage of not imposing equilibrium selection rules neither when estimating the fixed-cost parameters nor when computing counterfactual exercises.

4 Solution Method for Multi-Product Plant Location Games

We develop a solution approach for games with multiple binary decisions both within and across players. Our solution procedure bounds action profiles surviving IEDS. The key methodological contribution is to build on the method in Castro-Vincenzi et al. (2023), for single-agent settings, by incorporating both pairwise complementary and substitutable binary decisions in game-theoretic settings. Our algorithm allows for pairs of binary decisions within or across firms to be either (strategic) complements or (strategic) substitutes, and therefore relaxes the supermodularity or submodularity restrictions of pre-existing approaches.[7]

Our solution procedure consists of two steps. We first provide a first-step algorithm that provides an upper bound on the set of action profiles surviving IEDS in our game. Then, we show how to refine our algorithm with a novel branch-and-bound procedure that sharpens our bounds.

4.1 First-Step Algorithm

Our first-step procedure is an iterative algorithm that yields an upper bound on the set of action profiles that survive IEDS. Let π_f denote the payoff function of firm f, which maps an action profile $a \in A = \bigtimes_{f=1}^{F} A_f$ to a real number, where A_f is a set consisting all possible action profiles for firm f. The cardinality of A_f is $2^{|J_f|}$, where J_f is the set of binary decisions belonging to firm f. We also use $\mathcal{A}_f$ to denote the set consisting of all subsets of A_f, and denote by $\mathcal{A} = \bigtimes_{f=1}^{F} \mathcal{A}_f$ the Cartesian product of sets of action profiles across firms. Finally, we denote by $\mathcal{G} = (\mathcal{F}, A, \{\pi_f\})$ the game of interest, which consists of a tuple of players, possible (pure) strategy profiles, and payoff functions. To formally explain our procedure, we first introduce three key definitions.

First, we define a mapping that eliminates dominated strategies, which follows Bernheim (1984) and Pearce (1984).

Definition 1 (Best-Response and IEDS) *Define a game* $\mathcal{G} = (\mathcal{F}, A, \{\pi_f\})$. *Let* $BR_f(C) = \arg\max_{s \in A^f} \pi_f(s, C_{-f})$, *which is unique almost-surely, be firm* f*'s best response to* C_{-f}. *Define the mapping* $\lambda_{\mathcal{G}} : \mathcal{A} \to \mathcal{A}$, *by,* $\lambda_{\mathcal{G}}(B) = \bigtimes_{f \in \mathcal{F}} \bigcup_{s \in B} BR_f(s)$. *Then, starting from* A, *the set of*

---

[7] Jia (2008) builds on the literature on supermodular games in a complete-information setting. Sabal (2025) provides a solution algorithm for games under the assumption of submodularity, though in a private-information setting.

action profiles surviving IEDS is,

$$IEDS(\mathcal{G}) = \bigcap_{k=0}^{\infty} \lambda_{\mathcal{G}}^{k}(A),$$

where $\lambda_{\mathcal{G}}^{0} := \lambda_{\mathcal{G}}$, $\lambda_{\mathcal{G}}^{1} = \lambda_{\mathcal{G}} \circ \lambda_{\mathcal{G}}$, *...*, $\lambda_{\mathcal{G}}^{k} := \lambda_{\mathcal{G}}^{k-1} \circ \lambda_{\mathcal{G}}$.

Definition 1 formalizes the notion of iterated elimination of (pure) dominated strategies. Starting from the largest set of possible action profiles, $\lambda_{\mathcal{G}}$ iteratively eliminates all actions for each player that are never a best-response to some rival action profile.

We now define the marginal value of a binary action.

Definition 2 (Marginal Value) *For firm f, the marginal value of discrete choice $i \in J_f$ under C is given by,*

$$\Delta_i^f(C) := \pi_f(C^{i\to 1}, C^{-f}) - \pi_f(C^{i\to 0}, C^{-f}) \tag{21}$$

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

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

Definition 3 (Cross-Partial) *For any two choices $i \in J^f$ and $j \in J$ we define the cross-partial of i and j under C as*

$$\Delta_{ij}^f(C) := \Delta_j^f(C^{i\to 1}, C^{-f}) - \Delta_j^f(C^{i\to 0}, C^{-f}). \tag{22}$$

It follows that $\Delta_{ij}^f(C)$ is the difference in marginal value of coordinate i when j is set to 1 relative to when j is set to 0.

A sufficient assumption for our first-step algorithm is that, for all pairs of discrete choices i and j, the *sign* of $\Delta_{ij}^f(C)$ is known and independent of C. We state this assumption formally below.

Assumption 1 *For all $C \in A$, for all f, for any $i \in J_f$ and $j \in J$, the sign of $\Delta_{ij}^f(C)$ is known and independent of C.*

To develop our algorithm, we build on Castro-Vincenzi et al. (2023) and apply Tarski's fixed point theorem (Tarski 1955) on a lattice of sublattices of the J-dimensional binary vectors, $\mathcal{B}^{|J|} = A$. Denote by $S(\mathcal{B}^{|J|}) \subseteq \mathcal{A}$ the collection containing the empty set and all interval-sublattices of $\mathcal{B}^{|J|}$. That is, any sublattice $B \subseteq \mathcal{B}^{|J|}$ such that for every $x, y \in B$, it is the case that $\{z \in \mathcal{B}^J : x \wedge y \leq z \leq x \vee y\} \subseteq B$, where $x \wedge y$ is the infimum (meet) of x and y in B and $x \vee y$ is the supremum (join) of x and y in B.[8] If $J = 2$, then $\mathcal{B}^{|J|} = \{(0,0),(1,0),(0,1),(1,1)\}$ and

$$\begin{aligned} S(\mathcal{B}^{|J|}) = \big\{&\emptyset, \{(0,0)\}, \{(1,0)\}, \{(0,1)\}, \{(1,1)\}, \{(0,0),(1,0)\}, \{(0,0),(0,1)\}, \\ &\{(1,0),(1,1)\}, \{(0,1),(1,1)\}, \ldots, \{(0,0),(1,0),(0,1),(1,1)\}\big\}. \end{aligned}$$

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 not include all elements between $(0,0,0)$ and $(1,1,1)$ according to the partial order $\leq$.

In practice, Assumption 1 requires that, given model parameters, any two choices are known to be either complements or substitutes, independently of third own and rival choices. We refer to two binary choices i and j as being *complements* when $\Delta^f_{ij} > 0$ and as *substitutes* when $\Delta^f_{ij} < 0$. Whenever $i \in J^f$ and $j \in J^g$ for $f \neq g$, we say that i and j are either strategic complements or strategic substitutes.

We construct a set-increasing mapping that flexibly allows for both complements and substitutes across discrete choices in the game. Let $f(i)$ denote the firm that faces discrete choice i, i.e., $f(i) = g$ if and only if $i \in J_g$. Define, for any $\boldsymbol{C} \in S(\mathcal{B}^{|J|})$,

$$\overline{\Omega}(\boldsymbol{C}) = \{i \in J : \Delta_i^{f(i)}(\sup_i(\boldsymbol{C})) < 0\} \tag{23}$$

$$\underline{\Omega}(\boldsymbol{C}) = \{i \in J : \Delta_i^{f(i)}(\inf_i(\boldsymbol{C})) \geq 0\} \tag{24}$$

where, for each $i \in J$ and $j \in J$,

$$[\sup_i(\boldsymbol{C})]_j := \mathbb{1}[\Delta_{ij}^{f(i)} \geq 0][\sup(\boldsymbol{C})]_j + \mathbb{1}[\Delta_{ij}^{f(i)} < 0][\inf(\boldsymbol{C})]_j \tag{25}$$

$$[\inf_i(\boldsymbol{C})]_j := \mathbb{1}[\Delta_{ij}^{f(i)} \geq 0][\inf(\boldsymbol{C})]_j + \mathbb{1}[\Delta_{ij}^{f(i)} < 0][\sup(\boldsymbol{C})]_j. \tag{26}$$

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

$$F(\boldsymbol{C}) = \{C \in \boldsymbol{C} \colon C_i = 0 \text{ and } C_{i'} = 1 \, \forall i \in \overline{\Omega}(\boldsymbol{C}) \text{ and } \forall i' \in \underline{\Omega}(\boldsymbol{C})\}. \tag{27}$$

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 $i \in \overline{\Omega}(\boldsymbol{C})$ and $F(\boldsymbol{C})$ returns only vectors C with $C_i = 0$.

We prove that the mapping F is set-increasing.

---

[8] 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.

Theorem 1 (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 B. ■

Theorem 1 results from the monotonicity of the infimum and supremum together with Assumption 1. 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* $k = 1, 2, \ldots$ *until convergence.*

Proposition 1 states that Algorithm 1 converges monotonically to a set $\boldsymbol{C}^{fixed}$. To characterize $\boldsymbol{C}^{fixed}$, consider a "fake" game, identical to our game of interest, but in which each binary decision is undertaken by a separate player. That is, a game in which the set of players is J, and any two players with $f(i) = f(j) = g$ have the same payoff function, π_g. We denote this game as $\mathcal{G}_F$.

Proposition 1 *Algorithm 1 has the following properties:*

1. $\boldsymbol{C}_k = F(\boldsymbol{C}_{k-1}) \subseteq \boldsymbol{C}_{k-1}$ *for all* $k \in \mathbb{N}$.
2. *It converges to a fixed point* $\boldsymbol{C}^{fixed}$ *in a finite number of iterations* K *i.e.,* $F(\boldsymbol{C}_T) = \boldsymbol{C}^{fixed}$ *for all* $T \geq K$.
3. $F = \lambda_{\mathcal{G}_F}$.
4. $IEDS(\mathcal{G}_F) = \boldsymbol{C}^{fixed}$.
5. $IEDS(\mathcal{G}) \subseteq \boldsymbol{C}^{fixed}$.
6. *If a pure-strategy Nash equilibrium* C^* *exists, then it is an element of* $\boldsymbol{C}^{fixed}$.

Proof. See Appendix B. ■

Proposition 1 provides an economic interpretation of the action profiles contained in $\boldsymbol{C}^{fixed}$. If $|J_f| = 1$ for each firm f, then Algorithm 1 yields the object of interest: the set of action profiles surviving IEDS. Intuitively, under Assumption 1, Algorithm 1 operationalizes in a computationally feasible manner the procedure to iteratively rule out dominated strategies. Note that in the absence of Algorithm 1, ruling out dominated strategies is not generally computationally feasible. For each player f, one would have to check whether each of 1 or 0 is a best-response for some of the 2^{J-1} possible rival action profiles. Using knowledge of the sign of cross-partials, one does not have to loop over 2^{J-1} action profiles. Instead, it is necessary and sufficient to iteratively evaluate the returns to playing 1 (or 0) at known binary-action-specific bounds.

When players internalize more than a single binary action, our coordinate-specific bounds remain valid, but they may be slack, as stated in item 4 of Proposition 1. The reason is that our algorithm only eliminates dominated binary actions, but cannot eliminate dominated multi-coordinate actions by evaluating the returns of moving multiple binary actions at a time.

Finally, item 5 of Proposition 1 states that any existing pure strategy Nash equilibrium is contained in $\boldsymbol{C}^{fixed}$. Thus, even if the object of interest were Nash equilibria, our procedure provides a computationally feasible procedure to find a set that contains all Nash equilibria, and therefore to deal with equilibrium multiplicity. However, even under the possible non-existence of PSNE, $\boldsymbol{C}^{fixed}$ is an economically interpretable solution concept in our setting.

4.1.1 Interdependent Decisions in the Model

The plant and market entry game exhibits interdependent decisions across locations, products, and firms. We show that the sign of the interdependence across any pair of binary choices in the game is independent of third choices, and satisfies Assumption 1. That is, any pair of binary choices are always either (strategic) complements or (strategic) substitutes regardless of the value of (I, D). Below, we list the sign of the cross-partials that arises in our model.
Consider a production decision D_{mo} for firm f, product m, and origin o.

1. It is *complementary* to
    (a) I_{mn} for any market n, since adding an additional plant for product m weakly lowers its marginal cost in all markets.
2. It is *substitutable* to
    (a) $I_{m'n}$ for any $m' \neq m$, and $m' \in \Omega_f$, and any n, due to cannibalization: producing product m in country o reduces the cost of supplying it in market n, lowering the returns from selling product m' in that same market.
    (b) I_{kn} for any $k \in \Omega_{f'}$, with $f' \neq f$, and any n, since a weakly less costly model m lowers the benefit of selling rival products in the same market.
    (c) $D_{mo'}$ for any $o' \neq o$, because an additional plant in country o is less valuable when product m is already produced in many countries, as per minimum-cost sourcing.
    (d) $D_{m'n}$ for any $m' \neq m$, $m' \in \Omega_f$, and any n, since the decline in marginal cost for product m weakens the incremental sales gain from lowering the cost of product m'.
    (e) D_{kl} for any $k \in \Omega_{f'}$, with $f' \neq f$, and any l, because when other rival products become cheaper, the returns from lowering marginal costs decrease, due to cannibalization.

Then, consider a sales decision I_{mn} of firm f, product m, in market n.

1. It is *complementary* to
    (a) D_{mo} for any origin o, since an additional plant is more valuable when the product is sold in more markets.
2. It is *substitutable* to
    (a) $D_{m'o}$ for any $m' \neq m$, $m' \in \Omega_f$. When other products within the same firm become cheaper, the additional gains from selling product m decline, due to cannibalization.
    (b) D_{kl} for any $k \in \Omega_{f'}$, $f' \neq f$, any k, and any l, since less costly rival firms reduce the returns from selling product m.
    (c) $I_{m'n}$ for any $m' \neq m$, $m' \in \Omega_f$, due to cannibalization forces in market n.
    (d) I_{kn} for any $k \in \Omega_{f'}$, with $f' \neq f$, any model k, due to business stealing externalities across firms.
3. It is *independent* of
    (a) $I_{kn'}$ any model k, and any market $n' \neq n$, because demand is independent across markets.

We provide a formal proof of these properties in Appendix A.

4.2 Simple Example

We provide intuition for why our solution algorithm works. Consider a case with $F = 2$ firms, each with $J_1 = J_2 = 3$ binary decisions. Let $J_1 = \{1, 2, 3\}$, and $J_2 = \{4, 5, 6\}$. Suppose that for firm 1, we have,

$$\Delta_{12}(\pi_1(C)) \geq 0; \quad \Delta_{13}(\pi_1(C)) < 0; \quad \Delta_{23}(\pi_1(C)) \geq 0.$$

That is, coordinates 1 and 2 are complements; coordinates 1 and 3 are substitutes; and coordinates 2 and 3 are complements. Moreover, suppose that,

$$\Delta_{ii'}(\pi_1(C)) < 0 \qquad \text{for any } i \in \{1, 2, 3\}, \text{ and any } i' \in \{4, 5, 6\}.$$

Thus, binary decisions belonging to different firms are substitutes. Note that under Assumption 1, we can construct bounds on the marginal value of each of the binary actions in the game. For instance, for binary actions 1 and 2, we can construct lower bounds,

$$\begin{aligned} \underline{\Delta}_1 &= \Delta_1(\pi_1((\cdot, 0, 1), (1, 1, 1)), \\ \underline{\Delta}_2 &= \Delta_2(\pi_1((0, \cdot, 0), (1, 1, 1)). \end{aligned}$$

Note that we construct these lower bounds by evaluating the marginal value at a binary-action-specific vector C that sets all complements of i to 0 and all substitutes of i to 1. Analogously, we can construct upper bounds by evaluating the marginal values at coordinate-specific vectors that set complements to 1 and substitutes to 0.

If, for instance, $\underline{\Delta}_1 > 0$, our first-step algorithm says that *any* action profile that survives IEDS will have $C_1 = 1$. In other words, any action profile with $C_1 = 0$ is dominated, or not rationalizable.

Crucially, we can fix $C_1 = 1$ when evaluating the marginal-value bounds for other binary actions. Instead of evaluating $\underline{\Delta}_2 = \Delta_2(\pi_1((0, \cdot, 0), (1, 1, 1))$, the lower bound is now $\Delta_2(\pi_1((1, \cdot, 0), (1, 1, 1))$, which is larger than $\underline{\Delta}_1$, because 1 and 2 are complements. Therefore, updating the first coordinate to 1 favors updating the second coordinate to 1 in the next iteration. It does not, however, affect our ability to update the second coordinate to 0, since $\overline{\Delta}_2 = \Delta_2(\pi_1((1, \cdot, 1), (0, 0, 0))$ remains unchanged.

Applying the mapping F iteratively effectively fixes binary actions for which only 0 or 1 is rationalizable. The algorithm effectively returns bounds,

$$\overline{C}_i = \mathbb{1}\{\Delta_i^{f(i)}(\sup_i(\boldsymbol{C}^{fixed}) \geq 0)\} \geq C_i \geq \mathbb{1}\{\Delta_i^{f(i)}(\inf_i(\boldsymbol{C}^{fixed}) \geq 0)\} = \underline{C}_i,$$

for each $i \in J$.

4.3 Branch-and-Bound Algorithm for Games

While Proposition 1 guarantees that the first-step algorithm converges to a set containing $IEDS(\mathcal{G})$, this set is not generally sharp. To sharpen our bounds, we provide a novel branch-and-bound algorithm for games with many interdependent binary actions. Our iteratively applies a procedure across firms that builds on the branching algorithm for single-agent problems developed in Castro-Vincenzi et al. (2023). In single-agent problems, the branching algorithm in that paper iteratively applies the first-step algorithm by creating "branches" that fix to 0 and 1 decisions that are unsolved in the first-step. If in any of the branches the first-step algorithm does not converge, new branches are created and the first-step algorithm is again applied in each of those branches. In a single-agent setting, this algorithm terminates when convergence has been reached in all branches to either empty sets (i.e., a branch was not consistent with optimizing behavior), or to singleton sublattices that contain candidate global solutions (i.e., vectors for which there is no profitable single-coordinate deviation). Finally, the algorithm compares the payoff obtained at each of the singleton sublattices and returns the maximizer.

We develop a branching algorithm for games that enables us to eliminate action profiles that are included in $IEDS(\mathcal{G}_F)$, but not in $IEDS(\mathcal{G})$, which is the set of interest. Our procedure applies a branching procedure iteratively across firms. More precisely, we branch-and-bound by applying the following first-step mapping,

$$F_f(\boldsymbol{C}) = \{C \in \boldsymbol{C} : C_i = 0 \text{ and } C_{i'} = 1 \ \forall i \in \overline{\Omega}_f(\boldsymbol{C}) \text{ and } \forall i' \in \underline{\Omega}_f(\boldsymbol{C})\}, \tag{28}$$

where

$$\overline{\Omega}_f(\boldsymbol{C}) = \{i \in J_f : \Delta_i^{f(i)}(\sup_i(\boldsymbol{C})) < 0\}, \tag{29}$$

$$\underline{\Omega}_f(\boldsymbol{C}) = \{i \in J_f : \Delta_i^{f(i)}(\inf_i(\boldsymbol{C})) \geq 0\}. \tag{30}$$

The mapping F_f is identical to F but only updates coordinates corresponding to firm f. The mapping F_f still takes into account that the payoff for any of firm f's binary decisions are dependent on rival firms' actions, since $\boldsymbol{C}$ is a sublattice of action profiles in the entire game.

In Appendix C, we define the branching algorithm applied to a particular firm in a strategic setting. The branching procedure starts from $\boldsymbol{C}^{fixed}$, and fixes some unsolved decision for firm f to 0 and 1. In each of these two branches, it applies F_f iteratively until convergence. In any of these branches, there are two possibilities. The first possibility is that applying F_f until convergence yields a candidate solution for firm f, i.e., no decisions remain unsolved. In this game-theoretic setting, this amounts formally to obtaining a sublattice $\boldsymbol{C}_f$ in which all elements of the sublattice contain a single action for firm f, combined with all rival action profiles contained in $\boldsymbol{C}^{fixed}$. The second possibility in any of these branches is that even after applying F_f until convergence, some coordinates remain unsolved. If this is the case, the procedure generates two further branches within this branch and, in each of these branches, applies F_f iteratively until convergence. Once we have fixed more than one coordinate by branching, it is possible that F_f converges to an empty set. This happens whenever two or more coordinates fixed by branching are incompatible with firm f's best-response behavior. We show in Appendix C that applying the branching mapping iteratively yields a non-empty collection of sublattices $\bigcup_k \boldsymbol{C}_f^k$, where in each $\boldsymbol{C}_f^k$ the set of action profiles for firm f is a singleton.

In Proposition 2 of Appendix C, we prove that the collection $\bigcup_k \boldsymbol{C}_f^k$ has economic meaning. We show that any action profile for firm f surviving IEDS corresponds to some element $\boldsymbol{C}_f^k$ of the collection returned by branching. Crucially, we show that it is not possible for branching to miss any IEDS action profile for firm f. Still, this collection could contain actions for firm f that do not survive IEDS in the game. This is clearly the case because, so far, we have only solved for candidate IEDS profiles for firm f, but have not used this information (and common knowledge of rationality) to update other firms' decisions.

Theoretically, a limitation of this approach is that the collection $\bigcup_k \boldsymbol{C}_f^k$ may be very large and impossible to store in memory. In practice, we find that the first-step procedure combined with firm-by-firm branching rules out a large number of candidates, and it is feasible to store all candidate IEDS actions for firm f.

When applying branching to the next firm (after applying it to firm f), there are two approaches for using the information contained in $\bigcup_k \boldsymbol{C}_f^k$. One possibility is to loop over each of the candidate actions for firm f and apply branching to the next firm starting at each of $\boldsymbol{C}_f^k$. The issue with this approach is that it is challenging to scale it to cases with many firms, and one has to keep track of a theoretically combinatorial set of candidate IEDS actions. Instead, we follow a second approach, which is more computationally tractable, which we summarize in Algorithm 2.

Algorithm 2 *The combined solution algorithm proceeds as follows:*

1. *Start at* $\boldsymbol{C}_0 = \mathcal{B}^{|J|}$ *and apply the first-step algorithm* F *until convergence to* $\boldsymbol{C}^{fixed}$.
2. *If there remain firms with binary choices that are neither solved to* 0 *nor* 1, *impelement branching firm-by-firm. After branching for any such firm* f, *update the binary-action-specific bounds using the information in* $\bigcup_k \boldsymbol{C}_f^k$. *Update* $\boldsymbol{C}^{fixed}$ *to* $\boldsymbol{C}_2^{fixed}$. *Then, apply branching to the next firm starting from* $\{\boldsymbol{C}_2^{fixed}\}$.
3. *Apply step 2 iteratively across firms until convergence.*

Our approach leverages the fact that action profiles for firm f in $\bigcup_k \boldsymbol{C}_f^k$ may all have the same value for some decision i of firm f. That is, branching for firm f may have improved the bounds relative to the first-step procedure. Thus, when applying branching on the next firm, say firm g, we may start at a modified sublattice $\boldsymbol{C}_2^{fixed} \subseteq \boldsymbol{C}_{fixed}$ that contains only the actions for firm f contained in $\boldsymbol{C}^{fixed}$ compatible with the coordinates that were fixed by branching. We apply this procedure iteratively across firms until convergence, thus refining the first-step procedure in a computationally feasible manner. Algorithm 2 combines the first and second-step algorithms. In Appendix C, we show that the branching algorithm is also a set-increasing mapping. Since Algorithm 2 is a composition of set-increasing mappings, which is also set-increasing, the algorithm yields bounds on the set of interest, $IEDS(\mathcal{G})$.

4.4 Extensions

We discuss two extensions of our method in our setting. In a single-agent setting, Castro-Vincenzi et al. (2023) discusses related extensions at length. In this paper, we build on these extensions for strategic settings. First, we can solve for production location decisions by defining mapping F with a product-specific value function, defined by,

$$V_m(D_m; I_{-m}, D_{-m}) := \max_{I_m} \sum_{n=1}^{N} I_{mn} \Delta\pi_{fn}^{m}(D_m; I_{-m}, D_{-m}) - \sum_{n=1}^{N} D_{mn} F_{mn}^{p}. \tag{31}$$

We use this extension because it leverages the property of our model which says that sales decisions across countries are independent conditional on production location decisions. As such, when evaluating the marginal value of a production location decision, we obtain more informative bounds by solving for optimal sales decisions (which is quick due to independence across countries), rather than holding them fixed at bounds. In Appendix A.3, we show that V_m exhibits decreasing differences in D_m and in I_{-m} and D_{-m}, even in the presence of strategic interactions.

Second, whenever firms have payoff functions $\pi_f = \sum_{k=1}^{K} \pi_{f,k}$ that are additively separable in K components, Assumption 1 can be relaxed, even in the presence of strategic interactions. Instead of requiring constant-signed cross-partials for each firm's profit function (as in Assumption 1), the more general requirement is that this property holds for each additive component of each firm's payoff function. This holds because one can always modify F by redefining,

$$\overline{\Omega}(\boldsymbol{C}) = \{i \in J : \sum_k \Delta_i^{f(i),k}(\sup_{i,k}(\boldsymbol{C})) < 0\}$$
$$\underline{\Omega}(\boldsymbol{C}) = \{i \in J : \sum_k \Delta_i^{f(i),k}(\inf_{i,k}(\boldsymbol{C})) \geq 0\}.$$

That is, the modified F mapping fixes coordinates by computing bounds on the marginal value as the sum over bounds on coordinate-component-specific marginal values. While a caveat is that the "fake" game interpreation is now less straightforward, it is still the case that the fixed point of this mapping yields a candidate set of action profiles that is a superset of those surviving iterated elimination of dominated strategies.

4.5 Performance

Table 3 summarizes the performance of the solution algorithm across 10 draws of fixed-cost shocks, evaluated at the midpoint of the confidence set. The first-step squeezing procedure alone resolves the majority of the binary choices. On average, squeezing solves 36,741 of the 36,800 coordinates, leaving only 19 production and 40 sales decisions unresolved. This corresponds to 99.8 percent of all coordinates being pinned down after the first step. At the firm level, squeezing fully solves 35 of the 48 firms on average.

Branch-and-bound delivers an additional refinement. After the second step, the average number of solved coordinates rises to 36,757, while the number of unresolved production and sales choices falls to 15 and 28, respectively. Thus, the combined algorithm resolves 99.9 percent of all coordinates and fully solves 36 firms on average. In computational terms, the procedure is fast: squeezing on the full game takes 2.6 seconds per draw on average, while branch-and-bound adds only 0.4 seconds.

5 Estimation

We estimate the model in two steps. First, we recover the demand and cost primitives that determine second-stage variable profits under Cournot competition. This step delivers destination-specific price sensitivities, demand shifters, and the matrix of potential delivered marginal costs for every product-origin-destination triplet. Second, taking these primitives as given, we estimate the fixed-cost parameters that govern market entry, plant location, and the distribution of free locations. Because the first-stage game is incomplete and firms make many interdependent binary decisions, we do not target a unique equilibrium. Instead, we combine the bounds produced by our solution algorithm with monotone moment inequalities, so each candidate fixed-cost parameter vector is evaluated by whether it can rationalize the observed pattern of product offerings and production locations.

Table 3: Performance of the Solution Algorithm

	After Squeeze				After Branch & Bound				Time (s)	
Draw	Choices Solved	Unsolved Prod	Unsolved Sale	Firms Solved	Choices Solved	Unsolved Prod	Unsolved Sale	Firms Solved	Sq	B&B
1	36,727	14	59	28	36,751	11	38	29	3.1	0.9
2	36,750	17	33	39	36,759	15	26	39	2.0	0.1
3	36,787	6	7	46	36,787	6	7	46	2.0	0.0
4	36,759	16	25	36	36,768	13	19	36	1.9	0.1
5	36,712	29	59	31	36,728	25	47	31	2.1	1.2
6	36,694	28	78	29	36,736	17	47	32	2.0	0.5
7	36,746	24	30	38	36,758	18	24	38	2.0	0.1
8	36,721	27	52	31	36,743	20	37	31	2.0	0.3
9	36,741	15	44	32	36,757	14	29	32	1.9	0.3
10	36,774	10	16	44	36,784	8	8	44	1.9	0.0
Mean	36,741	19	40	35	36,757	15	28	36	2.1	0.4
Total	36,800	11,040	25,760	48	36,800	11,040	25,760	48		
%	99.8	0.2	0.2	72.9	99.9	0.1	0.1	75.0		

Notes: This table reports the performance of the solution algorithm across 10 draws of fixed-cost shocks. "After Squeeze" and "After B&B" report counts at the end of Step 1 (squeezing) and Step 2 (branch-and-bound), respectively. "Choices Solved" counts the coordinates resolved to 0 or 1; "Unsolved Prod" and "Unsolved Sale" decompose the remaining unsolved coordinates by choice type; "Firms Solved" counts firms with all coordinates resolved. The "Total" row gives the total counts of coordinates and firms, and the "%" row gives the mean as a fraction of total. The last two columns report computation times in seconds by step. Performance is evaluated at the midpoint of the confidence set reported in Table 6. Timings are from a single core of an Apple M1 Pro processor (8 cores, 32 GB RAM).

5.1 Demand and Marginal Costs

We estimate demand and marginal costs using the 12 destination markets for which we observe prices and product characteristics. The empirical strategy is tightly connected to the Cournot model in Section 3.1. First, we estimate the demand shifters and the destination-specific price sensitivities that enter utility. Second, using the Cournot markup equation, we invert observed prices and quantities to recover implied delivered marginal costs. Third, we decompose these implied costs into final-good tariffs, non-tariff trade costs, and model-origin production costs.

For demand, let $s_{mn} = q_{mn}/M_n$ be the market share of product m in destination n, and let $s_{0n} = 1 - \sum_{k\in\Omega_n} s_{kn}$ denote the outside-option share. Equation (3) implies

$$\log\left(\frac{s_{mn}}{s_{0n}}\right) = \phi_m + \phi_n + \beta_p\tilde{p}_{mn} + \beta_{py}\tilde{p}_{mn}\log y_n + X'_{mn}\beta + u_{mn}, \tag{32}$$

where p_{mn} is the price of product m in market n, y_n is destination log GDP per capita, ϕ_m is a product fixed effect, ϕ_n is a destination-year fixed effect, and X_{mn} includes interactions of log income with log horsepower, log weight, and log size, an indicator for the brand's home market, and log distance from the destination to the brand headquarters country. We estimate equation (32) by 2SLS, instrumenting price and price interacted with log GDP per capita using final-good tariffs, origin-destination distance, and their log-income interactions. The implied destination-specific price sensitivity is,

$$\hat{\alpha}_n = -\left(\hat{\beta}_p + \hat{\beta}_{py} \log y_n\right). \tag{33}$$

Given $\hat{\alpha}_n$, the Cournot markup equation (9) can be written in shares as

$$\widehat{\mu}_{fn} = \frac{1}{\hat{\alpha}_n}\left(1 + \frac{s_{fn}}{s_{0n}}\right), \qquad s_{fn} \equiv \sum_{k \in \Omega_{fn}} s_{kn}. \tag{34}$$

The additive markup is common across products sold by firm f in destination n. We therefore recover an implied delivered marginal cost for each model-origin-destination observation as,

$$\widehat{c}_{mn} = p_{mn} - \widehat{\mu}_{fn}. \tag{35}$$

Next, we project the implied delivered marginal costs onto fixed effects and observables. Let o denote the observed production origin serving destination n with product m. Because $\widehat{c}_{mn}$ includes final-good tariffs, we net out the final-good tariff factor and estimate,

$$\log \widehat{c}_{mn} - \log t_{on} = \lambda_{mo} + \lambda_n + \kappa_1 \log dist_{on} + \kappa_2 \mathbb{1}\{o = n\} + \varepsilon_{mon}. \tag{36}$$

The model-origin fixed effect λ_{mo} captures the marginal production cost of product m in origin o, and the destination fixed effect absorbs destination-specific price wedges and measurement differences. The coefficient on log distance and the same-country indicator recover the non-tariff trade-cost component τ_{on}.

In the final step, we project the estimated model-origin fixed effects on origin and brand indicators, distance between headquarters and origin, parts tariffs, and product characteristics:

$$\begin{aligned}\widehat{\lambda}_{mo} = \kappa_o + \kappa_b + \kappa_3 \log dist_{oh(b)} + \kappa_4 \mathbb{1}\{o = h(b)\} + \kappa_5 \log(1 + \text{parts tariff}_{h(b)o}) \\ + \kappa_6 \log HP_m + \kappa_7 \log Weight_m + \kappa_8 \log Size_m + \eta_{mo}.\end{aligned} \tag{37}$$

We weight this regression by the number of destination markets served by each product-origin pair. We use the fitted values from equations (36) and (37) to build the full matrix of marginal costs for every potential product, production origin, and destination that enters the structural model.

Table 4 summarizes the baseline demand and marginal-cost estimates. The price coefficient is allowed to vary with destination income: the coefficient on price is positive, while the interaction between price and log GDP per capita is negative, so the implied price sensitivity is lower for low-income countries. The demand estimates also imply that, conditional on product and destination fixed effects, heavier, more powerful, and larger SUVs are valued more in richer countries. The home-market coefficient is large and positive, consistent with strong preference for brands in their headquarters country.

The marginal-cost estimates are based on the implied marginal costs from the Cournot inversion.

Table 4: Demand and Marginal Cost Parameter Estimates

	Parameter	**Standard Error**
Demand ($N = 2,599$)		
Price (rescaled)	3.313^{a}	0.985
Price × (log) GDPpc	-0.465^{a}	0.118
(log) Horsepower × (log) GDPpc	0.125	0.078
(log) Weight × (log) GDPpc	0.753^{a}	0.201
(log) Size × (log) GDPpc	0.057^{a}	0.019
Home Market	1.207^{a}	0.169
(log) Distance to Brand HQ	0.041	0.062
Marginal Costs (log) ($N = 2,569$)		
Trade Costs		
(log) Distance to Destination (κ_1)	0.037^{a}	0.007
$\mathbb{1}\{o = n\}$ (κ_2)	-0.065^{a}	0.022
Production Costs		
(log) Distance to Brand HQ (κ_3)	0.038^{a}	0.013
$\mathbb{1}\{o = h(b)\}$ (κ_4)	0.059	0.039
(log) Parts Tariff (κ_5)	3.554^{b}	1.565
(log) Horsepower (κ_6)	0.581^{a}	0.049
(log) Weight (κ_7)	1.638^{a}	0.076
(log) Size (κ_8)	0.078^{a}	0.026

Notes: The demand equation is estimated by IV-2SLS, instrumenting price and price interacted with log GDP per capita using tariffs, bilateral distance, and their log-income interactions. Model and destination-year fixed effects are included. The trade-cost regression is estimated by OLS on positive implied marginal costs and includes destination-country and model-origin fixed effects. The production-cost projection is estimated by OLS using the estimated model-origin fixed effects, ISO3 production-origin fixed effects, brand fixed effects, and number-of-served-market weights. Robust standard errors are reported for all regressions. $^{\mathrm{a}}$ denotes 1% significance, $^{\mathrm{b}}$ denotes 5% significance, and $^{\mathrm{c}}$ denotes 10% significance.

Non-tariff trade costs increase with origin-destination distance, and local production lowers delivered costs conditional on distance and tariffs. Our production-cost projection, shows that marginal costs are increasing in distance from the brand headquarters country, in auto-parts tariffs, and in vehicle horsepower, weight, and footprint. These estimates are used to predict production costs for model-origin pairs outside the observed price sample and to construct potential (counterfactual) marginal costs used in the subsequent analyses.

Table 5: Distribution of Own-Price Elasticity and Price-Cost Ratio, Positive Marginal Costs

	Own-price elasticity						Price-cost ratio					
	Mean	Median	P10	P25	P75	P90	Mean	Median	P10	P25	P75	P90
Panel A: By Destination Market												
United States	-6.65	-6.32	-10.15	-7.29	-4.99	-4.53	1.21	1.19	1.12	1.16	1.26	1.30
Germany	-7.10	-6.43	-11.08	-7.84	-5.14	-4.27	1.21	1.19	1.10	1.16	1.26	1.33
Australia	-5.20	-4.10	-8.37	-5.87	-3.70	-3.10	1.31	1.33	1.14	1.21	1.40	1.48
United Kingdom	-6.54	-5.68	-9.93	-7.33	-4.96	-4.10	1.23	1.23	1.12	1.17	1.29	1.36
France	-5.70	-5.56	-7.77	-6.33	-4.11	-3.84	1.27	1.23	1.16	1.20	1.35	1.38
Italy	-5.36	-4.87	-8.64	-5.63	-3.97	-3.62	1.29	1.28	1.13	1.22	1.36	1.41
Japan	-4.37	-3.92	-6.80	-5.19	-3.42	-2.01	1.45	1.35	1.17	1.25	1.42	2.04
Spain	-5.27	-4.97	-7.36	-5.71	-3.81	-3.51	1.29	1.26	1.16	1.22	1.38	1.42
Mexico	-2.40	-2.00	-4.33	-3.75	-1.31	-1.29	2.71	2.00	1.30	1.37	4.22	4.44
Mainland China	-3.64	-2.69	-7.21	-4.33	-1.97	-1.49	2.42	1.61	1.17	1.32	2.05	3.16
Brazil	-2.32	-1.84	-3.36	-2.70	-1.76	-1.70	2.05	2.24	1.43	1.60	2.38	2.47
India	-1.69	-1.43	-2.93	-1.54	-1.13	-1.09	6.77	3.52	1.52	3.05	11.46	12.77
Panel B: Top Parent Companies												
Volkswagen	-6.39	-5.30	-10.51	-7.56	-4.47	-3.51	1.27	1.24	1.11	1.16	1.32	1.43
Renault-Nissan-Mitsubishi	-4.53	-4.11	-6.98	-5.07	-3.31	-2.79	1.47	1.35	1.17	1.26	1.46	1.62
Stellantis	-5.28	-5.01	-7.27	-6.18	-4.17	-3.33	1.31	1.26	1.16	1.21	1.35	1.44
Toyota	-6.13	-5.68	-8.85	-7.26	-4.40	-3.66	1.26	1.22	1.13	1.17	1.32	1.39
Hyundai	-4.05	-4.32	-5.85	-5.10	-2.88	-1.72	2.20	1.31	1.21	1.26	1.54	2.41
General Motors	-6.87	-6.72	-10.68	-8.43	-4.74	-4.21	1.23	1.19	1.11	1.14	1.28	1.34
Ford	-6.25	-6.09	-8.79	-7.14	-4.95	-3.99	1.25	1.20	1.13	1.17	1.26	1.34
Honda	-4.83	-4.78	-7.26	-5.29	-2.78	-2.69	1.40	1.27	1.16	1.24	1.58	1.61
Great Wall	-2.32	-2.35	-2.94	-2.42	-2.09	-1.63	2.03	1.80	1.55	1.76	2.02	2.84
Geely	-4.32	-2.46	-10.22	-7.33	-1.97	-1.08	5.86	1.72	1.11	1.17	2.10	26.04

Table 5 reports different moments of the distribution of own-price elasticities and price-cost ratios implied by our demand estimates. Overall, the median own-price elasticity is around -5.5, while the median price-cost ratio is close to 1.23. The table also shows substantial heterogeneity across destinations. High-income markets such as the United States and Germany have elasticities around -6 to -7, whereas lower-income markets such as India, Brazil, and Mexico are less elastic and have larger price-cost ratios. This heterogeneity follows from the estimated destination-specific price sensitivity in equation (33) and the Cournot markup factor in equation (34).

Figure 1 decomposes the origin component of marginal cost for firms headquartered in the United States, China and Germany. Each panel reports a multiplicative cost component, normalized to the corresponding value for a U.S.-headquartered firm producing in the United States. Values above one therefore indicate a higher origin component than the U.S.-HQ/U.S.-origin benchmark, while values below one indicate a lower component. Panel (a) combines the estimated production-origin fixed effect, the headquarters-origin geography terms, and the parts-tariff com-

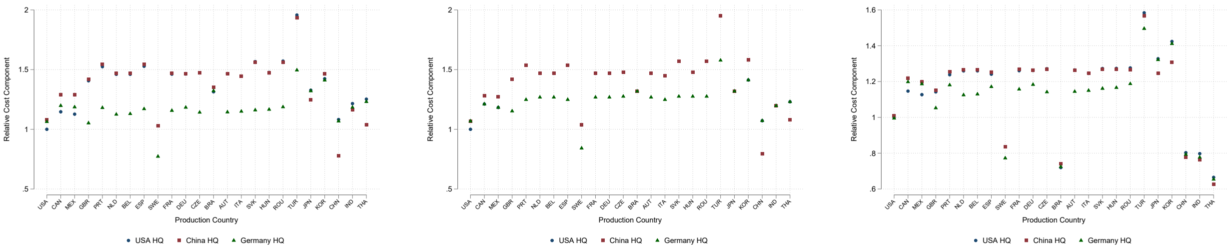

(a) Full origin component **(b) Origin FE and parts tariffs** **(c) Origin FE and distance**

Figure 1: Marginal cost origin components, each panel normalized to its own U.S.-HQ/U.S.-origin value

ponent. Panel (b) removes the geography terms and shows how the origin fixed effect and parts-tariff exposure vary across production countries and headquarters countries. Panel (c) removes the parts-tariff term and shows the contribution of the origin fixed effect together with headquarters-origin geography. The comparison across panels shows that origin cost differences do not come from a single source: a production country can look attractive because of its estimated origin fixed effect, because it is close to a firm's headquarters country, or because its parts-tariff exposure is favorable for that headquarters-origin pair.

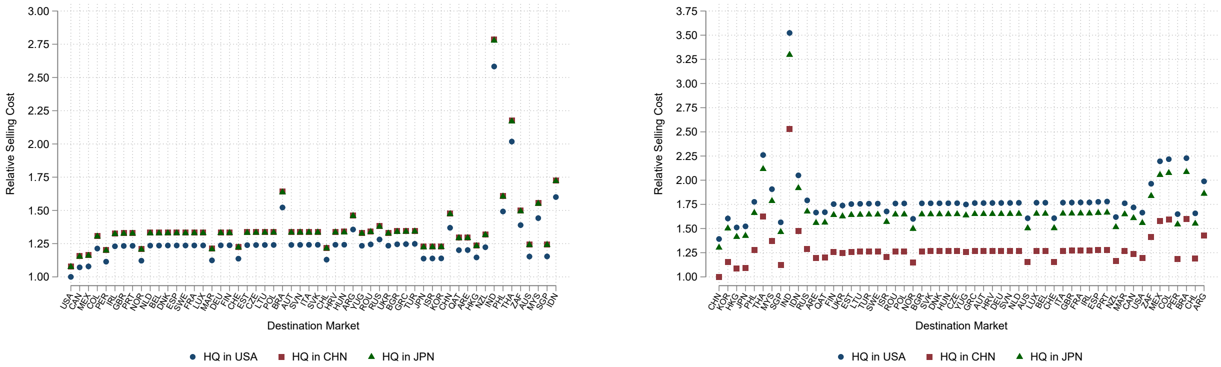

(a) Produced in U.S. **(b) Produced in China**

Figure 2: Relative marginal selling cost by destination for brands headquartered in the United States, China and Japan, and producing in the United States and Japan.

Figures 2 and 3 translate the estimated marginal-cost equation into delivered-cost comparisons. These figures hold fixed model characteristics and brand productivity, so they should be read as the component of the delivered marginal costs coming from the tariffs on final goods, transport costs, origin, headquarters-origin geography, and parts-tariff terms rather than as costs for any particular brand-model pair. Figure 2 fixes the production origin, first at the United States and then at China, and compares the relative cost of serving different destination markets for firms headquartered in the United States, China, and Japan. Within each panel, costs are normalized by the lowest plotted value, so the vertical axis measures how much more expensive a destination is relative to the cheapest destination-origin-HQ combination.

Figure 3 fixes the destination market, first at the United States and then at China, and compares

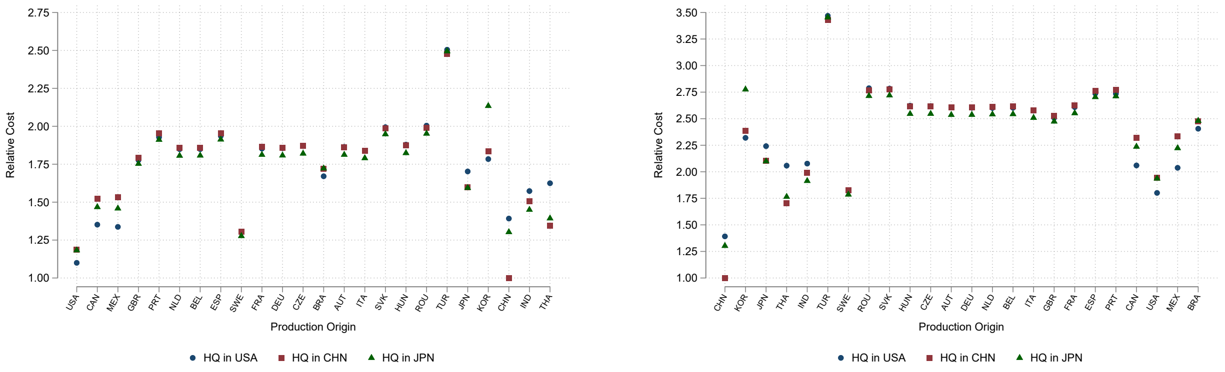

(a) Sold in U.S. **(b) Sold in China**

Figure 3: Relative marginal cost by production origin for selling in the United States and China for firms headquartered in the United States, China and Japan.

the relative cost of sourcing from different production origins for the same set of headquarters countries. This comparison corresponds to the sourcing decision made by the firms in our model: for a given destination, a firm compares feasible production origins and chooses the one with the lowest delivered marginal cost.

5.2 Fixed Costs

Models with strategic entry are generally incomplete. Given a vector of model parameters, entry games typically have multiple equilibria. In the absence of mixed strategies, not even existence is generally guaranteed. In such settings, without additional model restrictions (e.g., symmetric payoffs/interdependencies, or order-of-entry restrictions), the model cannot be simulated, rendering standard SMM or MLE techniques infeasible.

In such a setting, the literature on entry games has turned to moment inequalities for estimation of fixed-cost parameters. With multiple binary decisions both within and across-firms, an additional challenge emerges: solving for all the equilbria of the model is not computationally feasible. Therefore, moment-inequality approaches in the style of Ciliberto and Tamer (2009), which require the computation of all Nash equilbria at each candidate parameter vector, are not feasible.

We develop a novel moment inequality procedure for entry games with many interdependent binary decisions. Our approach generalizes the moment inequalities in Fan and Yang (2025) by incorporating common knowledge of rationality and iteratively eliminating dominated strategies, rather than only keeping strategies that survive level-1 rationality. Therefore, our approach provides parameter bounds that are less slack than those in Fan and Yang (2025).

We leverage our solution algorithm to build moment inequalities. Letting $P_F(\nu, \gamma)$ denote the set of action profiles that survive both our first-step and second-step procedures given a vector of fixed-cost realizations ν and a vector of fixed cost parameters γ, define:

$$\overline{\Delta}_i^{f(i)}(\nu, \gamma) = \max_{C \in P_F(\nu,\gamma)} \Delta_i^{f(i)}(C; \nu, \gamma), \quad \text{and} \quad \underline{\Delta}_i^{f(i)}(\nu, \gamma) = \min_{C \in P_F(\nu,\gamma)} \Delta_i^{f(i)}(C; \nu, \gamma). \tag{38}$$

Since the model predicts that any observed action profile C^* satisfies $C^* \in P_F(\nu,\gamma)$, equation (38) implies that, for any $i \in J$,

$$\mathbb{1}\{\overline{\Delta}_i^{f(i)}(\nu,\gamma) \geq 0\} - C_i^* \geq 0, \qquad \text{and} \qquad \mathbb{1}\{\underline{\Delta}_i^{f(i)}(\nu,\gamma) \geq 0\} - C_i^* \leq 0. \tag{39}$$

Equation (39) relates IEDS to inequalities that hold at any $i \in J$, and at any realization of fixed cost shocks ν. Economically, these inequalities say that if $C_i^* = 1$ in the data, then necessarily the largest marginal value of such a choice across all rationalizable action profiles in the game must be weakly positive. Similarly, if $C_i^* = 0$ in the data, then necessarily the smallest marginal value of choice i across all rationalizable profiles must be weakly negative.

As $P_F(\nu,\gamma) \subseteq \mathcal{B}^{|J|}$ for all (ν,γ), we have

$$\overline{\Delta}_i^{f(i)}(\nu,\gamma) \leq \max_{C \in \mathcal{B}^{|J|}} \Delta_i^{f(i)}(C,\nu,\gamma) \qquad \text{and} \qquad \underline{\Delta}_i^{f(i)}(\nu,\gamma) \geq \min_{C \in \mathcal{B}^{|J|}} \Delta_i^{f(i)}(C,\nu,\gamma).$$

Thus, our moment inequalities will be tighter than those obtained following the procedure in Fan and Yang (2025), which use the maximum and minimum marginal value of any binary decision, across all possible action profiles (and not only the rationalizable subset), to construct their moment inequalities.

Denote by,

$$\overline{C}_i(\nu,\gamma) = \mathbb{1}\{\overline{\Delta}_i^{f(i)}(\nu,\gamma) \geq 0\} \qquad \text{and} \qquad \underline{C}_i(\nu,\gamma) = \mathbb{1}\{\underline{\Delta}_i^{f(i)}(\nu,\gamma) \geq 0\}. \tag{40}$$

and group into $\overline{C}_f(\nu,\gamma)$ all elements $\overline{C}_i(\nu,\gamma)$ that correspond to firm f actions. Then, we can rewrite the inequalities in equation (39) as,

$$\overline{C}_i(\nu,\gamma) - C_i^* \geq 0 \qquad \text{and} \qquad \underline{C}_i(\nu,\gamma) - C_i^* \leq 0.$$

and, given an increasing function $g : \mathcal{B}^{|J|} \to \mathbb{R}$, we can write the following inequalities,

$$\sum_{f \in \mathcal{F}} \{R^{-1} \sum_{r=1}^{R} g(\underline{C}_f(\nu^r,\theta))\} \leq \sum_{f \in \mathcal{F}} g(C_f^*) \leq \sum_{f \in \mathcal{F}} \{R^{-1} \sum_{r=1}^{R} g(\overline{C}_f(\nu^r,\theta))\}, \tag{41}$$

with ν_f^r a random draw from the distribution of fixed-cost shocks.[9] Thus, we can apply moment inequalities that are monotone in firms' binary choices.

We implement (41) using moment functions that are weakly increasing in the product-level production and sales decisions. The unit of observation is the product m, and the moments are organized into five groups.

---

[9] Recall that we assume that the distribution of market entry and plant location fixed costs is a mixture of a log-normal distribution and a multinomial distribution determining the free location.

1. Total markets entered.

$$g_m^1(\boldsymbol{I}_m, \boldsymbol{D}_m) = \sum_{n=1}^{N} I_{mn}. \tag{42}$$

This moment captures the overall propensity of a product to be sold across destination markets and is therefore informative about the market entry fixed cost. If entry fixed costs are too low, the simulated lower bound tends to exceed the observed moment; if they are too high, the simulated upper bound tends to fall short of it.

2. Entered destinations in destination-profitability bin k.

$$g_m^{2,k}(\boldsymbol{I}_m, \boldsymbol{D}_m) = \sum_{n=1}^{N} I_{mn}\, \mathbb{1}\{\beta_{mn}^d = k\}, \qquad k \in \{2, 3, 4, 5, 6, 7, 8\}. \tag{43}$$

The index β_{mn}^d assigns each product-destination pair to a profitability octile. Specifically, we rank product-destination pairs by the marginal variable profit of sales entry evaluated at a reference parameter vector γ^{ref}, chosen as the midpoint of the confidence set reported in Castro-Vincenzi et al. (2023):

$$\bar{\Delta}_{mn}^e(\gamma^{\text{ref}}) = \frac{1}{S^{\text{ref}}} \sum_{s=1}^{S^{\text{ref}}} \frac{1}{n_\xi} \sum_{k=1}^{n_\xi} \Delta\pi_{n,ks}^m. \tag{44}$$

$\Delta\pi_{n,ks}^m$ denotes the Cournot variable-profit gain in destination market n from switching product m's sales decision from 0 to 1, evaluated at delivered marginal cost $mc_{mn,o_{sk}^*(m)}$, where $o_{sk}^*(m)$ is the lowest-cost plant active for product m in the lower-bound action profile under γ^{ref} and fixed-cost and demand-shock draws indexed by s and k, $\underline{C}^{sk}(\gamma^{\text{ref}})$. All other own and rival choices are held fixed at that same lower-bound action profile. We use $S^{\text{ref}} = 200$ to compute the bins.

Intuitively, more profitable destinations should be more likely to be entered. The extent to which observed entry aligns with this profitability ranking is informative about the dispersion parameter σ^e: when unobserved entry fixed costs are highly dispersed, observed product offerings appear less closely tied to demand and cost fundamentals.

3. Number of non-HQ production countries.

$$g_m^3(\boldsymbol{I}_m, \boldsymbol{D}_m) = \sum_{o=1}^{L} D_{mo}\, \mathbb{1}\{o \neq h(m)\}. \tag{45}$$

This moment captures the overall propensity to establish production outside headquarters and is informative about the level of production fixed costs. If production plant fixed costs are low, the model is more likely to generate production in non-headquarters locations, so the simulated lower bound of this moment tends to exceed its observed counterpart. If production plant fixed costs are high, the model generates too little non-headquarters production, so the simulated upper bound tends to fall short of the data.

4. Non-HQ production in origin-profitability bin k.

$$g_m^{4,k}(\boldsymbol{I}_m, \boldsymbol{D}_m) = \sum_{o=1}^{L} D_{mo}\, \mathbb{1}\{o \neq h(m)\}\, \mathbb{1}\{\beta_{mo}^{o} = k\}, \qquad k \in \{2, 3, 4, 5, 6, 7, 8\}. \tag{46}$$

The index β_{mo}^{o} assigns each product-origin pair to an origin-profitability octile. Specifically, we rank product-origin pairs by the marginal variable profit of plant entry evaluated at the same reference parameter vector γ^{ref},

$$\bar{\Delta}_{mo}^{p}(\gamma^{\text{ref}}) = \frac{1}{S^{\text{ref}}} \sum_{s=1}^{S^{\text{ref}}} \frac{1}{n_\xi} \sum_{k=1}^{n_\xi} \sum_{n:\, \underline{I}_{mn}^{sk}=1} \Delta\pi_{n,ks}^{mo}. \tag{47}$$

Here, $\Delta\pi_{n,ks}^{mo}$ denotes the Cournot variable-profit gain in destination market n from switching model m's production location from $o_{sk}^{*}(m)$ to o whenever $mc_{mn,o} < mc_{mn,o_s^*(m)}$, and zero otherwise, where $o_{sk}^{*}(m)$ is the lowest-cost plant active for product m in the lower-bound action profile $\underline{C}^{sk}(\gamma^{\text{ref}})$. More profitable non-headquarters origins should be more likely to be selected. As the dispersion of production fixed costs increases, observed production choices become less closely aligned with this ranking. These moments are therefore informative about the dispersion of production fixed costs.

5. Production in headquarters.

$$g_m^{5}(\boldsymbol{I}_m, \boldsymbol{D}_m) = D_{m,h(m)}. \tag{48}$$

This moment measures whether product m is produced in the headquarters country. It is primarily informative about γ_3, which governs the extent to which production is concentrated at headquarters. More negative values of γ_3 make non-headquarters production locations less attractive and therefore increase the model's tendency to produce at home.

In each of the two binned sets of moments, we omit the first octile to avoid a mechanical linear dependence with the corresponding count moment. Altogether, this yields 17 monotone moments and hence 34 unconditional moment inequalities. Each moment is weakly increasing in firms' binary choices, so that the lower and upper action bounds produced by our solution algorithm map directly into lower and upper bounds on the corresponding sample moments.

Table 6 reports the confidence set for the five fixed-cost parameters. The market-entry fixed cost is tightly identified: on the grid, the only accepted values are $\gamma_2 = 3.0$ and $\sigma^e = 3.0$. In contrast, the production-side parameters are less sharply identified, with $\gamma_1 \in [7.0, 15.0]$ and $\sigma^p \in [2.5, 6.5]$. We are continuing to explore alternative moments to sharpen identification along these dimensions. The free-plant parameter is estimated to lie in $\gamma_3 \in [-6.0, -4.0]$, implying a strong home bias in the location of the free plant: headquarters is between $e^4 \approx 55$ and $e^6 \approx 403$ times more likely than any given non-headquarters country to be the free production location; with 24 potential production countries, this corresponds to a headquarters free-plant probability between roughly 70 and 95 percent.

Table 6: Confidence Set for Fixed-Cost Parameters

Parameter	Meaning	Confidence Set
γ_2	Entry FC location	$[3.0,\ 3.0]$
σ^e	Entry FC scale	$[3.0,\ 3.0]$
γ_1	Production FC location	$[7.0,\ 15.0]$
γ_3	Free-plant HQ preference	$[-6.0,\ -4.0]$
σ^p	Production FC scale	$[2.5,\ 6.5]$

Notes: 95% confidence set constructed using the Andrews–Soares (2010) procedure with brand-level clustering. The grid spans $5^5 =$ 3,125 points centered at the CMMS reference point with equal spacing.

Table 7 reports the fit of non-targeted moments. The model matches several broad coverage moments reasonably well. In particular, it comes fairly close to the average number of sales destinations and production locations per model, the number of models sold per market, and the share of the top parent firm in a market. These margins suggest that the estimated fixed costs capture the broad scale of firms' product portfolios and the extent of cross-market participation. At the same time, the model overpredicts internationalization and market concentration. Relative to the data, firms in the model sell and produce in too many countries, export too much, and exhibit too little home bias.

6 The Welfare Effects of U.S. Tariffs on SUVs

In March 2025, the Trump administration invoked Section 232 of the Trade Expansion Act of 1962 to impose a 25 percent ad-valorem tariff on all foreign-produced vehicles on national-security grounds. The concurrent Liberation Day tariffs, announced on April 2, extended broad import duties across virtually all goods, but automobiles were excluded since they had already been covered under Section 232. Still, these actions constitute the most sweeping universal tariff increase on imported cars in decades.

In addition to these global tariffs, the U.S. has experienced recurring trade tensions, particularly with the European Union, in the automobile sector. During Trump's first term, the administration initiated a Section 232 national-security investigation targeting auto imports from the E.U., Japan, and South Korea, and threatened duties of up to 25 percent on European-produced cars. The threat ultimately served as leverage in negotiations rather than being fully enacted. The pattern continued in the second term: although the 2025 Section 232 tariff initially applied universally, the E.U. secured a separate negotiated rate of 15 percent in August 2025, below the universal 25 percent rate applied to other countries. However, in May 2026, Trump threatened to raise E.U.-specific duties back to 25 percent, citing non-compliance. More broadly, the possibility of E.U.-specific automobile duties distinct from the universal Section 232 rate has repeatedly been raised, negotiated, and contested across both administrations.

Motivated by these policies, we consider two tariff counterfactuals. In the first, the United States

Table 7: Model Fit Based on Fixed-Cost Parameters

	Data		Model	
	Unw.	Wtd.	Unw.	Wtd.
Panel A: Coverage				
Models (parent)	9.3	20.6	[9.6, 9.6]	[19.5, 19.9]
Sales countries (model)	12.3	28.4	[11.1, 11.6]	[26.7, 28.2]
Production countries (model)	1.4	2.2	[1.3, 1.6]	[1.8, 2.7]
Models sold (market)	102.3	185.5	[90.8, 95.0]	[181.1, 209.3]
Sales countries (parent)	19.4	46.0	[30.3, 30.9]	[52.7, 53.0]
Prod. countries (parent)	2.9	7.2	[4.5, 5.6]	[10.4, 11.5]
Share exported (model-origin)	32.1%	44.7%	[53.6%, 61.8%]	[60.3%, 67.4%]
HQ production share (parent)	75.3%	49.0%	[53.6%, 63.7%]	[20.4%, 35.4%]
HQ sales share (parent)	67.4%	34.2%	[23.7%, 27.2%]	[15.0%, 18.9%]
Panel B: Market structure and concentration				
Models sold (market)	102.3	185.5	[90.8, 95.0]	[181.1, 209.3]
Parent firms (market)	17.4	26.5	[26.0, 26.5]	[34.1, 36.0]
Brands (market)	33.8	53.0	[45.6, 46.7]	[66.9, 71.7]
Top parent share (market)	27.8%	21.0%	[21.6%, 24.0%]	[21.7%, 24.3%]
Top 3 parent share (market)	58.8%	45.6%	[50.2%, 53.5%]	[49.9%, 54.0%]
Top model share (market)	11.5%	6.2%	[20.2%, 23.4%]	[18.6%, 21.9%]
Top 5 model share (market)	34.8%	22.8%	[55.7%, 61.3%]	[51.5%, 56.8%]

Notes: This table compares the 2019 SUV cross-section ("Data") to the structural baseline averaged across shock draws ("Model"). Each row gives the cross-sectional mean of the listed metric across the unit of observation in parentheses (parent, model, or market). "Unw." is the unweighted mean; "Wtd." weights by quantity sold for parent- and model-level rows and by market size for market-level rows. Bracketed entries $[\underline{\cdot}, \overline{\cdot}]$ are resolution bounds on the model statistic, obtained by setting unresolved (0.5) entries to 0 (lower bound) and 1 (upper bound). Panel A reports coverage of products, sales destinations, and production locations; Panel B reports within-market concentration.

imposes an additional 100 percentage-point ad valorem tariff on all foreign-produced SUVs sold in the U.S. market. In the second, the same increase applies only to E.U.-produced SUVs. Comparing the two counterfactuals isolates how the breadth of a tariff shapes the reorganization of production and offerings in the industry. In theory, a universal tariff tends to push firms to reshore production toward U.S. plants, whereas a discriminatory tariff leaves an escape valve, allowing sourcing to shift toward unaffected foreign locations. As a result, discriminatory tariffs may dampen both the reshoring response and the aggregate welfare effects.

We are interested not only in how tariffs affect outcomes such as sales quantities, prices, producer profits, consumer surplus, and tariff revenue, but also in how much of these adjustments occur through endogenous production relocation. To isolate the contribution of production relocation, we solve each counterfactual under two regimes. In the first regime, firms are allowed to reoptimize both production and sales decisions under the new tariff schedule. In the second regime, firms can adjust only their sales decisions, while production locations are held fixed at their baseline lower-bound levels. Comparing the two regimes reveals how much of the overall response reflects firms' ability to reallocate production across countries.

6.1 Effects of Global Tariffs

Table 8 reports the effects on U.S. households of a global 100% tariff. The counterfactual bounds reflect three sources of uncertainty. First, they capture the residual ambiguity in the IEDS solution arising from binary decisions for which neither 0 nor 1 can be ruled out as rationalizable. Second, they reflect estimation uncertainty in the fixed-cost parameters. In particular, we evaluate the counterfactual at each accepted point in the estimation grid search, and construct bounds by taking the maximum upper bound and minimum lower bound across all accepted parameter values. Third, the bounds also incorporate simulation noise from shock draws. We describe the exact bounding method for each component in Appendix D.

When production relocation is allowed, U.S. consumer surplus falls by \$302–\$569 per household, as the tariff raises the consumer price of imported cars and reduces the set of products available in the U.S. market. Average prices rise from \$28.0–\$29.2 thousand to \$31.8–\$32.3 thousand, and the number of products sold in the United States falls from 238–250 to 165–166 (Table 9).

How the change in variable profits affects U.S. households depends on who captures them: households that hold equity in U.S.-headquartered firms benefit from profits accruing to those firms regardless of where production occurs; households benefit from U.S.-brand profits if brand identity proxies for domestic economic activity; and local investment links more directly to profits earned at U.S.-located plants. We therefore report profit changes under three attribution schemes: variable profits of all U.S.-headquartered firms, of all U.S. brands, and of U.S.-located plants.

The gains are largest for U.S. plants (\$472–\$902 per household) and smaller for U.S. brands (\$184–\$417) and U.S.-headquartered firms (\$128–\$440). This ranking reflects a composition effect: U.S.-headquartered firms and U.S. brands also operate plants abroad, and the tariff reduces the variable profits of those foreign plants by suppressing their U.S. sales. In contrast, the U.S.-plant

Table 8: U.S. 100% Tariff on All Imports: Effect on U.S. Households

	Initial Level	Change With Prod. Relocation	Change Without Prod. Relocation
Panel A: Consumer Surplus per Household (USD)			
	—	[-569, -302]	[-830, -294]
Panel B: Total Variable Profits per Household (USD)			
Firms with US HQs	[1,856, 2,161]	[128, 440]	[-132, 512]
US Brand	[1,672, 1,898]	[184, 417]	[-19, 470]
US Plants	[1,416, 1,815]	[472, 902]	[25, 739]
Panel C: Total Variable Costs of US-Produced Cars per Household (USD)			
	[4,207, 5,479]	[1,472, 2,833]	[168, 2,425]
Panel D: Total Labor Costs of US-Produced Cars per Household (USD)			
	[841, 1,096]	[294, 567]	[34, 485]
Panel E: US Government Tariff Revenue per Household (USD)			
	[43, 68]	[38, 65]	[198, 344]
Panel F: Net Effect per Household (USD), Summing A+B+D+E			
Firms with US HQs	—	[-109, 771]	[-732, 1,046]
US Brand	—	[-53, 748]	[-619, 1,004]
US Plants	—	[235, 1,233]	[-574, 1,273]
Panel G: Share of US-Sold Models Sourced from the US			
	[22.1%, 28.5%]	[13.3pp, 20.2pp]	[3.8pp, 17.8pp]

Notes: This table reports the effect of an additional 100 percentage-point ad-valorem U.S. tariff on every car produced outside the U.S., averaged across 10 shock draws. "Initial Level" reports bounds on the baseline level; "Change With/Without Prod. Relocation" reports bounds on the difference between the counterfactual and the baseline. Panels A–E are per U.S. household (U.S.D); Panel F sums changes in Panels A, B, D, and E to report the net household effect (three rows, one per profit attribution in Panel B); Panel G is the share of U.S.-sold models sourced from U.S. plants (initial in % level, changes in percentage points). Cost and profit aggregates in Panels B–D are net of tariffs. Bounding rules in Appendix D.

measure captures the full reshoring gains accruing to all plants located on U.S. soil, including those owned by foreign-headquartered producers.

The global tariffs induce production reshoring: the share of U.S.-sold products sourced from the United States rises by 13.3–20.2 percentage points. This increase is driven primarily by a sharp contraction in foreign-sourced products, which fall from 174–187 to 95–97, rather than by a large expansion in domestic production: the number of U.S.-sourced products rises only modestly, from 53–71 to 69–70. This net increase in U.S.-sourced products reflects two forces. On the one hand, some firms relocate sourcing of existing products from foreign plants to U.S. plants. On the other hand, the exit of foreign-sourced products weakens competition in the U.S. market, making entry of new U.S.-sourced products profitable at the margin.

Average markups barely change, remaining near \$7.0–\$7.1 thousand. In our model, each firm's markup depends on its aggregate market share. The tariff compresses the market share of foreign-sourced products while expanding that of U.S.-sourced ones, leaving average firm-level market

Table 9: U.S. 100% Tariff on All Imports: Effect on U.S. Prices

	Initial Level	Level With Prod. Relocation	Level Without Prod. Relocation
Panel A: All US Sold Products (USD thousands)			
Num. Products	[238.0, 250.2]	[165.2, 166.0]	[161.0, 176.7]
Avg. Price	[28.0, 29.2]	[31.8, 32.3]	[31.9, 36.0]
Avg. Markup	[7.0, 7.4]	[7.0, 7.1]	[6.7, 7.0]
Avg. Marginal Cost	[20.5, 21.8]	[18.5, 18.7]	[18.3, 19.4]
Panel B: US Sold-US Sourced (USD thousands)			
Num. Products	[52.9, 71.2]	[69.4, 69.9]	[52.1, 70.3]
Avg. Price	[35.2, 39.9]	[37.8, 38.3]	[37.9, 38.7]
Avg. Markup	[8.1, 9.4]	[8.5, 8.6]	[8.2, 8.7]
Avg. Marginal Cost	[27.1, 30.5]	[29.3, 29.7]	[29.5, 30.2]
Panel C: US Sold-Foreign Sourced (USD thousands)			
Num. Products	[174.2, 187.1]	[95.6, 96.6]	[106.4, 108.9]
Avg. Price	[24.0, 26.6]	[27.3, 28.1]	[27.4, 34.8]
Avg. Markup	[6.5, 7.2]	[5.9, 6.0]	[5.9, 6.2]
Avg. Marginal Cost	[17.1, 19.6]	[10.6, 10.9]	[10.6, 14.2]

Notes: This table reports per-model averages in the U.S. market under the same counterfactual as Table 8, averaged across 10 shock draws. Panel A covers all U.S.-sold models; Panel B restricts to U.S.-sold and U.S.-sourced; Panel C to U.S.-sold and foreign-sourced. "Avg. Price" is the tariff-inclusive consumer price; "Avg. Marginal Cost" is the producer cost paid by the firm, tariff exclusive. All values in U.S.D thousands except "Num. Products", which is a model count. Bounding rules in Appendix D.

shares relatively stable. As a result, the increase in U.S. consumer prices primarily reflects changes in marginal costs rather than changes in markups.

These cost changes operate through two channels. First, for each surviving foreign-sourced product, consumer prices rise by the full tariff payment. Second, the price increase from tariff passthroughs is partially moderated by a composition effect: only the more cost-competitive foreign models remain in the U.S. market, pulling down the average price. In particular, the average tariff-exclusive marginal cost falls from \$20.5–\$21.8 thousand to \$18.5–\$18.7 thousand, driven primarily by this composition change.

Another effect of the global tariffs on U.S. households operates through labor income. Since our model tracks total variable costs but does not separately identify their labor component, we take a back-of-the-envelope approach, assuming labor accounts for 20 percent of total variable costs. Under this assumption, the tariff increases labor income by \$294–\$567 per household (Panel D of Table 8), driven by the large expansion of U.S. plant production. By contrast, tariff revenue rises by only \$38–\$65 thousand per household (Panel E), because endogenous relocation shrinks the volume of imports that continue to pay the tariff.

To assess the net effect of the global tariffs on U.S. households, Panel F of Table 8 aggregates the consumer surplus loss with the gains in variable profits, labor income, and tariff revenue. When production relocation is allowed, the net effect is unambiguously positive under the U.S.-plants attribution of variable profits: the average U.S. household gains \$235–\$1,233, because the reshoring-driven gains in profits (\$472–\$902) and labor income (\$294–\$567) more than offset the consumer surplus loss (\$302–\$569). When profits are instead attributed based on ownership of U.S.-headquartered firms or U.S. brands, the resulting intervals include zero, although both remain predominantly positive.

When production relocation is shut down in the last column of both tables, the gains in profits and labor income shrink substantially. U.S. plants gain only \$25–\$739 per household in variable profits, compared to \$472–\$902 when relocation is allowed, while labor income gains fall from \$294–\$567 to \$34–\$485. Panel G confirms the mechanism: the sourcing share of U.S. plants rises by only 3.8–17.8 percentage points, far less than the 13.3–20.2 percentage point increase observed when firms are allowed to relocate production.

Tariff revenue moves in the opposite direction. Because imports remain an important source of supply, the government collects \$198–\$344 per household, roughly four to five times the \$38–\$65 collected when production relocation is allowed. This partially compensates U.S. households for the weaker profit and labor income gains, but not enough to pin down the sign of the overall welfare effect. Panel F shows that the net effect remains ambiguous under all profit-attribution schemes.

6.2 Effects of E.U.-Targeted Tariffs

Table 10 reports the effects of a 100% tariff on E.U.-produced cars on U.S. households. The loss in consumer surplus is much smaller than under the global tariff, as E.U.-produced cars account for only about one-fifth of models sold in the United States. When production relocation is allowed, the change in consumer surplus per household ranges from \$-281 to nearly zero. In some cases, U.S. consumers are essentially unharmed and may even experience a small gain. The intuition is that E.U. firms respond by relocating production to the United States, so these products are priced according to U.S. marginal cost rather than foreign marginal cost inclusive of the tariff. As a result, consumer prices rise only modestly, by about \$200–\$600 per car, compared with \$3,800–\$6,800 under the global tariff.

In addition, among the products that remain foreign-sourced, tariff-induced selection removes relatively high-cost varieties, which further pushes down the average price. Without relocation, by contrast, E.U. cars bear the full tariff burden, surviving imports become more expensive, and more products exit the U.S. market. Consumer surplus therefore falls unambiguously, by \$-504 to \$-30 per household. Relative to the global-tariff case, the sourcing response is also modest: the share of U.S.-sold products sourced from U.S. plants increases by at most 6.6 percentage points and may change little at all.

The tariff-revenue result is nuanced. When E.U. firms relocate production, the U.S. government collects less tariff revenue than before the policy change: Panel E reports a change of \$-45 to \$-13 per household, which is unambiguously negative. The mechanism is relocation-driven erosion of the tariff base. The discriminatory tariff induces E.U. producers to shift assembly to the United States; once production becomes domestic, these cars no longer cross the border and therefore no longer pay import duties. The government thus applies a higher tariff rate to a shrinking import flow. Without relocation, by contrast, tariff revenue increases to \$90–\$219 per household, because E.U. imports remain an important source of E.U.-brand supply.

The declines in tariff revenue are partly offset by gains in profits and labor income. When relocation is allowed, variable profits attributable to U.S. plants increase by \$98–\$585 per household, as E.U. firms shift assembly to existing U.S. facilities and expand output there. Labor income also rises, by \$73–\$363 per household. Even so, the overall effect in Panel F remains ambiguous under all three profit-attribution schemes.

Table 12 shows that Canada could be a third-party beneficiary of the U.S.–E.U. tariff when production relocation is allowed. The share of U.S.-sold models sourced from Canadian plants rises unambiguously by 0.7–2.6 percentage points (Panel G): as E.U.-produced cars exit the U.S. market and some E.U. firms relocate production to U.S. plants, Canadian plants—which face no new tariff—absorb part of the displaced demand. Canadian plant variable profits increase by as much as \$1,228 per household and labor income by as much as \$785 (Panels B and D), although both effects remain ambiguously signed because their lower bounds are slightly negative.

Table 10: U.S. 100% Tariff on E.U. Imports: Effect on U.S. Households

	Initial Level	Change With Prod. Relocation	Change Without Prod. Relocation
Panel A: Consumer Surplus per Household (USD)			
	—	[-281, 9]	[-504, -30]
Panel B: Total Variable Profits per Household (USD)			
Firms with US HQs	[1,856, 2,161]	[-69, 255]	[-258, 365]
US Brand	[1,672, 1,898]	[-10, 234]	[-155, 307]
US Plants	[1,416, 1,815]	[98, 585]	[-169, 489]
Panel C: Total Variable Costs of US-Produced Cars per Household (USD)			
	[4,207, 5,479]	[363, 1,817]	[-516, 1,581]
Panel D: Total Labor Costs of US-Produced Cars per Household (USD)			
	[841, 1,096]	[73, 363]	[-103, 316]
Panel E: US Government Tariff Revenue per Household (USD)			
	[43, 68]	[-45, -13]	[90, 219]
Panel F: Net Effect per Household (USD), Summing A+B+D+E			
Firms with US HQs	—	[-322, 614]	[-774, 870]
US Brand	—	[-262, 593]	[-671, 812]
US Plants	—	[-155, 944]	[-685, 994]
Panel G: Share of US-Sold Models Sourced from the US			
	[22.1%, 28.5%]	[-0.5pp, 6.6pp]	[-4.7pp, 7.3pp]

Notes: This table reports the effect of an additional 100 percentage-point ad-valorem U.S. tariff on every car produced in the E.U., averaged across 10 shock draws. "Initial Level" reports bounds on the baseline level; "Change With/Without Prod. Relocation" reports bounds on the difference between the counterfactual and the baseline. Panels A–E are per U.S. household (U.S.D); Panel F sums changes in Panels A, B, D, and E to report the net household effect (three rows, one per profit attribution in Panel B); Panel G is the share of U.S.-sold models sourced from U.S. plants (initial in % level, changes in percentage points). Cost and profit aggregates in Panels B–D are net of tariffs. Bounding rules in Appendix D.

Table 11: U.S. 100% Tariff on E.U. Imports: Effect on U.S. Prices

	Initial Level	Level With Prod. Relocation	Level Without Prod. Relocation
Panel A: All US Sold Products (USD thousands)			
Num. Products	[238.0, 250.2]	[227.8, 228.5]	[217.6, 238.0]
Avg. Price	[28.0, 29.2]	[28.2, 28.6]	[28.1, 30.2]
Avg. Markup	[7.0, 7.4]	[7.0, 7.1]	[6.8, 7.0]
Avg. Marginal Cost	[20.5, 21.8]	[20.5, 20.7]	[20.2, 21.3]
Panel B: US Sold-US Sourced (USD thousands)			
Num. Products	[52.9, 71.2]	[64.1, 65.3]	[51.7, 69.9]
Avg. Price	[35.2, 39.9]	[38.8, 39.9]	[38.2, 38.9]
Avg. Markup	[8.1, 9.4]	[8.8, 9.1]	[8.5, 8.9]
Avg. Marginal Cost	[27.1, 30.5]	[30.0, 30.8]	[29.6, 30.1]
Panel C: US Sold-Foreign Sourced (USD thousands)			
Num. Products	[174.2, 187.1]	[162.6, 164.4]	[160.9, 168.1]
Avg. Price	[24.0, 26.6]	[23.7, 24.5]	[23.7, 27.6]
Avg. Markup	[6.5, 7.2]	[6.2, 6.4]	[6.2, 6.3]
Avg. Marginal Cost	[17.1, 19.6]	[16.5, 16.9]	[16.1, 18.5]

Notes: This table reports per-model averages in the U.S. market under the same counterfactual as Table 10, averaged across 10 shock draws. Panel A covers all U.S.-sold models; Panel B restricts to U.S.-sold and U.S.-sourced; Panel C to U.S.-sold and foreign-sourced. "Avg. Price" is the tariff-inclusive consumer price; "Avg. Marginal Cost" is the producer cost paid by the firm, tariff exclusive. All values in U.S.D thousands except "Num. Products", which is a model count. Bounding rules in Appendix D.

Canadian household consumer surplus is also ambiguous, ranging from \$-133 to \$114 per household (Panel A). This reflects offsetting forces. On the one hand, Canadian consumers may benefit from the reallocation of production to North America, which can lower the trade cost. On the other hand, they may be harmed if the E.U.-produced varieties become less available.

Taking stock, the global and E.U.-specific tariff experiments operate through distinct channels and have very different consequences. Under the global 100% tariff, large-scale reshoring unambiguously benefits U.S. households when profits are attributed to U.S. plants. By contrast, the E.U.-only tariff induces more targeted relocation, which nearly eliminates consumer harm but also erodes the tariff base, leaving the overall net effect ambiguous.

More broadly, the comparison between the two scenarios shows that broad tariffs work mainly through large sourcing reallocations and product exit, generating substantial reshoring and sizable shifts in domestic profits and labor income. An E.U.-specific tariff, by contrast, primarily reallocates sourcing across foreign locations and has much smaller overall effects on U.S. households. Finally, the effect on Canada show that even countries not directly targeted by the tariff can experience non-trivial spillovers, highlighting the wider cross-border consequences of bilateral trade policy.

Table 12: U.S. 100% Tariff on E.U. Imports: Effect on Canada Households

	Initial Level	Change With Prod. Relocation	Change Without Prod. Relocation
Panel A: Consumer Surplus per Household (USD)			
	—	[-133, 114]	[-243, 202]
Panel B: Total Variable Profits per Household (USD)			
Canadian Plants	[173, 1,228]	[-229, 1,228]	[-1,013, 790]
Panel C: Total Variable Costs of Canadian-Produced Cars per Household (USD)			
	[386, 3,491]	[-129, 3,926]	[-3,017, 2,685]
Panel D: Total Labor Costs of Canadian-Produced Cars per Household (USD)			
	[77, 698]	[-26, 785]	[-603, 537]
Panel E: Canadian Government Tariff Revenue per Household (USD)			
	[75, 106]	[-27, 5]	[-30, 19]
Panel F: Net Effect per Household (USD), Summing A+B+D+E			
Canadian Plants	—	[-415, 2,132]	[-1,888, 1,549]
Panel G: Share of US-Sold Models Sourced from Canada			
	[1.6%, 3.2%]	[0.7pp, 2.6pp]	[-1.1pp, 1.8pp]

Notes: This table reports the effect of the same counterfactual as Table 10 on Canadian households, averaged across 10 shock draws. Panels A–E are per Canadian household (U.S.D), with Panel B summing variable profits attributable to Canadian plants, Panels C–D restricting cost aggregates to Canadian-produced cars, and Panel E reporting Canadian-government tariff revenue collected from Canada-destined cars; Panel F sums changes in Panels A, B, D, and E to report the net household effect; Panel G is the share of U.S.-sold models sourced from Canadian plants (initial in % level, changes in percentage points). Cost and profit aggregates in Panels B–D are net of tariffs. Bounding rules in Appendix D.

7 Conclusion

We develop a quantitative framework to study how policies shape the global organization of production and offerings in multinational oligopolies. Our methodological contribution is to provide a solution method for entry games that may not be neither supermodular nor submodular; we only require that for any pair binary actions in the game, the researcher can establish whether they are (strategic) substitutes or complements, independepently of third choices, for each of the additive components of firms' payoffs. Using our framework, we quantify the importance of production reallocation in shaping the global consequences of U.S. tariffs in the SUV market. A few caveats remain, which we plan to address in future research: choices under dynamics and uncertainty or economies of scope, and strategic policies and retaliation.

References

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

Bernheim, B. Douglas, “Rationalizable Strategic Behavior,” *Econometrica*, 1984, *52* (4), 1007–28.

Brander, James A., “Strategic Trade Policy,” in Gene M. Grossman and Kenneth Rogoff, eds., *Handbook of International Economics*, Vol. 3, Elsevier, 1995, chapter 27, pp. 1395–1455.

Bresnahan, Timothy F. and Peter Reiss, “Entry and Competition in Concentrated Markets,” *The Journal of Political Economy*, 1991, *99* (5), 977–1009.

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

_ , Eugenia Menaguale, Eduardo Morales, and Alejandro Sabal, “Market Entry and Plant Location in Multiproduct Firms,” *mimeo*, 2023.

Ciliberto, Federico and Elie Tamer, “Market Structure and Multiple Equilibria in the Airline Markets,” *Econometrica*, 2009, *77*, 1791–1828.

Coşar, Kerem A., Paul L. E. Grieco, Shengyu Li, and Felix Tintelnot, “What Drives Home Market Advantage?,” *Journal of International Economics*, 2018, *110*, 135–150.

de Paula, Aureo, “Econometric Analysis of Games with Multiple Equilibria,” *Annual Review of Economics*, 2013, *5*, 107–131.

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.

_ and _ , “Estimating Discrete Games with Many Firms and Many Decisions: An Application to Merger and Product Variety,” *Journal of Political Economy*, June 2025, *133* (6), 1886–1944.

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*, 05 2024, *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.

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

Pearce, David G., "Rationalizable Strategic Behavior and the Problem of Perfection," *Econometrica*, 1984, *52* (4), 1029–1050.

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.

A Interdependence Properties

In this section, we first study equilibrium derivatives and cross-partials in our multinomial logit Cournot model. We then apply these results to prove that Assumption 1 holds in our model.

A.1 Equilibrium Variable Profits Under Multi-Product MNL Cournot

Write total variable profits for firm f, omitting country subscripts.

$$\pi_f = \sum_{j \in \Omega_f} (p_j(q) - c_j) q_j,$$

where under multinomial logit demand, the inverse demand function is,

$$p_j(q) = \frac{1}{\alpha}[\delta_j - \log(q_j) + \log(M - Q)]. \tag{49}$$

and $Q := \sum_{j \in \Omega} q_j$. The first-order condition with respect to q_j is,

$$\begin{aligned} p_j(q^*) - c_j &= \frac{1}{\alpha}[\frac{1}{q_j^*}]q_j^* + \frac{1}{\alpha} \sum_{k \in \Omega_f} \frac{q_k^*}{M - Q^*} \\ &= \frac{1}{\alpha}[1 + \frac{Q_f^*}{M - Q^*}]. \end{aligned} \tag{50}$$

This implies the same absolute markup charged by the firm across its products. Due to this constant markup result, we obtain the following expression for equilibrium profits as a function of equilibrium aggregates Q_f^* and Q^*,

$$\pi_f^*(Q_f^*, Q^*) = \frac{1}{\alpha}[Q_f^* + \frac{(Q_f^*)^2}{M - Q^*}]. \tag{51}$$

Combining equations (49) and (50) yields,

$$q_j^* = \exp(\delta_j - \alpha c_j)(M - Q^*) \exp(-1 - \frac{Q_f^*}{M - Q^*}). \tag{52}$$

Summing across all $j \in \Omega_f$ obtains,

$$\frac{Q_f^*}{M - Q^*} \exp(\frac{Q_f^*}{M - Q^*}) = K_f, \tag{53}$$

where $K_f = \sum_{j \in \Omega_f} \exp(\delta_j - \alpha c_j - 1)$. The equations in (53) solve for firm total quantities given their quality-cost index K_f. Given the solution for firm quantities and product-specific marginal costs, we obtain equilibrium prices. Moreover, note that equation (53) can be simplified using the

Lambert W inverse function,

$$x_f := \frac{Q_f^*}{M - Q^*} = W(K_f). \tag{54}$$

Next, define $R := M - Q^*$, and $S := \sum_h x_h$. Then $Q_f^* = Rx_f$ and $Q^* = RS = M - R$, and so $R = \frac{M}{1+S}$. It follows that,

$$\pi_f^* = \frac{1}{\alpha} R x_f [1 + x_f] = \frac{M}{\alpha} \frac{x_f(1 + x_f)}{1 + S}. \tag{55}$$

We now study how changes in K_f and K_g (for $g \neq f$) affect firm f's equilibrium profits, and the curvature of equilibrium profits.

First, we record the first derivative of x_f with respect to K_f, which yields,

$$\frac{\partial x_f}{\partial K_f} = \frac{x_f}{K_f[1 + x_f]} > 0.$$

Differentiating again with respect to K_f, we obtain,

$$\frac{\partial^2 x_f}{\partial K_f^2} = -\frac{x_f^2(x_f + 2)}{K_f^2(1 + x_f)^3} < 0. \tag{56}$$

Also, note that, conveniently, $\partial x_f / \partial K_g = 0$ for all $g \neq f$. Introduce further notation and let $n_f := x_f(1 + x_f)$. Then,

$$\begin{aligned}
\frac{\partial \pi_f^*}{\partial K_f} &= \frac{M}{\alpha} \frac{\partial x_f}{\partial K_f} \Big[\frac{n_f'}{1 + S} - \frac{n_f}{(1 + S)^2}\Big] && (57)\\
&= \frac{M}{\alpha} \frac{\partial x_f}{\partial K_f} \Big[\frac{(1 + 2x_f)(1 + x_f + T) - x_f(1 + x_f)}{(1 + S)^2}\Big] && (58)\\
&= \frac{M}{\alpha} \frac{\partial x_f}{\partial K_f} \Big[\frac{(1 + x_f)^2 + (1 + 2x_f)T}{(1 + S)^2}\Big] > 0, && (59)
\end{aligned}$$

where $T := \sum_{h \neq f} x_h$. Moreover, for $g \neq f$,

$$\frac{\partial \pi_f^*}{\partial K_g} = -\frac{M}{\alpha} \frac{x_f(1 + x_f)}{(1 + S)^2} \frac{\partial x_g}{\partial K_g} < 0, \tag{60}$$

since $\frac{\partial S}{\partial K_g} = \frac{\partial x_g}{\partial K_g}$.

We are now ready to study the second derivatives. First, let's consider the cross-partial $\frac{\partial^2 \pi_f^*}{\partial K_f \partial K_g}$. Starting from equation (59) and differentiating with respect to K_g, we obtain,

$$\frac{\partial^2 \pi_f^*}{\partial K_f \partial K_g} = \frac{M}{\alpha} \frac{\partial x_f}{\partial K_f} \frac{\partial x_g}{\partial K_g} \frac{(1 + 2x_f)(1 + S) - 2[(1 + x_f)^2 + (1 + 2x_f)T]}{(1 + S)^3}, \tag{61}$$

where we have used the fact that $\partial x_f/\partial K_f$ is a constant with respect to K_g, and $\partial S/\partial K_g = \partial x_g/\partial K_g$. Simplifying equation (61) further, we obtain,

$$\begin{aligned}
\frac{\partial^2\pi_f^*}{\partial K_f\partial K_g} &= \frac{M}{\alpha}\frac{\partial x_f}{\partial K_f}\frac{\partial x_g}{\partial K_g}\frac{(1+2x_f)(1+x_f-T)-2(1+x_f)^2}{(1+S)^3} \\
&= \frac{M}{\alpha}\frac{\partial x_f}{\partial K_f}\frac{\partial x_g}{\partial K_g}\frac{x_f(1+x_f)-(1+x_f)^2-(1+2x_f)T}{(1+S)^3} \\
&= -\frac{M}{\alpha}\frac{\partial x_f}{\partial K_f}\frac{\partial x_g}{\partial K_g}\frac{1+x_f+(1+2x_f)T}{(1+S)^3} < 0. \qquad (62)
\end{aligned}$$

Next, we consider the second derivative $\frac{\partial^2\pi_f^*}{\partial K_f^2}$. For this purpose, it is useful to first record the first and second derivatives of π_f^* with respect to x_f. Differentiating equation (55) with respect to x_f, yields,

$$\frac{\partial\pi_f^*}{\partial x_f} = \frac{M}{\alpha}\frac{(1+2x_f)(1+S)-x_f(1+x_f)}{(1+S)^2} = \frac{M}{\alpha}\frac{(1+2x_f)(1+T)+x_f^2}{(1+S)^2} > 0. \qquad (63)$$

Differentating once more obtains,

$$\frac{\partial^2\pi_f^*}{\partial x_f^2} = 2\frac{M}{\alpha}\frac{(1+S)^2-x_f^2-(1+T)-2x_f-2x_fT}{(1+S)^3} = 2\frac{M}{\alpha}\frac{T(1+T)}{(1+S)^3} > 0. \qquad (64)$$

Next, we use the chain rule to write,

$$\frac{\partial^2\pi_f^*}{\partial K_f^2} = \frac{\partial^2\pi_f^*}{\partial x_f^2}\left(\frac{\partial x_f}{\partial K_f}\right)^2 + \frac{\partial\pi_f^*}{\partial x_f}\frac{\partial^2 x_f}{\partial K_f^2}. \qquad (65)$$

Combining equation (65) with equations (56), (59), (63), and (64), we obtain,

$$\begin{aligned}
\frac{\partial^2\pi_f^*}{\partial K_f^2} &= \frac{M}{\alpha}\frac{x_f^2[2T(1+T)(1+x_f)-((1+2x_f)(1+T)+x_f^2)(x_f+2)(1+S)]}{K_f^2(1+S)^3(1+x_f)^3} \qquad (66) \\
&= \frac{Mx_f^2}{\alpha}\frac{N}{K_f^2(1+S)^3(1+x_f)^3}. \qquad (67)
\end{aligned}$$

Simplifying the numerator term N, we obtain,

$$\begin{aligned}
N &= 2T(1+T)(1+x_f)-(1+x_f+x_f)(1+T)(x_f+2)(1+x_f+T)-x_f^2(x_f+2)(1+x_f+T) \\
&= 2T(1+T)(1+x_f)-T(1+x_f)(1+T)(x_f+2)-(1+x_f)(1+T)(x_f+2)(1+x_f) \\
&\quad -x_f(1+T)(x_f+2)(1+x_f)-Tx_f(1+T)(x_f+2)-x_f^2(x_f+2)(1+x_f+T) \\
&= -Tx_f(1+x_f)(1+T)-(1+x_f)(1+T)(x_f+2)(1+x_f) \\
&\quad -x_f(1+T)(x_f+2)(1+x_f)-Tx_f(1+T)(x_f+2)-x_f^2(x_f+2)(1+x_f+T) < 0.
\end{aligned}$$

Since $N < 0$, it follows that

$$\frac{\partial^2 \pi_f^*}{\partial K_f^2} < 0. \tag{68}$$

A.2 Interdependence Properties under $\{\Pi_f\}$

To prove the sign of the interdependence across any pair of decisions in our model, we leverage the representation of equilibrium variable profits in destination country n as a function of sufficient $\{K_{gn}\}_{g\in\mathcal{F}}$,

$$\pi_{fn}(K_{fn}, K_{-fn}) = \frac{M_n}{\alpha_n} \frac{W(K_{fn})(1+W(K_{fn}))}{1+\sum_{g\in\mathcal{F}} W(K_{gn})}, \tag{69}$$

where $K_{fn} = K_{fn}(I_{fn}, D_f) = \sum_{m\in\Omega_f} I_{mn} \exp(\delta_{mn} - \alpha_n c_{mn}(D_f) - 1)$. We use the property that $K_{fn}(I_{fn}, D_f)$ is increasing in I_{fn} and D_f.

A.2.1 (I_f, D_f) and (I_g, D_g) are substitutes for $f \neq g$

Consider the marginal value of decision I_{mn} for firm f with $m \in \Omega_f$,

$$\begin{aligned}\Delta^f_{I_{mn}}(I_f, D_f, I_{-f}, D_{-f}) &= \pi_{fn}(I_{fn}^{m\to 1}, D_f, I_{-fn}, D_{-f}) - \pi_{fn}(I_{fn}^{m\to 0}, D_f, I_{-fn}, D_{-f}) - F^e_{mn} \\ &= \pi_{fn}(K_{fn}(I_{fn}^{m\to 1}, D_f), K_{-fn}) - \pi_{fn}(K_{fn}(I_{fn}^{m\to 0}, D_f), K_{-fn}) - F^e_{mn} \\ &= \pi_{fn}(\overline{K}_{fn}, K_{-fn}) - \pi_{fn}(\underline{K}_{fn}, K_{-fn}) - F^e_{mn}, \end{aligned} \tag{70}$$

where in equality (70), $\overline{K}_{fn} := K_{fn}(I_{fn}^{m\to 1}, D_f) \geq K_{fn}(I_{fn}^{m\to 0}, D_f) =: \underline{K}_{fn}$. Note that K_{gn} is also increasing in D_g and I_{gn}. To prove that I_{mn} exhibits negative cross-partials with I_{-fn} and D_{-f}, it is therefore sufficient to have that $\frac{\partial^2 \pi_{fn}}{\partial K_{fn} \partial K_{gn}} < 0$ for $f \neq g$, which holds due to equation (62). Clearly, equation (70) also shows that I_{mn} is independent of any I_{kt} for any product k and any $t \neq n$.

Next, consider the marginal value of decision D_{mo} for firm f with $m \in \Omega_f$,

$$\begin{aligned}\Delta^f_{D_{mo}}(I_f, D_f, I_{-f}, D_{-f}) &= \sum_{n\in\mathcal{N}} [\pi_{fn}(I_{fn}, D_f^{mo\to 1}, I_{-fn}, D_{-f}) - \pi_{fn}(I_{fn}, D_f^{mo\to 0}, I_{-fn}, D_{-f})] - F^p_{mo} \\ &= \sum_{n\in\mathcal{N}} [\pi_{fn}(K_{fn}(I_{fn}, D_f^{mo\to 1}), K_{-fn}) - \pi_{fn}(K_{fn}(I_{fn}, D_f^{mo\to 0}), K_{-fn})] - F^p_{mo} \\ &= \sum_{n\in\mathcal{N}} [\pi_{fn}(\overline{K}_{fn}, K_{-fn}) - \pi_{fn}(\underline{K}_{fn}, K_{-fn})] - F^p_{mo}, \end{aligned} \tag{71}$$

where in equality (71), $\overline{K}_{fn} := K_{fn}(I_{fn}, D_f^{mo\to 1}) \geq K_{fn}(I_{fn}, D_f^{mo\to 0}) =: \underline{K}_{fn}$. Note that K_{gn} is also increasing in D_g and I_{gn}. To prove that I_{mn} exhibits negative cross-partials with I_{-fn} and D_{-f}, it is therefore sufficient to have that $\frac{\partial^2 \pi_{fn}}{\partial K_{fn} \partial K_{gn}} < 0$ for $f \neq g$, which holds due to equation (62). This concludes the proof.

A.2.2 I_{mn} and D_{mo} are complements for any $n, o \in \mathcal{N}$

To prove this property, we write,

$$\begin{aligned}\Delta^f_{I_{mn}}(I_f, D_f, I_{-f}, D_{-f}) &= \pi_{fn}(I_{fn}^{m\to 1}, D_f, I_{-fn}, D_{-f}) - \pi_{fn}(I_{fn}^{m\to 0}, D_f, I_{-fn}, D_{-f}) - F^e_{mn}\\ &= \pi_{fn}(I_{fn}^{m\to 1}, D_f, K_{-fn}) - \pi_{fn}(I_{fn}^{m\to 0}, D_f, K_{-fn}) - F^e_{mn} \qquad (72)\end{aligned}$$

Consider now the change in $\Delta^f_{I_{mn}}(I_f, D_f, I_{-f}, D_{-f})$ when a plant for product m is opened in origin o, i.e.,

$$\begin{aligned}&\Delta^f_{I_{mn}}(I_f, D_f^{mo\to 1}, I_{-f}, D_{-f}) - \Delta^f_{I_{mn}}(I_f, D_f^{mo\to 0}, I_{-f}, D_{-f})\\ &\quad= \pi_{fn}(I_{fn}^{m\to 1}, D_f^{mo\to 1}, K_{-fn}) - \pi_{fn}(I_{fn}^{m\to 0}, D_f^{mo\to 1}, K_{-fn})\\ &\quad\quad- [\pi_{fn}(I_{fn}^{m\to 1}, D_f^{mo\to 0}, K_{-fn}) - \pi_{fn}(I_{fn}^{m\to 0}, D_f^{mo\to 0}, K_{-fn})]\\ &\quad= \pi_{fn}(I_{fn}^{m\to 1}, D_f^{mo\to 1}, K_{-fn}) - \pi_{fn}(I_{fn}^{m\to 1}, D_f^{mo\to 0}, K_{-fn}) \qquad (73)\\ &\quad= \pi_{fn}(\overline{K}_{fn}, K_{-fn}) - \pi_{fn}(\underline{K}_{fn}, K_{-fn}) \qquad (74)\end{aligned}$$

where equality (73) follows from the fact that if $I_{mn} = 0$, then π_{fn} is independent of production location decisions for product m. In equality (74), $\overline{K}_{fn} \geq \underline{K}_{fn}$ follows from the fact that K_{fn} is increasing in D_f. Therefore, it suffices to show that π_{fn} is increasing in K_{fn}, which holds due to equation (59). This concludes the proof.

A.2.3 I_{mn} and I_{kn} are substitutes for $m, k \in \Omega_f$

To prove this property, we build on equation (72) and write,

$$\begin{aligned}\Delta^f_{I_{mn}}(I_f, D_f, I_{-f}, D_{-f}) &= \pi_{fn}(I_{fn}^{m\to 1}, D_f, K_{-fn}) - \pi_{fn}(I_{fn}^{m\to 0}, D_f, K_{-fn}) - F^e_{mn}\\ &= \pi_{fn}(K_{fn}(I_{fn}^{m\to 1}, D_f), K_{-fn}) - \pi_{fn}(K_{fn}(I_{fn}^{m\to 0}, D_f), K_{-fn}) - F^e_{mn}\\ &= \pi_{fn}(\overline{K}_{fn}, K_{-fn}) - \pi_{fn}(\underline{K}_{fn}, K_{-fn}) - F^e_{mn}, \qquad (75)\end{aligned}$$

where in equality (75), $\overline{K}_{fn} := K_{fn}(I_{fn}^{m\to 1}, D_f) \geq K_{fn}(I_{fn}^{m\to 0}, D_f) =: \underline{K}_{fn}$. Recall that, $K_{fn}(I_{fn}, D_f) = \sum_{m\in\Omega_f} I_{mn} \exp(\delta_{mn} - \alpha_n c_{mn}(D_f) - 1)$. Therefore, holding equal production choices, the increase in K_{fn} that results from adding a product offering (i.e., switching from $I_{mn} = 0$ to $I_{mn} = 1$) is constant and equal to $\exp(\delta_{mn} - \alpha_n c_{mn} - 1)$. Therefore, establishing whether I_{mn} and I_{kn} are substitutes (for $m, k \in \Omega_f$) is equivalent to asking whether the change in equilibrium variable profits that arises from a constant change in K_{fn} is larger or smaller, when K_{fn} is larger or smaller. Due to equation (68), we have that $\frac{\partial^2 \pi_{fn}}{\partial K_{fn}^2} < 0$, which shows that $\Delta^f_{I_{mn}}$ is decreasing in I_{fn}. This concludes the proof.

A.2.4 I_{mn} and D_{ko} are substitutes for $m, k \in \Omega_f$, and any $o, n \in \mathcal{N}$

This proof is identical to the proof in subsection A.2.3. The change in K_{fn} that results from switching from $I_{mn} = 0$ to $I_{mn} = 1$ is constant and equal to $\exp(\delta_{mn} - \alpha_n c_{mn} - 1)$. The returns from switching K_{fn} by such a constant amount are smaller when K_{fn} is (weakly) larger due to the switch from $D_{ko} = 0$ to $D_{ko} = 1$. Therefore, the result follows from $\frac{\partial^2 \pi_{fn}}{\partial K_{fn}^2} < 0$.

A.2.5 D_{mo} and D_{mt} are substitutes for $o \neq t$

To prove this property, we rewrite,

$$\begin{aligned}
\Delta^f_{D_{mo}}(I_f, D_f, I_{-f}, D_{-f}) &= \sum_{n\in\mathcal{N}} [\pi_{fn}(I_{fn}, D_f^{mo\to 1}, I_{-fn}, D_{-f}) - \pi_{fn}(I_{fn}, D_f^{mo\to 0}, I_{-fn}, D_{-f})] - F^p_{mo} \\
&= \sum_{n\in\mathcal{N}} [\pi_{fn}(K_{fn}(I_{fn}, D_f^{mo\to 1}), K_{-fn}) - \pi_{fn}(K_{fn}(I_{fn}, D_f^{mo\to 0}), K_{-fn})] - F^p_{mo} \\
&= \sum_{n\in\mathcal{N}} \Delta^{mo}_n(I_{fn}, D_f, K_{-fn}) - F^p_{mo}, \qquad (76)
\end{aligned}$$

whre in equality (76),

$$\Delta^{mo}_n(I_{fn}, D_f, K_{-fn}) := \pi_{fn}(K_{fn}(I_{fn}, D_f^{mo\to 1}), K_{-fn}) - \pi_{fn}(K_{fn}(I_{fn}, D_f^{mo\to 0}), K_{-fn}).$$

Next, consider the difference-in-differences $\Delta^f_{D_{mo}}(I_f, D_f^{mt\to 1}, I_{-f}, D_{-f}) - \Delta^f_{D_{mo}}(I_f, D_f^{mt\to 0}, I_{-f}, D_{-f})$, which is the sum across all markets of terms $\Delta^{mo}_n(I_{fn}, D_f^{mt\to 1}, K_{-fn}) - \Delta^{mo}_n(I_{fn}, D_f^{mt\to 0}, K_{-fn})$. We prove that for each $n \in \mathcal{N}$, $\Delta^{mo}_n(I_{fn}, D_f^{mt\to 1}, K_{-fn}) - \Delta^{mo}_n(I_{fn}, D_f^{mt\to 0}, K_{-fn}) \leq 0$. The inequality clearly becomes an equality whenever $I_{mn} = 0$. Thus, consider the case in which $I_{mn} = 1$. There are multiple cases:

1. Neither o nor t are in the top 2 minimum-cost origins for model m in destination n. In this case, we have equality.
2. Origin o is ranked lowest-cost, t is not ranked second lowest-cost. In this case, we have equality.
3. Origin t is ranked lowest-cost, o is not ranked second-lowest cost. In this case, we have equality.
4. Origin o is ranked lowest-cost, origin t is second lowest-cost. In this case, the increase in K_{fn} when $D_{mt} = 0$ is larger than when $D_{mt} = 1$, so the difference-in-differences is strictly negative.
5. Origin t is ranked lowest-cost, and origin o is second lowest-cost. In this case, the first term in the difference-in-differences is zero, while the second term is positive, so the inequality is strict.

This concludes the proof.

A.3 Interdependence Properties Under V_m

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

$$V_m(D_m; I_{-m}, D_{-m}) := \max_{I_m} \sum_{n=1}^{N} I_{mn} \Delta\pi^m_{fn}(D_m; I_{-m}, D_{-m}) - \sum_{n=1}^{N} D_{mn} F^p_{mn}, \tag{77}$$

Throughout this section, we use subscript m to denote the decisions by firm $f(m)$ for product m, i.e., D_m is the vector of all production location decisions for product m. Similarly, we use subscript $-m$ to refer to *all* non-m decisions, within and across firms, i.e., I_{-m} refers to all sales decisions for products $k \neq m$ sold by firm $f(m)$ and all rival firms. Finally, we let $\Delta\pi^m_{fn}$ denote the marginal value of selling product m in destination market n.

As in Castro-Vincenzi et al. (2023), we first consider the single-product firm case. We then prove properties that hold in the general case with oligopolistic multi-product firms.

A.3.1 Substitution Across Production Locations for the Single-Product Firm

Consider the mapping,

$$V_m(D_m) := \max_{I_m} \sum_{n=1}^{N} I_{mn} \Delta\pi^m_{fn}(D_m) - \sum_{n=1}^{N} D_{mn} F^p_{mn}. \tag{78}$$

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

$$V_m((1,1), D_m) - V_m((0,1), D_m) \leq V_m((1,0), D_m) - V_m((0,0), D_m). \tag{79}$$

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

Introducing notation, let $I^{(i,j)}_{mn}$ 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 $((i,j), D_m)$ 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 D_m.

The first step of the proof is to realize that,

$$I^{(1,1)}_{mn} = I^{(1,0)}_{mn} + I^{(0,1)}_{mn} - I^{(1,0)}_{mn} I^{(0,1)}_{mn}. \tag{80}$$

Equation (80) 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,

$$I_{mn}^{(1,1)} \geq \max\{I_{mn}^{(1,0)}, I_{mn}^{(0,1)}\}, \tag{81}$$

and

$$I_{mn}^{(0,0)} \leq \min\{I_{mn}^{(1,0)}, I_{mn}^{(0,1)}\} \tag{82}$$

Inequality (81) says that if m is sold in n under either $(1,0)$ or $(0,1)$, then it is necessarily sold under $(1,1)$. Inequality (82) says that if m is sold in n under $(0,0)$, it must be sold either under $(1,0)$ or $(0,1)$. Combining inequalities (80)-(82), we obtain,

$$I_{mn}^{(1,1)} - I_{mn}^{(0,1)} = I_{mn}^{(1,0)} - I_{mn}^{(1,0)} I_{mn}^{(0,1)} \leq I_{mn}^{(1,0)} - I_{mn}^{(0,0)}. \tag{83}$$

We are now ready to prove the proposition. Expanding inequality (79), the goal is to prove that

$$\begin{gathered}\sum_n \left(I_{mn}^{(1,1)} - I_{mn}^{(0,1)}\right) \Delta\pi_{fn}^m(1,1) + \sum_n I_{mn}^{(0,1)} \left[\Delta\pi_{fn}^m(1,1) - \Delta\pi_{fn}^m(0,1)\right] \\ \leq \\ \sum_n \left(I_{mn}^{(1,0)} - I_{mn}^{(0,0)}\right) \Delta\pi_{fn}^m(1,0) + \sum_n I_{mn}^{(0,0)} \left[\Delta\pi_{fn}^m(1,0) - \Delta\pi_{fn}^m(0,0)\right]\end{gathered}$$

or equivalently,

$$\begin{aligned}\sum_n \left(I_{mn}^{(1,1)} - I_{mn}^{(0,1)}\right) \Delta\pi_{fn}^m(1,1) &+ \sum_n (I_{mn}^{(0,1)} - I_{mn}^{(0,0)}) \left[\Delta\pi_{fn}^m(1,1) - \Delta\pi_{fn}^m(0,1)\right] \\ &+ \sum_n I_{mn}^{(0,0)} [\Delta\pi_{fn}^m(1,1) - \Delta\pi_{fn}^m(0,1)] \\ &\leq \\ \sum_n \left(I_{mn}^{(1,0)} - I_{mn}^{(0,0)}\right) \Delta\pi_{fn}^m(1,0) &+ \sum_n I_{mn}^{(0,0)} \left[\Delta\pi_{fn}^m(1,0) - \Delta\pi_{fn}^m(0,0)\right].\end{aligned} \tag{84}$$

Due to substitution across production locations for a given product (subsection A.2.5), it must be that,

$$\Delta\pi_{fn}^m(1,1) - \Delta\pi_{fn}^m(0,1) \leq \Delta\pi_{fn}^m(1,0) - \Delta\pi_{fn}^m(0,0), \tag{85}$$

Combining inequalities (84)-(85), it follows that a sufficient condition to prove is,

$$\sum_n (I_{mn}^{(1,1)} - I_{mn}^{(0,0)}) \Delta\pi_{fn}^m(1,1) \leq \sum_n (I_{mn}^{(0,1)} - I_{mn}^{(0,0)}) \Delta\pi_{fn}^m(0,1) + \sum_n (I_{mn}^{(1,0)} - I_{mn}^{(0,0)}) \Delta\pi_{fn}^m(1,0). \tag{86}$$

Note that the LHS of inequality (86) sums across the variable profits in all countries for which the firm 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.4 V_m exhibits decreasing differences in D_m

A corollary of the result proven in subsection A.3.1 is that the function V_m, defined in equation (77) for general multi-product firms, exhibits decreasing differences in D_m. The proof follows immediately from the arguments in section A.3.1. Note that all of the arguments apply even in the case with multi-product firms competing oligopolistically, given that they do not depend on the values of (I_{-m}, D_{-m}).

A.4.1 V_m exhibits decreasing differences in D_{-m} and I_{-m}

To prove this result, consider the difference,

$$\Delta V^o_m(D_m; I_{-m}, D_{-m}) = V_m(D_m^{o\to 1}; I_{-m}, D_{-m}) - V_m(D_m^{o\to 0}; I_{-m}, D_{-m}) - F^p_{mo}.$$

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

$$\begin{aligned}
\Delta \underline{V}^o_m(D_m) - \Delta \overline{V}^o_m(D_m) = {} & \sum_n \underline{I}^1_{mn}[\Delta \underline{\pi}^m_{fn}(D_m^{o\to 1}) - \Delta \overline{\pi}^m_{fn}(D_m^{o\to 1})] \\
& + \sum_n (\underline{I}^1_{mn} - \overline{I}^1_{mn}) \Delta \overline{\pi}^m_{fn}(D_m^{o\to 1}) \\
& - \sum_n \underline{I}^0_{mn}[\Delta \underline{\pi}^m_{fn}(D_m^{o\to 0}) - \Delta \overline{\pi}^m_{fn}(D_m^{o\to 0})] \\
& - \sum_n (\underline{I}^0_{mn} - \overline{I}^0_{mn}) \Delta \overline{\pi}^m_{fn}(D_m^{o\to 0}),
\end{aligned} \tag{87}$$

where under and over-lines are used to denote dependence on $\underline{D}_{-m}(\underline{I}_{-m})$ and $\overline{D}_{-m}(\overline{I}_{-m})$, 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 (87) and write,

$$\Delta \underline{V}^o_m(D_m) - \Delta \overline{V}^o_m(D_m) = \Sigma_{G_m} + \Sigma_{E_m} + \Sigma_{L_m} \tag{88}$$

where Σ_{G_m} denotes the components of the sum in equation (87) 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,

$$\Sigma_{G_m} = \sum_{n \in G_m} \Delta \underline{\pi}^m_{fn}(D^{o \to 1}_m) \geq 0, \tag{89}$$

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 (i) $\underline{I}^1_{mn} = \overline{I}^1_{mn} = \underline{I}^0_{mn} = \overline{I}^0_{mn} = 0$, in which case the difference-in-differences is $\Sigma_{E_m} = 0$, (ii) $\underline{I}^1_{mn} = \overline{I}^1_{mn} = \underline{I}^0_{mn} = \overline{I}^0_{mn} = 1$, or (iii) $\underline{I}^1_{mn} = \underline{I}^0_{mn} = 1$ and $\overline{I}^1_{mn} = \overline{I}^0_{mn} = 0$. Thus,

$$\Sigma_{E_m} = \sum_{n \in E_m : \underline{I}^1_{mn} = \overline{I}^1_{mn} = \underline{I}^0_{mn} = \overline{I}^0_{mn} = 1} [(\Delta \underline{\pi}^m_{fn}(D^{o \to 1}_m) - \Delta \overline{\pi}^m_{fn}(D^{o \to 1}_m)) - (\Delta \underline{\pi}^m_{fn}(D^{o \to 0}_m) - \Delta \overline{\pi}^m_{fn}(D^{o \to 0}_m))] \tag{90}$$

$$+ \sum_{n \in E_m : \underline{I}^1_{mn} = \underline{I}^0_{mn} = 1 \text{ and } \overline{I}^1_{mn} = \overline{I}^0_{mn} = 0} [\Delta \overline{\pi}^m_{fn}(D^{o \to 1}_m) - \Delta \overline{\pi}^m_{fn}(D^{o \to 0}_m)] \geq 0 \tag{91}$$

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 well as the fact that producing product m is complementary with selling product m in country n. These are all properties established in Section A.2.

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,

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

where the inequality follows from the properties proven in Section A.2 and from the fact that for $n \in L_m$, $\overline{I}^0_{mn} = 0 \implies \Delta \overline{\pi}^m_{fn}(D^{o \to 0}_m) < 0$.

We have therefore proven that V_m exhibits decreasing differences in D_{-m} and I_{-m}.

B Method Proofs

B.1 Proof of Theorem 1

For $\boldsymbol{C} \subseteq \boldsymbol{C'}$, because the sup (inf) is increasing (decreasing) in the optimizing set, it must be that for all i, $\Delta_i^{f(i)}(\sup_i(\boldsymbol{C})) \leq \Delta_i^{f(i)}(\sup_i(\boldsymbol{C'}))$ and $\Delta_i^{f(i)}(\inf_i(\boldsymbol{C})) \geq \Delta_i^{f(i)}(\inf_i(\boldsymbol{C'}))$. Thus, $\overline{\Omega}(\boldsymbol{C'}) \subseteq \overline{\Omega}(\boldsymbol{C})$

and $\underline{\Omega}(\boldsymbol{C'}) \subseteq \underline{\Omega}(\boldsymbol{C})$, implying $F(\boldsymbol{C}) \subseteq F(\boldsymbol{C'})$. It follows that F is a set-increasing mapping.

B.2 Proof of Proposition 1

Start at $\boldsymbol{C}_0 = \mathcal{B}^{|J|}$ and iteratively apply F. By Theorem 1, since F is set-increasing, $\boldsymbol{C}_1 = F(\boldsymbol{C}_0) \subseteq \boldsymbol{C}_0$ implies that $F(\boldsymbol{C}_1) \subseteq F(\boldsymbol{C}_0)$. Therefore, letting $\boldsymbol{C}^{fixed}$ be the supremum of the set of fixed points of F, we have that $F(\boldsymbol{C}^{fixed}) = \boldsymbol{C}^{fixed} \subseteq F^K(\boldsymbol{C}_0)$ for any finite K. If there existed K such that, strictly, $\boldsymbol{C}^{fixed} \subset F^K(\boldsymbol{C}_0) = F^{K+1}(\boldsymbol{C}_0)$, then we would have found a larger fixed point, which is a contradiction. Therefore, there exists K such that $F^K(\boldsymbol{C}_0) = \boldsymbol{C}^{fixed}$, and $K \leq J$, since J is the largest number of coordinates that can be fixed. We have therefore proven items 1 and 2.

We now prove item 3. We first prove that $\lambda_{\mathcal{G}_F}(B) \subseteq F(B)$, for any $B \in S(\mathcal{B}^{|J|})$. We prove this by contraprosition. Suppose that $C \notin F(B)$. Thus, there exists some j such that either (i) $C_j = 0$ and $\Delta_j^{f(j)}(\inf_j(B)) \geq 0$, or (ii) $C_j = 1$ and $\Delta_j^{f(j)}(\sup_j(B)) < 0$. In either case, due to Assumption 1, it follows that C_j can never be a best response to *any* combination of opponent actions in B_{-j}, and therefore $C \notin \lambda_{\mathcal{G}_F}(B)$. To prove the other direction, suppose that $C \notin \lambda_{\mathcal{G}_F}(B)$. Then, there exists some j such that either (i) $C_j = 0$ and for all $s \in B$, $BR_j(s) = 1$, or (ii) $C_j = 1$ and for all $s \in B$, $BR_j(s) = 0$. Case (i) implies that $\Delta_j^{f(j)}(\inf_j(B)) \geq 0$. Case (ii) implies that $\Delta_j^{f(j)}(\sup_j(B)) < 0$. Therefore, $C \notin F(B)$. This concludes the proof.

Item 4 follows as a corollary of item 3. Indeed, by applying $F(\mathcal{B}^{|J|}) = \lambda_{\mathcal{G}_F}(\mathcal{B}^{|J|})$ iteratively until convergence, we obtain $IEDS(\mathcal{G}_F) = \boldsymbol{C}^{fixed}$.

To prove item 5, it we first show that $\lambda_{\mathcal{G}}(B) \subseteq \lambda_{\mathcal{G}_F}(B)$ for any $B \in S(\mathcal{B}^{|J|})$. Pick $C \in \lambda_{\mathcal{G}}(B)$. Then $C_f = BR_f(\tilde{C})$ for some $\tilde{C} \in B$. Let $\tilde{C}_{-i} = (C_{-i,f}, \tilde{C}_{-f}) \in B$ and notice that $C_i \in BR_i(\tilde{C})$. It follows that $C \in \lambda_{\mathcal{G}_F}(B)$ and, therefore, $\lambda_{\mathcal{G}}(B) \subseteq \lambda_{\mathcal{G}_F}(B)$ for any $B \in S(\mathcal{B}^{|J|})$. Then, starting at $\mathcal{B}^{|J|}$, and applying $\lambda_{\mathcal{G}}$, we obtain,

$$IEDS(\mathcal{G}) := \bigcap_{k=0}^{\infty} \lambda_{\mathcal{G}}^k(\mathcal{B}^{|J|}) \subseteq \bigcap_{k=0}^{\infty} \lambda_{\mathcal{G}}^k \circ \lambda_{\mathcal{G}_F}(\mathcal{B}^{|J|}) \subseteq \bigcap_{k=0}^{\infty} \lambda_{\mathcal{G}_F}^k(\mathcal{B}^{|J|}) =: IEDS(\mathcal{G}_F),$$

where the first inclusion follows from monotonicity of $\lambda_{\mathcal{G}}$ and the second inclusion follows from $\lambda_{\mathcal{G}} \subseteq \lambda_{\mathcal{G}_F}$. We have therefore proven item 5.

Finally, item 6 follows from the standard result (e.g., Bernheim 1984) that any PSNE is contained in $IEDS(\mathcal{G})$, and therefore $\boldsymbol{C}^{fixed} = IEDS(\mathcal{G}_F)$.

C Firm-by-Firm Branching Algorithm

This section adapts the branching algorithm from Castro-Vincenzi et al. (2023) to games.

Let $\boldsymbol{C} \in S(\mathcal{B}^{|J|})$ be a sublattice. Fix a firm f. Let $Z_f(\boldsymbol{C}) = \{i \in J_f : [\sup(\boldsymbol{C})]_i = 0\}$ be the set of firm-f coordinates that are solved to be 0 at $\boldsymbol{C}$ and $O_f(\boldsymbol{C}) = \{i \in J_f : [\inf(\boldsymbol{C})]_i = 1\}$ be the set of coordinates that are solved to be 1 at $\boldsymbol{C}$. Let $U_f(\boldsymbol{C}) = Z_f(\boldsymbol{C})^c \cap O_f(\boldsymbol{C})^c$. Since there are a finite number of elements in $U_f(\boldsymbol{C})^c$, we can define an ordering of coordinates in $U_f(\boldsymbol{C})^c$ as a function of

$\boldsymbol{C}$. Let $W_f : S(\mathcal{B}^{|J|}) \to J_f$ be a function mapping a sublattice $\boldsymbol{C}$ to a coordinate $i \in U_f(\boldsymbol{C})$. We are now ready to define the branching mapping.

Definition 4 (Branching Mapping) *The branching mapping is defined by function* $B_f : \mathcal{P}(S(\mathcal{B}^{|J|})) \to \mathcal{P}(S(\mathcal{B}^{|J|}))$ *as,*

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

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

We can now define the branching algorithm applied to firm f coordinates.

Definition 5 (Branching Algorithm for Firm f) *The branching algorithm applied to firm* f *is defined as follows:*

1. *Start at* $\boldsymbol{C}_0$ *with* $IEDS(\mathcal{G}) \subseteq \boldsymbol{C}_0$*. Obtain* $\boldsymbol{P}_1 = B_f(\{\boldsymbol{C}_0\})$*.*
2. *Compute* $\boldsymbol{P}_2 = B_f(\boldsymbol{P}_1)$
3. *Iterate until convergence to fixed point* $\boldsymbol{P}^*_f$*. By construction,* $\boldsymbol{P}^*_f$ *is a collection of sublattices of* $\mathcal{B}^{|J|}$*, which we denote by* $\bigcup_k \boldsymbol{C}^k_f$*.*

We now prove the key properties of the branching algorithm for firm f.

Proposition 2 *The branching algorithm for firm* f *has the following properties:*

1. B_f *is monotone-increasing, and applying* B_f *starting from* $\{\boldsymbol{C}_0\}$ *with* $IEDS(\mathcal{G}) \subseteq \boldsymbol{C}_0$ *converges to an upper bound of the set of fixed points of* B_f*.*
2. *Each non-empty sublattice* $\boldsymbol{C}^k_f$ *contains a unique action for firm* f*, denoted by* $C^k_f(\boldsymbol{P}^*_f)$*.*
3. *If* $(C_f, C_{-f}) \in IEDS(\mathcal{G})$*, then* $C_f = C^k_f(\boldsymbol{P}^*_f)$ *for some* k*.*

Proof. The proof of item 1 is identical ot the proof in Castro-Vincenzi et al. (2023), and follows from an application of Tarski (1955).

The proof of item 2 follows by definition of the branching algorithm. Indeed, if a sublattice contained more than a single action for firm f, then the next iteration of B_f would mechanically eliminate such a sublattice by further branching. Therefeore, any non-empty sublattice returned by branching must contain a single action for firm f.

If $(C_f, C_{-f}) \in IEDS(\mathcal{G})$, then necessarily the sublattice $\boldsymbol{C}_f$ with C_f fixed for firm f is a fixed point of F_f. Indeed, for any $i \in J_f$ it must hold that if $C_i = 1$, $\Delta_i^{f(i)}(\inf_i(\boldsymbol{C}_f)) \geq 0$, and if $C_i = 0$, then $\Delta_i^{f(i)}(\sup_i(\boldsymbol{C}_f)) < 0$. Otherwise, we would obtain a contradiction to the fact that C_f is rationalizable. Thus, it suffices to show that any fixed point $\boldsymbol{C}_f$ of F_f with a unique action for firm f, is contained $\boldsymbol{P}_f^*$. Suppose, by way of contradiction, that this were not the case. Since, $\boldsymbol{C}_f \subseteq \boldsymbol{C}^{fixed}$, this can only happen if at some branch $\overline{\boldsymbol{C}}_f$ containing $\boldsymbol{C}_f$, $\boldsymbol{C}_f \nsubseteq F_{f,*}(\overline{\boldsymbol{C}}_f)$. This is not possible, since $F_{f,*}(\boldsymbol{C}_f) = \boldsymbol{C}_f \subseteq F_{f,*}(\overline{\boldsymbol{C}}_f)$, since F_f is set-increasing. Therefore, $F_{f,*}$ never "skips" a fixed point. Moreover, since at there is at least one action for firm f that survives IEDS (by compactness of the action set and continuity of the payoff functions, see Bernheim 1984), it follows that $\bigcup_k \boldsymbol{C}_f^k$ is well-defined and each $\boldsymbol{C}_f^k$ is indeed non-empty. This concludes the proof.

■

D Bounding the Counterfactuals

This appendix demonstrates how we construct the outer bounds on the counterfactual outcomes reported in Section 6. For each fixed-cost shock draw s, the squeezing and branching algorithm returns a partial solution $C^{(s)} = (I^{(s)}, D^{(s)})$ in which a small set $\mathcal{U}^{(s)}$ of sales coordinates I_{mn} and production coordinates D_{mo} remain unresolved; let $\mathbf{C}^{(s)} \subseteq \{0,1\}^J$ denote the set of full $\{0,1\}$-vectors compatible with $C^{(s)}$.

D.1 Consumer Surplus

From the demand system, ex-post consumer surplus in destination n equals

$$\begin{aligned} CS_n(C) &= \frac{M_n}{\alpha_n} \ln\Big(1 + \sum_{f=1}^{F} \sum_{m \in \Omega_{fn}} \exp(u_{mn})\Big) \\ &= \frac{M_n}{\alpha_n} \ln\Big(1 + \sum_{f=1}^{F} \sum_{m \in \Omega_{fn}} \exp(\delta_{mn} + \xi_{mn} - \alpha_n c_{mn})\Big) \\ &= \frac{M_n}{\alpha_n} \ln\Big(1 + \sum_{f=1}^{F} W(K_{fn})\Big), \end{aligned}$$

with $\Omega_{fn} \equiv \{m \in \Omega_f \mid I_{mn} = 1\}$, $K_{fn} \equiv \sum_{m \in \Omega_{fn}} \exp(\delta_{mn} + \xi_{mn} - \alpha_n c_{mn} - 1)$, and $c_{mn} = \min_{o:\, D_{mo}=1}\{t_{on}\tau_{on}c_{mo}\}$. We drop the additive Euler constant from the Type-I extreme value normalization since it cancels in every comparison we report.

CS_n is increasing in all sales and production choices: a higher I_{mn} adds a model to the choice set, while a higher D_{mo} adds a plant to the minimum source set, weakly lowering the delivered cost c_{mn} and so weakly raising CS_n. Therefore, we construct the bounds on CS_n as:

$$\overline{CS}_n^{(s)} \equiv \sup_{C \in \mathbf{C}^{(s)}} CS_n(C) \qquad \text{and} \qquad \underline{CS}_n^{(s)} \equiv \inf_{C \in \mathbf{C}^{(s)}} CS_n(C),$$

which are attained at the two corners of $\mathcal{U}^{(s)}$ by setting every unresolved coordinate to 1 (for $\overline{CS}$) or 0 (for $\underline{CS}$). We report $S^{-1}\sum_s \underline{CS}_n^{(s)}$ and $S^{-1}\sum_s \overline{CS}_n^{(s)}$ as the bounds on consumer surplus.

D.2 Total Variable Profits

Firm f's variable profit in destination n is

$$\pi_{fn}(C) = \frac{M_n}{\alpha_n}\,\frac{W(K_{fn})\,(1+W(K_{fn}))}{1+\sum_{f'=1}^{F} W(K_{f'n})},$$

which is increasing in K_{fn} and decreasing in $K_{f'n}$ for $f' \neq f$. Equivalently, π_{fn} is increasing in firm f's own coordinates $(I_{f\cdot}, D_{f\cdot})$ and decreasing in its rivals'.

We report three aggregations of variable profit. The first sums the variable profits of U.S.-headquartered firms across all destination markets:

$$\Pi^{\text{U.S.-HQ}}(C) = \sum_{f:\,\text{HQ}_f=\text{US}} \sum_n \pi_{fn}(C).$$

For each U.S.-HQ firm f, we define the asymmetric corner states

$$\begin{aligned}\overline{C}_f^{(s)} &\equiv \big(\overline{I}_{f\cdot}^{(s)},\, \overline{D}_{f\cdot}^{(s)},\, \underline{I}_{-f\cdot}^{(s)},\, \underline{D}_{-f\cdot}^{(s)}\big),\\ \underline{C}_f^{(s)} &\equiv \big(\underline{I}_{f\cdot}^{(s)},\, \underline{D}_{f\cdot}^{(s)},\, \overline{I}_{-f\cdot}^{(s)},\, \overline{D}_{-f\cdot}^{(s)}\big),\end{aligned}$$

where $\overline{C}_f^{(s)}$ raises f's unresolved entries and lowers its rivals', and $\underline{C}_f^{(s)}$ flips the asymmetry. The outer bounds on $\Pi^{\text{U.S.-HQ}}$ at draw s are

$$\begin{aligned}\overline{\Pi}^{\text{U.S.-HQ},(s)} &= \sum_{f:\,\text{HQ}_f=\text{US}} \sum_n \pi_{fn}\big(\overline{C}_f^{(s)}\big),\\ \underline{\Pi}^{\text{U.S.-HQ},(s)} &= \sum_{f:\,\text{HQ}_f=\text{US}} \sum_n \pi_{fn}\big(\underline{C}_f^{(s)}\big).\end{aligned}$$

The other two aggregations are based on U.S.-headquartered brands and U.S. plants:

$$\begin{aligned}\Pi^{\text{U.S.-Brand}}(C) &= \sum_{(f,m):\,\text{HQ}_{b(m)}=\text{US}} \sum_n \pi_{fmn}(C),\\ \Pi^{\text{U.S.-Plants}}(C) &= \sum_{f,m,n:\,o^*(f,m,n;C)=\text{US}} \pi_{fmn}(C),\end{aligned}$$

Since each model's contribution flows through a different firm, we need to work with model-level variable profit, i.e. $\pi_{fmn}(C) = \mu_{fn}(C)\,q_{fmn}(C)$. We bound this object by separately bounding the markup μ_{fn} and the sales quantity q_{fmn}, then taking the product as an outer bound: $\sup_{\mathbf{C}^{(s)}} \pi_{fmn} \leq (\sup_{\mathbf{C}^{(s)}} \mu_{fn})(\sup_{\mathbf{C}^{(s)}} q_{fmn})$.

First, the per-unit markup $\mu_{fn} \equiv (1+W(K_{fn}))/\alpha_n$ is constant across models within firm f's market and depends only on f's own coordinates. It is increasing in f's own coordinates and does not depend on rivals' coordinates. As a result, its bounds are attained at the same corner states $\overline{C}_f^{(s)}$ and $\underline{C}_f^{(s)}$ used for $\Pi^{\text{U.S.-HQ}}$:

$$\overline{\mu}_{fn,k}^{(s)} = \mu_{fn,k}\big(\overline{C}_f^{(s)}\big), \qquad \underline{\mu}_{fn,k}^{(s)} = \mu_{fn,k}\big(\underline{C}_f^{(s)}\big).$$

Next, model m's sales quantity in market n is

$$q_{fmn} \;=\; \frac{M_n \exp\!\big(\delta_{mn} + \xi_{mn} - \alpha_n c_{mn} - 1 - W(K_{fn})\big)}{1 + W(K_{fn}) + \sum_{f' \neq f} W(K_{f'n})}.$$

is increasing in m's own K-contribution but decreasing in every other K-contribution (same-firm cannibalization and rival presence). Bounds therefore come from a model-level asymmetric resolution by setting m's own unsolved coordinates to the upper bound and all other models' to the lower bound. Specifically, for each model m, define the model-level corner states

$$\begin{aligned}\overline{C}_m^{(s)} &\equiv \big(\overline{I}_{m,\cdot}^{(s)},\ \overline{D}_{m,\cdot}^{(s)},\ \underline{I}_{-m,\cdot}^{(s)},\ \underline{D}_{-m,\cdot}^{(s)}\big),\\ \underline{C}_m^{(s)} &\equiv \big(\underline{I}_{m,\cdot}^{(s)},\ \underline{D}_{m,\cdot}^{(s)},\ \overline{I}_{-m,\cdot}^{(s)},\ \overline{D}_{-m,\cdot}^{(s)}\big),\end{aligned}$$

where the $-m$ subscript collects every other model, spanning both firm f's remaining models and rivals' models. Quantity bounds at draw s are

$$\begin{aligned}\overline{q}_{fmn}^{(s)} &= q_{fmn}\big(\overline{C}_m^{(s)}\big),\\ \underline{q}_{fmn}^{(s)} &= q_{fmn}\big(\underline{C}_m^{(s)}\big).\end{aligned}$$

Finally, combining the firm-level markup bound with the model-level quantity bound gives an outer bound on the model-level variable profit:

$$\begin{aligned}\overline{\pi}_{fmn}^{(s)} &\equiv \overline{\mu}_{fn}^{(s)}\,\overline{q}_{fmn}^{(s)} = \mu_{fn}\big(\overline{C}_f^{(s)}\big)\,q_{fmn}\big(\overline{C}_m^{(s)}\big),\\ \underline{\pi}_{fmn}^{(s)} &\equiv \underline{\mu}_{fn}^{(s)}\,\underline{q}_{fmn}^{(s)} = \mu_{fn}\big(\underline{C}_f^{(s)}\big)\,q_{fmn}\big(\underline{C}_m^{(s)}\big).\end{aligned}$$

Summing over U.S.-branded (f, m) pairs:

$$\begin{aligned}\overline{\Pi}^{\text{U.S.-Brand},(s)} &= \sum_{(f,m):\,\text{HQ}_{b(m)}=\text{US}} \sum_n \overline{\pi}_{fmn}^{(s)},\\ \underline{\Pi}^{\text{U.S.-Brand},(s)} &= \sum_{(f,m):\,\text{HQ}_{b(m)}=\text{US}} \sum_n \underline{\pi}_{fmn}^{(s)}.\end{aligned}$$

Analogous constructions yield the bounds on $\Pi^{\text{U.S.-Plants}}$, with one more complication described in the next subsection.

D.3 Bounding Sourcing-Restricted Aggregates

The construction so far bounds $\Pi^{\text{U.S.-HQ}}$ and $\Pi^{\text{U.S.-Brand}}$, whose summation sets are pinned down by firm and brand identifiers invariant to the solution. However, $\Pi^{\text{U.S.-Plants}}$, together with the any sourcing-based outcomes, introduces a new complication: which (f, m, n) should be included depends on the cost-minimizing source for (m, n), and that source can flip with the unresolved D coordinates. The summation set is therefore uncertain.

We handle this by classifying each (f, m, n) according to the sourcing status of a candidate plant location ℓ (here $\ell = \text{US}$). Writing $c_{m,l,n} \equiv t_{ln}\tau_{ln}c_{ml}$ for the landed cost in market n delivered from plant l:

$$\ell \textit{ always cheapest } \equiv D^{(s)}_{m\ell} = 1 \text{ and } c_{m,l,n} > c_{m,\ell,n} \text{ for every } l \neq \ell \text{ with } D^{(s)}_{ml} \geq 0.5,$$
$$\ell \textit{ possibly cheapest } \equiv D^{(s)}_{m\ell} \geq 0.5 \text{ and } c_{m,l,n} \geq c_{m,\ell,n} \text{ for every } l \neq \ell \text{ with } D^{(s)}_{ml} = 1.$$

"Always cheapest" says ℓ is active and strictly undercuts every potentially active plant, so ℓ is the source for every $C \in \mathbf{C}^{(s)}$. "Possibly cheapest" says ℓ might be active and is no worse than any certainly-active plant, so ℓ is the source for at least one $C \in \mathbf{C}^{(s)}$. The former implies the latter.

For any non-negative outcome Θ_{fmn} with per-model bounds $\overline{\Theta}^{(s)}_{fmn} \geq \underline{\Theta}^{(s)}_{fmn} \geq 0$ that vanishes whenever (f, m, n) is not sourced from ℓ, replacing the exact sourcing condition $o^*(f, m, n; C) = \ell$ in the summation set with the appropriate sourcing class yields ordered aggregate bounds:

$$\overline{\Theta}^{\ell,(s)} = \sum_{(f,m,n):\, \ell \text{ possibly cheapest}} \overline{\Theta}^{(s)}_{fmn},$$
$$\underline{\Theta}^{\ell,(s)} = \sum_{(f,m,n):\, \ell \text{ always cheapest}} \underline{\Theta}^{(s)}_{fmn}.$$

Set ordering and non-negativity give $\overline{\Theta}^{\ell,(s)} \geq \underline{\Theta}^{\ell,(s)}$. We apply this construction with $\Theta^{(s)}_{fmn} = \pi^{(s)}_{fmn}$ to bound $\Pi^{\text{U.S.-Plants}}$, and analogously below for total variable costs.

D.4 Total Variable Cost

For a country ℓ, we report the total variable cost incurred at ℓ's plants, summed across every destination market in which ℓ-plant output is sold. The relevant sum runs over all (f, m, n) triples whose active plant is ℓ:

$$\text{TVC}^{\ell}(C) = \sum_{n} \sum_{f,m:\, I_{mn}=1 \,\wedge\, o^*(f,m,n;C)=\ell} \tau_{\ell n}\, c_{m\ell}\, q_{fmn}(C),$$

where $\tau_{\ell n}\, c_{m\ell}$ is the per-unit producer cost (CIF) of supplying market n from ℓ's plant, net of tariffs. For each (f, m, n), the producer cost is determinant, so the cost contribution is bounded above by $\tau_{\ell n}\, c_{m\ell}\, \overline{q}^{(s)}_{fmn}$ and below by $\tau_{\ell n}\, c_{m\ell}\, \underline{q}^{(s)}_{fmn}$. Applying the same quantity bounds above:

$$\overline{\mathrm{TVC}}^{\ell,(s)} = \sum_{n} \sum_{(f,m):\,\ell \text{ possibly cheapest at } n} \tau_{\ell n}\, c_{m\ell}\, \overline{q}^{(s)}_{fmn},$$
$$\underline{\mathrm{TVC}}^{\ell,(s)} = \sum_{n} \sum_{(f,m):\,\ell \text{ always cheapest at } n} \tau_{\ell n}\, c_{m\ell}\, \underline{q}^{(s)}_{fmn}.$$

The sourcing classification is computed per (f, m, n), not per m, because the cost ordering between ℓ and competitor plants can flip across destinations: the U.S. plant might be the cheapest source for the U.S. market but lose to Mexico for the E.U. market. We apply this construction with $\ell = \mathrm{US}$ for the U.S.-households table and $\ell = \mathrm{CAN}$ for the Canada-households table. Total labor cost at ℓ's plants is $\sigma_L \cdot \mathrm{TVC}^{\ell}$ with $\sigma_L = 0.20$, reported per household of country ℓ. Reported intervals are S^{-1}-averages of the per-draw endpoints.

D.5 Sourcing Share

The fraction of n-sold models actually sourced from country ℓ is

$$\mathrm{share}^{\ell}_{n}(C) = \frac{\#\{(f,m) : I_{mn} = 1 \wedge o^{*}(f,m,n;C) = \ell\}}{\#\{(f,m) : I_{mn} = 1\}}.$$

The numerator and denominator share the same C, so bounding them separately gives a strictly loose envelope. We instead treat each (f, m) as a joint choice in $\{(0,0), (0,1), (1,1)\}$, where the pair is (numerator-contribution, denominator-contribution): $(0,0)$ if the model is not sold, $(0,1)$ if sold but sourced elsewhere, and $(1,1)$ if sold and ℓ-sourced.

Following this construction, each (f, m) with $I^{(s)}_{mn} \geq 0.5$ falls into one of six cells:

	Sale state ($I^{(s)}_{mn}$)	Sourcing state at ℓ
ν_A	1	always cheapest
ν_B	1	possibly (not always) cheapest
ν_D	1	never cheapest
ν_E	0.5	always-if-sold
ν_F	0.5	possibly-if-sold (not always)
ν_G	0.5	never-if-sold

Maximizing the share within $\mathbf{C}^{(s)}$ admits every candidate that can be made $(1, 1)$ and excludes every candidate that can only be made $(0, 1)$; minimizing reverses both choices. The bounds are

$$\overline{\mathrm{share}}^{\ell,(s)}_{n} = \frac{\nu_A + \nu_B + \nu_E + \nu_F}{\nu_A + \nu_B + \nu_D + \nu_E + \nu_F},$$
$$\underline{\mathrm{share}}^{\ell,(s)}_{n} = \frac{\nu_A}{\nu_A + \nu_B + \nu_D + \nu_F + \nu_G}.$$

We apply $(\ell, n) = (\mathrm{US}, \mathrm{US})$ for the U.S.-households table and $(\mathrm{CAN}, \mathrm{CAN})$ for the Canada-households table. Reported intervals are S^{-1}-averages of the per-draw endpoints.

D.6 Tariff Revenue

Destination n's government collects an ad-valorem tariff at rate $t_{ln} - 1$ on each unit imported from origin l, applied to the per-unit producer cost $\tau_{ln} c_{ml}$. Total tariff revenue at configuration C is

$$TR_n(C) \;=\; \sum_f \sum_{m \in \Omega_{fn}} (t_{o^*(C),n} - 1)\, \tau_{o^*(C),n}\, c_{m,o^*(C)}\, q_{fmn}(C),$$

where $o^*(f, m, n; C) \equiv \arg\min_{o:\, D_{mo}=1} c_{m,o,n}$ is the cheapest active plant supplying (m, n), $c_{m,l,n} \equiv t_{ln} \tau_{ln} c_{ml}$ is its landed cost, and $\Omega_{fn} \equiv \{m \in \Omega_f : I_{mn} = 1\}$.

Substituting the model-level quantity $q_{fmn}(C) = M_n\, s_{f,n}(C)\, \kappa_m(C)/K_{fn}(C)$ and pulling firm-level factors out of the inner sum,

$$TR_n(C) \;=\; \sum_f \frac{Q_{fn}(C)}{K_{fn}(C)}\, H_{fn}(C),$$

where $Q_{fn}(C) \equiv M_n\, s_{f,n}(C)$ is firm f's total quantity in n, $\kappa_m(C) \equiv \exp(\delta_{mn} + \xi_{mn} - \alpha_n c_{m,o^*(C),n} - 1)$, $K_{fn}(C) \equiv \sum_{m \in \Omega_{fn}} \kappa_m(C)$, and

$$H_{fn}(C) \;\equiv\; \sum_{m \in \Omega_{fn}} (t_{o^*(C),n} - 1)\, \tau_{o^*(C),n}\, c_{m,o^*(C)}\, \kappa_m(C).$$

We bound Q_{fn}/K_{fn} and H_{fn} separately.

Bound on Q_{fn}/K_{fn}. Let $u_f(C) \equiv W(K_{fn}(C))$. The Lambert-$W$ identity $W(K)/K = \exp(-W(K))$ gives the closed form

$$\frac{Q_{fn}(C)}{K_{fn}(C)} \;=\; \frac{M_n\, e^{-u_f(C)}}{1 + \sum_g u_g(C)},$$

which is decreasing in every firm's $u_g = W(K_{gn})$, including own. The supremum over $\mathbf{C}^{(s)}$ is therefore attained at the global state $\underline{C}^{(s)}$ that minimizes every K_{gn} jointly – every unresolved (I, D) coordinate set to 0 across all firms – and the infimum at the all-strong corner $\overline{C}^{(s)}$. Letting $\underline{K}^{(s)}_{gn} \equiv K_{gn}(\underline{C}^{(s)})$ and $\overline{K}^{(s)}_{gn} \equiv K_{gn}(\overline{C}^{(s)})$,

$$\overline{(Q/K)}^{(s)}_{fn} \;\equiv\; \frac{M_n\, \exp\!\big(-W(\underline{K}^{(s)}_{fn})\big)}{1 + \sum_g W(\underline{K}^{(s)}_{gn})}, \qquad \underline{(Q/K)}^{(s)}_{fn} \;\equiv\; \frac{M_n\, \exp\!\big(-W(\overline{K}^{(s)}_{fn})\big)}{1 + \sum_g W(\overline{K}^{(s)}_{gn})}.$$

Bound on H_{fn}. The summand $h_{fmn}(l) \equiv (t_{ln} - 1)\, \tau_{ln}\, c_{ml}\, \exp(\delta_{mn} + \xi_{mn} - \alpha_n c_{m,l,n} - 1)$ is not monotone in landed cost: the leading factor $(t_{ln} - 1)\, \tau_{ln}\, c_{ml}$ varies across plants independently of

$c_{m,l,n}$, so the plant maximizing $h_{fmn}(l)$ need not be the cheapest one. We therefore bound h_{fmn} per model over the *potential sourcing set*

$$\mathcal{L}^{\text{poss}}_{fmn} \equiv \Big\{ l : D^{(s)}_{ml} \geq 0.5,\ c_{m,l,n} \leq \min_{l' :\, D^{(s)}_{ml'} = 1} c_{m,l',n} \Big\},$$

i.e., every potentially-active plant ($D^{(s)}_{ml} \geq 0.5$) not strictly dominated in landed cost by some confirmed-active plant. Plants outside $\mathcal{L}^{\text{poss}}_{fmn}$ can never be the active source $o^*(C)$ for any $C \in \mathbf{C}^{(s)}$. Define $h^{(s)}_{\max,fmn} \equiv \max_{l \in \mathcal{L}^{\text{poss}}_{fmn}} h_{fmn}(l)$ and $h^{(s)}_{\min,fmn}$ analogously. Aggregating across models,

$$\overline{H}^{(s)}_{fn} = \sum_{m :\, I^{(s)}_{mn} \geq 0.5} h^{(s)}_{\max,fmn}, \qquad \underline{H}^{(s)}_{fn} = \sum_{m :\, I^{(s)}_{mn} = 1} h^{(s)}_{\min,fmn}.$$

Tariff-revenue bounds. Combining the two,

$$\overline{TR}^{(s)}_{n} = \sum_f \overline{(Q/K)}^{(s)}_{fn}\, \overline{H}^{(s)}_{fn}, \qquad \underline{TR}^{(s)}_{n} = \sum_f \underline{(Q/K)}^{(s)}_{fn}\, \underline{H}^{(s)}_{fn}.$$

We compute these per shock draw and report the S^{-1}-average of the per-draw endpoints. The reported TR_n is for $n =$ US in the U.S.-households table (Panel E) and $n =$ CAN in the Canada-households table.

D.7 Average Price, Markup, and Marginal Cost

Table 3 reports per-model averages within three panels: Panel A (all U.S.-sold), Panel B (U.S.-sold and U.S.-sourced), and Panel C (U.S.-sold and foreign-sourced). For each panel $\mathcal{P}$, restrict the potential sourcing set $\mathcal{L}^{\text{poss}}_{fmn}$ from Section D.6 to panel-admissible plants:

$$\mathcal{L}^{\text{poss},\mathcal{P}}_{fm} \equiv \big\{ l \in \mathcal{L}^{\text{poss}}_{fmn} : l \text{ admissible under } \mathcal{P} \big\},$$

where Panel A admits any plant, Panel B fixes $l =$ US, and Panel C requires $l \neq$ US. The bounds on the landed cost ($\underline{c}^{\mathcal{P}}_{mn}$, $\overline{c}^{\mathcal{P}}_{mn}$) and on the producer cost are taken as the minimum and maximum across plants in the panel-admissible set $\mathcal{L}^{\text{poss},\mathcal{P}}_{fm}$. By construction, the upper bound for landed cost coincides with the cheapest-confirmed-plant value used in Section D.1. For producer cost the coincidence breaks when tariff rates vary across origins, since the producer-cost ordering across plants need not match the landed-cost ordering. Markup bounds reuse the firm-level corners of Section D.2: $\overline{\mu}^{(s)}_{fn} = \mu_{fn}\big(\overline{C}^{(s)}_f\big)$ and $\underline{\mu}^{(s)}_{fn} = \mu_{fn}\big(\underline{C}^{(s)}_f\big)$. Consumer-price bounds compose: $\overline{p}^{\mathcal{P},(s)}_{fmn} = \overline{c}^{\mathcal{P}}_{mn} + \overline{\mu}^{(s)}_{fn}$ and $\underline{p}^{\mathcal{P},(s)}_{fmn} = \underline{c}^{\mathcal{P}}_{mn} + \underline{\mu}^{(s)}_{fn}$.

The set of (f, m) pairs included in each panel is determined by the sourcing classification of Section D.3:

Panel $\mathcal{P}$	$\underline{\mathcal{P}}$ (always in)	$\overline{\mathcal{P}}$ (possibly in)
A: all U.S.-sold	$I^{(s)}_{m,\mathrm{US}} = 1$	$I^{(s)}_{m,\mathrm{US}} \geq 0.5$
B: U.S.-sourced	$I^{(s)}_{m,\mathrm{US}} = 1$, U.S. always cheapest	$I^{(s)}_{m,\mathrm{US}} \geq 0.5$, U.S. possibly cheapest
C: foreign-sourced	$I^{(s)}_{m,\mathrm{US}} = 1$, U.S. never cheapest	$I^{(s)}_{m,\mathrm{US}} \geq 0.5$, U.S. not always cheapest

$\overline{\mathcal{P}}$ is the larger (possibly-in) set; $\underline{\mathcal{P}}$ the always-in subset. For $x \in \{p, \mu, c\}$, the panel average $\overline{x}^{\mathcal{P}} \equiv |\mathcal{P}|^{-1} \sum_{(f,m)\in\mathcal{P}} x_{fmn}$ is bounded by pushing numerator and denominator to opposite extremes:

$$\overline{\overline{x}}^{\mathcal{P},(s)} = \frac{\sum_{(f,m)\in\overline{\mathcal{P}}} \overline{x}^{\mathcal{P},(s)}_{fmn}}{|\underline{\mathcal{P}}|}, \qquad \underline{x}^{\mathcal{P},(s)} = \frac{\sum_{(f,m)\in\underline{\mathcal{P}}} \underline{x}^{\mathcal{P},(s)}_{fmn}}{|\overline{\mathcal{P}}|}.$$

Product-count columns report $|\underline{\mathcal{P}}|$ and $|\overline{\mathcal{P}}|$ directly.