

Combinatorial Discrete Choice: Theory and Application to Multinational Production

Costas Arkolakis, Fabian Eckert, and Rowan Shi*

March 29, 2026

In general, finding the optimal combination among N items requires evaluating all 2^N possibilities. We develop a method applicable whenever items are complements or substitutes that uses this structure to discard suboptimal combinations without evaluation and extends to heterogeneous-agent environments. In a calibrated general equilibrium model of multinational firms choosing interdependent production locations, we show the method is orders of magnitude faster than incumbent approaches and avoids error from discretizing firm heterogeneity. Quantitatively, whether production locations are complements or substitutes determines whether welfare gains from multinational production concentrate in the most productive countries or accrue more broadly.

*Arkolakis: Yale University (costas.arkolakis@yale.edu); Eckert: University of California, San Diego (fpe@ucsd.edu); Shi: Toronto Metropolitan University (rowanxshi@torontomu.ca). We are indebted to Junyuan Chen for help with the Julia code. We are grateful to Pablo Fajgelbaum and Eduardo Morales for excellent discussions at the Princeton IES Summer Trade Workshop. We thank Stanislav Atanasov, Jake DiCicco, Tra Nguyen, Andrew Salmon, Benjamin Tong, Yunus Tuncbilek, and Nardeen Abdulkareem for outstanding research assistance. We also thank Chrysanthos Gounaris, David Atkin, Arnaud Costinot, Dave Donaldson, Penny Goldberg, Navin Kartik, Samuel Kortum, Nitya Pandalai-Nayar, and Steve Redding for comments and suggestions. Shi gratefully acknowledges support from the Princeton IES. Code implementing the techniques introduced in this paper can be found at <https://github.com/rowanxshi/CDCP.jl>.

A combinatorial discrete choice problem (CDCP) is the problem of finding the combination of items that maximizes an objective function when the value of each item depends on which others are also chosen. The decision problem of a multinational firm selecting global production locations is an example: setting up production sites in Germany and the United States may generate scale economies that raise the value of both locations, while setting up such sites in Germany and France may lead to mutual cannibalization in the European market. Without imposing more structure on the objective function, the decision maker must evaluate all possible combinations, the number of which grows exponentially: with just twenty potential locations, there are already over one million combinations.

This paper develops a method to solve CDCPs whose objective function satisfies a single crossing condition—weaker than requiring items to be complements or substitutes—that discards suboptimal combinations without computing their payoffs. The method extends to heterogeneous-agent settings, always recovers the global optimum and, in our quantitative application, is orders of magnitude faster than evaluating every combination. Using it, we calibrate a general equilibrium model of multinational production in which firms choose foreign production locations. When locations are strongly complementary, gains from multinational production mainly accrue to the most productive firms and their headquarter countries, while some less productive countries even experience welfare losses; when they are substitutes, gains are positive and more even across countries.

Our method builds on [Jia \[2008\]](#), which shows that complementarity among choices enables an approach that discards suboptimal combinations without computing their payoffs. The key insight is that complementarity induces monotonicity: an item’s value is highest when all others are chosen and lowest when it stands alone. If an item is valuable even in the least favorable environment—when chosen in isolation—it must always be included, and all combinations excluding it can be discarded; if it is not valuable even in the most favorable environment—when all other items are present—it must always be excluded, and all combinations including it can be discarded.

We show that this elimination logic extends to any CDCP satisfying a *single crossing differences* condition, either from below or from above. Single crossing differences from below generalizes complementarity: if an item’s value is non-negative, it remains so as more items are added, though it may increase or decrease. Single crossing differences from above captures the symmetric case: if an item’s value is non-negative, it remains so as items are removed, generalizing substitutability. For this latter case, the incumbent method is brute-force evaluation of all combinations, since the method of [Jia \[2008\]](#) does not apply.

Iteratively applying the elimination logic reduces a CDCP’s domain by discarding combinations of items that cannot be optimal, but the reduced domain may

still contain multiple candidate combinations. In such cases, Jia [2008] applies the brute-force approach to the reduced domain, evaluating all remaining combinations. We introduce a *branching* procedure that applies the elimination logic as far as possible before resorting to a brute-force step: when elimination can no longer make progress, it fixes one item’s status, reapplies the elimination logic to each resulting case, and compares across cases to find the global optimum.

Heterogeneous agents compound the computational challenge: since each agent type may have a different optimal combination, solving the problem requires tackling a separate CDCP for every type in the population. For such settings, we extend our method to recover the exact *policy function* mapping each type to its optimal combination. This extension applies to CDCPs satisfying an additional single crossing condition: if including an item is worthwhile for a given agent type, it remains so for all higher types. When this condition holds, we can solve for the cutoffs at which the policy function changes value directly by using an indifference condition, avoiding the need to solve CDCPs for every agent type individually.

We apply our method to a general equilibrium model of multinational production in which each firm solves a CDCP when selecting production locations. Fixed costs of operating in each location prevent firms from producing everywhere, making the location choice combinatorial. Two counteracting channels determine whether locations are net complements or substitutes: an additional location increases firm-level scale economies but simultaneously cannibalizes demand across the firm’s existing locations. We show that our single crossing differences conditions hold in the model, with locations being net complements when the scale parameter exceeds the cannibalization parameter, and net substitutes otherwise.

Our method makes it possible to calibrate the model to data on 32 OECD and European countries when locations are complements or substitutes: without it, each firm’s location choice problem would require evaluating billions of combinations. The calibration targets bilateral trade flows, foreign affiliate sales, and affiliate counts across all country pairs, together with country-level moments on production, firm entry, and firm survival. Because production location sets vary across firms, the model does not yield aggregate gravity equations, so we use indirect inference, embedding our method inside the calibration loop to solve for policy functions at every parameter guess.

In benchmarking exercises using the calibrated model, our method outperforms existing approaches along three dimensions: computational speed, precision, and robustness across the range of complementarity and substitutability. We benchmark our method for up to 256 countries, sampling from our calibrated model to generate synthetic countries. Our method is up to four orders of magnitude faster than the brute-force approach and an order of magnitude faster than Jia [2008] with discretized firm heterogeneity when locations are complements, and retains

its speed when locations are substitutes. On precision, we show that the incumbent approach of solving CDCPs for a discrete number of types and interpolating between them can induce substantial error in aggregate outcomes, while our approach is exact. Finally, our method is equally fast in both the complements and substitutes calibrations.

Whether locations are complements or substitutes determines whether welfare gains from multinational production concentrate in the most productive countries or accrue more broadly, a finding we establish by calibrating the model across more than 30 degrees of complementarity and substitutability. The stronger the complementarity, the more the distribution of welfare gains becomes uneven: countries home to large multinational firms, such as the United States, experience substantial gains, while some less productive countries that primarily host foreign production sites even lose. With substitutability, gains are positive and more even across countries since cross-location cannibalization limits the scale economies attainable by the most productive firms.

Modeling discrete choices when the value of an item depends on which others are chosen has long posed a challenge, compounded in settings with rich agent heterogeneity. The classic random utility framework of [McFadden \[1974\]](#) sidesteps it by restricting agents to choosing a single item, which rules out interactions among items but permits flexible heterogeneity through item-specific shocks and tractable aggregation. Extending the framework to allow agents to choose *multiple* items is tractable when items are independent [see, e.g., [Hendel, 1999](#)], but when they are interdependent, existing approaches restrict attention to settings with few items [see, e.g., [Train et al., 1987](#), [Gentzkow, 2007](#)] due to the exponential growth in the number of combinations.

Our paper contributes to a growing literature on discrete choice problems with interdependent items outside the random utility framework. [Jia \[2008\]](#) introduces a solution method when choices are complements, which has been applied to store expansion, firm sourcing, and multinational production settings [see [Antràs et al., 2017](#), [Alfaro-Ureña et al., 2024](#), [Antràs et al., 2024a](#)].¹ We extend this line of work in two main ways. First, we provide a method that applies whenever items are substitutes rather than complements, and more generally whenever the objective function satisfies a single crossing condition. Second, we develop an exact heterogeneous-agent extension that avoids the approximation error induced by discretizing heterogeneity [see [Tintelnot, 2017](#), [Antràs et al., 2024a](#)].

¹A separate strand of this literature pursues solution methods that do not rely on monotonicity of choices. Computation-oriented approaches to CDCPs have used linear objective functions to apply integer programming [e.g., [Head et al., 2026](#)], greedy algorithms that do not guarantee global optimality [e.g., [Fan and Yang, 2020](#)], or machine learning methods to approximate heterogeneous-agent policy functions [e.g., [Kulesza, 2024](#)].

Relative to existing papers in the literature on multinational production, our tools make incorporating cross-location interactions and fixed costs possible, which we use to study how the welfare implications of multinational production depend on the presence of these forces. Existing quantitative work on multinationals has avoided CDCPs [see Ramondo and Rodríguez-Clare, 2013, Ramondo, 2014, Arkolakis et al., 2018], has solved small-scale CDCPs using the brute-force approach [see Tintelnot, 2017, Dyrda et al., 2024], or has restricted attention to cases with complementarities so that the method of Jia [2008] applies [see Antràs et al., 2017, Alfaro-Ureña et al., 2024, Antràs et al., 2024a]. Our counterfactual experiments suggest that interdependencies among production locations, such as scale economies and cannibalization effects, are central to understanding who gains and who loses from multinational production, with implications for the broader debate on the distributional consequences of globalization.

1. Defining and Solving Combinatorial Discrete Choice Problems

In this section, we formally define CDCPs and show how to solve them when the objective function satisfies two single crossing conditions. Let L denote a finite set of items indexed by ℓ , and $f : \mathcal{P}(L) \rightarrow \mathbb{R}$ the objective function, mapping the set of all possible combinations of items, $\mathcal{P}(L)$, to a real number. The Online Appendix contains a formal mathematical treatment of all statements and results.

1.1. Defining CDCPs

Consider a decision-maker who chooses a set $\mathcal{L} \in \mathcal{P}(L)$ to maximize f and denote the optimal decision set by $\mathcal{L}^* \equiv \arg \max_{\mathcal{L} \in \mathcal{P}(L)} f(\mathcal{L})$. Our example throughout is a profit-maximizing firm selecting a set of foreign production locations \mathcal{L} from $\mathcal{P}(L)$. Such decision problems are straightforward to solve when the value of any location $\ell \in L$ does not depend on the decision set \mathcal{L} : the firm can evaluate each location in isolation, reducing the problem to $|L|$ independent binary decisions. However, when the value of location ℓ depends on which other locations are in \mathcal{L} , the firm must evaluate combinations of locations jointly, making the problem *combinatorial*.

To formalize the notion of interdependence among locations, we introduce the following marginal value operator:

Definition 1 (Marginal Value Operator). For all $\ell \in L$ and all $\mathcal{L} \in \mathcal{P}(L)$, the marginal value operator D_ℓ is

$$D_\ell f(\mathcal{L}) \equiv f(\mathcal{L} \cup \{\ell\}) - f(\mathcal{L} \setminus \{\ell\}).$$

We say that location choices are interdependent when the marginal value of at least a single $\ell \in L$ varies with the decision set \mathcal{L} . We now formally define CDCPs.

Definition 2 (Combinatorial Discrete Choice Problem). A combinatorial discrete choice problem is the maximization problem

$$\max_{\mathcal{L} \in \mathcal{P}(L)} f(\mathcal{L}),$$

where there exist $\ell \in L$ and $\mathcal{L}, \mathcal{L}' \in \mathcal{P}(L)$ such that $D_\ell f(\mathcal{L}) \neq D_\ell f(\mathcal{L}')$.

Without additional structure on the objective function f , no known algorithm exists to solve the generic CDCP in Definition 2 in polynomial time.² The brute-force approach evaluates f at all $2^{|L|}$ location combinations in $\mathcal{P}(L)$.

1.2. Solving CDCPs: A Single Crossing Differences Approach

In this section, we introduce a class of CDCPs whose objective function satisfies a single crossing condition and develop an iterative solution method that uses this structure to solve them without evaluating all combinations.

Single Crossing Differences in Choices Our solution method applies to CDCPs whose objective functions satisfy the following single crossing differences in choices (SCD-C) property.

Definition 3 (SCD-C). The objective function f satisfies single crossing differences in choices from above if, for every $\ell \in L$ and for all decision sets $\mathcal{L}, \mathcal{L}' \in \mathcal{P}(L)$ such that $\mathcal{L} \subset \mathcal{L}'$,

$$D_\ell f(\mathcal{L}') \geq 0 \quad \Rightarrow \quad D_\ell f(\mathcal{L}) \geq 0,$$

and single crossing differences in choices from below if

$$D_\ell f(\mathcal{L}) \geq 0 \quad \Rightarrow \quad D_\ell f(\mathcal{L}') \geq 0.$$

²The generic CDCP is NP-hard: no algorithm is known that solves all instances in time polynomial in $|L|$.

SCD-C restricts the marginal value of a location ℓ so that, along any sequence of expanding decision sets, it can cross zero at most once. SCD-C from below ensures that a location with a non-negative marginal value for some decision set \mathcal{L} retains a non-negative marginal value as additional locations are added. SCD-C from above, conversely, ensures that a location with a non-negative marginal value for some \mathcal{L} retains a non-negative marginal value as locations are removed.

Figure 1 illustrates the two SCD-C conditions. Both the solid and the dashed lines in the left panel show examples of how the marginal value of a location ℓ can vary as the location set expands from \mathcal{L}_1 to \mathcal{L}_4 when SCD-C from below holds. While the marginal value may change or even decrease, as with the solid line, it crosses zero at most once from below. The right panel shows two marginal value functions satisfying SCD-C from above, which can cross zero at most once from above.³

The properties of supermodularity and submodularity capture stronger notions of complementarity and substitutability and serve as sufficient conditions for SCD-C. A function f is supermodular if marginal values are increasing, that is, if $D_\ell f(\mathcal{L}) \leq D_\ell f(\mathcal{L}')$ for all $\mathcal{L} \subseteq \mathcal{L}'$ and all ℓ , which implies SCD-C from below. Submodularity reverses the inequality and implies SCD-C from above. While both the solid and dashed lines in Figure 1 depict marginal value functions consistent with SCD-C, only the dashed lines are consistent with super- or submodularity.

We illustrate the difference between SCD-C from below and the stronger supermodularity condition with an example. Consider a multinational firm that operates plants in Germany and the United States. Supermodularity requires that adding a plant in Canada raises the marginal value of both the German and US plants. In reality, however, the Canadian plant may lower the marginal value of the US plant—by cannibalizing its sales to the Canadian market—while raising the marginal value of the German plant by enhancing the firm’s internal economies of scale. Such a scenario violates supermodularity, but it remains consistent with SCD-C from below as long as the US plant’s marginal value stays weakly positive when the Canadian plant is added.

The Squeezing Procedure The squeezing procedure applies to CDCPs whose objective function satisfies either form of SCD-C. With SCD-C, a location’s marginal

³Milgrom and Shannon [1994] introduce single crossing conditions into economics to derive comparative statics in settings without differentiability. There is a terminological discrepancy, acknowledged in Milgrom [2004], between Milgrom and Shannon [1994], which uses *single crossing condition*, and Milgrom [2004], which uses *single crossing differences*. The condition we refer to as SCD-C is similar to, but weaker than, the quasi-supermodularity condition of Milgrom and Shannon [1994]. We formally establish the relationship between these conditions in the Online Appendix. We adopt the single crossing *differences* terminology, to emphasize that the *marginal value* of the choice changes sign at most once.




FIGURE 1: MARGINAL VALUE FUNCTIONS SATISFYING SCD-C

Note: Both panels show the marginal value of location ℓ evaluated at a sequence of expanding decision sets. In the left panel, the solid black line satisfies SCD-C from below but not supermodularity; the dashed line satisfies both. In the right panel, the solid black line satisfies SCD-C from above but not submodularity; the dashed line satisfies both. With SCD-C, marginal values can increase or decrease arbitrarily along the sequence as long as they cross zero at most once.

value crosses zero at most once, so that evaluating it at the extreme decision sets—the largest and smallest possible—can be sufficient to infer whether the location is optimal without evaluating any other decision sets. The procedure uses this logic to iteratively reduce the CDCP’s domain, removing non-optimal decision sets from consideration without explicitly evaluating each one.

We introduce a pair of *bounding sets*, $[\underline{\mathcal{L}}, \overline{\mathcal{L}}]$, comprising a lower and an upper bounding set satisfying $\underline{\mathcal{L}} \subseteq \mathcal{L}^* \subseteq \overline{\mathcal{L}}$. Each pair defines a subdomain of the original problem, and the original domain is represented by the *domain bounding sets*, $[\emptyset, L]$. We say that a pair of bounding sets $[\underline{\mathcal{K}}, \overline{\mathcal{K}}]$ is *tighter* than $[\underline{\mathcal{L}}, \overline{\mathcal{L}}]$ if $\underline{\mathcal{L}} \subseteq \underline{\mathcal{K}}$ and $\overline{\mathcal{K}} \subseteq \overline{\mathcal{L}}$.

To solve a CDCP, we introduce a mapping that, starting with the domain bounding sets, produces progressively tighter bounding sets around the optimal decision set \mathcal{L}^* , eliminating non-optimal sets from the domain. We refer to this mapping as the *squeezing step*.

Definition 4 (Squeezing step). For all $\mathcal{L}, \mathcal{L}' \in \mathcal{P}(L)$, the squeezing step is the mapping defined by:

$$S([\mathcal{L}, \mathcal{L}']) \equiv [\inf\{\Phi(\mathcal{L}), \Phi(\mathcal{L}')\}, \sup\{\Phi(\mathcal{L}), \Phi(\mathcal{L}')\}],$$

where

$$\Phi(\mathcal{L}) \equiv \{\ell \in L \mid D_{\ell}f(\mathcal{L}) \geq 0\}.$$

At the heart of the squeezing step is the mapping Φ , which identifies the set of locations $\ell \in L$ with non-negative marginal value given the decision set \mathcal{L} . Importantly, $\Phi(\mathcal{L}^*) = \mathcal{L}^*$ by construction, so that the optimal decision set is a fixed point of the squeezing step.⁴ Section 1.3 provides a detailed discussion of Jia [2008], which first introduced Φ to the economics literature.

When the objective function satisfies SCD-C, iteratively applying the squeezing step to the domain bounding sets generates a sequence of progressively tighter bounding sets that converges in at most $|L|$ steps. The following theorem formalizes this result.

Theorem 1. *If the objective function f exhibits SCD-C from above or below, iteratively applying the squeezing step to the domain bounding sets $[\underline{\mathcal{L}}, \overline{\mathcal{L}}]$ yields a sequence of bounding sets*

$$\emptyset = \underline{\mathcal{L}}^{(0)} \subseteq \dots \subseteq \underline{\mathcal{L}}^{(k)} \subseteq \underline{\mathcal{L}}^{(k+1)} \subseteq \mathcal{L}^* \subseteq \overline{\mathcal{L}}^{(k+1)} \subseteq \overline{\mathcal{L}}^{(k)} \subseteq \dots \subseteq \overline{\mathcal{L}}^{(0)} = L,$$

where k denotes the iteration count. Iteration on the squeezing step converges in $K \leq |L|$ applications.

We now sketch a constructive proof of Theorem 1.⁵ The mapping Φ is monotone if and only if the objective function f satisfies SCD-C, where the type of SCD-C determines the direction of the monotonicity. Consider two decision sets $\mathcal{L}, \mathcal{L}' \in \mathcal{P}(L)$ such that $\mathcal{L} \subset \mathcal{L}'$. With SCD-C from above, if a location has a non-negative marginal value in the larger set \mathcal{L}' , it must also have a non-negative marginal value in any nested set \mathcal{L} . As a result, $\mathcal{L} \subset \mathcal{L}'$ implies $\Phi(\mathcal{L}') \subseteq \Phi(\mathcal{L})$: the mapping Φ is *order-reversing*. With SCD-C from below, the argument is similar: if a location has a non-negative marginal value in the smaller set \mathcal{L} , then it must have a non-negative marginal value in the larger set. As a result, $\mathcal{L} \subset \mathcal{L}'$ implies $\Phi(\mathcal{L}) \subseteq \Phi(\mathcal{L}')$ and the mapping Φ is *order-preserving*.⁶

When the type of SCD-C that f satisfies is known, the monotonicity of Φ simplifies the mapping S . With SCD-C from above, Φ is order-reversing: applying Φ to the upper bounding set yields a decision set nested in the one obtained from

⁴In the Online Appendix, we show that the set of fixed points of Φ contains at least one optimal decision set of f .

⁵The squeezing step tightens the bounding sets around the optimal decision set \mathcal{L}^* that includes locations for which the firm is indifferent. In the Online Appendix, we provide a way to check whether the identified optimum is unique, and a method to recover *all* the optima if it is not.

⁶We show the converse for the order-preserving case. Suppose Φ is order-preserving, and let $\mathcal{L} \subset \mathcal{L}'$ be arbitrary decision sets. Let ℓ be an arbitrary location. Now suppose $D_{\ell}f(\mathcal{L}) \geq 0$. Then, $\ell \in \Phi(\mathcal{L}) \subseteq \Phi(\mathcal{L}')$ since Φ is order-preserving; but $\ell \in \Phi(\mathcal{L}')$ implies $D_{\ell}f(\mathcal{L}') \geq 0$ by definition of Φ . As a result, Φ being order-preserving implies that $D_{\ell}f(\mathcal{L}) \geq 0 \Rightarrow D_{\ell}f(\mathcal{L}') \geq 0$. The proof in the order-reversing case is similar.

applying Φ to the lower bounding set. The squeezing step therefore simplifies to $S([\underline{\mathcal{L}}, \overline{\mathcal{L}}]) = [\Phi(\overline{\mathcal{L}}), \Phi(\underline{\mathcal{L}})]$. With SCD-C from below, Φ is order-preserving, and the squeezing step becomes $S([\underline{\mathcal{L}}, \overline{\mathcal{L}}]) = [\Phi(\underline{\mathcal{L}}), \Phi(\overline{\mathcal{L}})]$.

The monotonicity of Φ also ensures that the sets returned when applying the squeezing step to a pair of bounding sets are themselves always a valid pair of bounding sets. In particular, consider the bounding sets $[\underline{\mathcal{L}}, \overline{\mathcal{L}}]$, so that $\underline{\mathcal{L}} \subseteq \mathcal{L}^* \subseteq \overline{\mathcal{L}}$. If f satisfies SCD-C from above, applying Φ reverses this order, so $\Phi(\overline{\mathcal{L}}) \subseteq \Phi(\mathcal{L}^*) \subseteq \Phi(\underline{\mathcal{L}})$. Moreover, since $\Phi(\mathcal{L}^*) = \mathcal{L}^*$ by construction, this expression simplifies to $\Phi(\overline{\mathcal{L}}) \subseteq \mathcal{L}^* \subseteq \Phi(\underline{\mathcal{L}})$ so that $[\Phi(\overline{\mathcal{L}}), \Phi(\underline{\mathcal{L}})]$ forms a new pair of bounding sets. Similarly, with SCD-C from below, the order-preserving nature of Φ implies $\Phi(\underline{\mathcal{L}}) \subseteq \mathcal{L}^* \subseteq \Phi(\overline{\mathcal{L}})$.

The monotonicity of Φ guarantees that, starting from the domain bounding sets $[\emptyset, L]$, iteratively applying the squeezing step produces a sequence of bounding sets that weakly tighten with each iteration. To see why, suppose f satisfies SCD-C from above, so that Φ is order-reversing. Applying the squeezing step to the domain bounding sets yields

$$\emptyset \subseteq \mathcal{L}^* \subseteq L \quad \Rightarrow \quad \Phi(L) \subseteq \mathcal{L}^* \subseteq \Phi(\emptyset).$$

The new pair of bounding sets $[\Phi(L), \Phi(\emptyset)]$ is vacuously tighter than the domain bounding sets $[\emptyset, L]$. Applying the squeezing step again produces:

$$\Phi(L) \subseteq \mathcal{L}^* \subseteq \Phi(\emptyset) \quad \Rightarrow \quad \Phi(\Phi(\emptyset)) \subseteq \mathcal{L}^* \subseteq \Phi(\Phi(L)).$$

The order-reversing property of Φ guarantees that the new upper bounding set $\Phi(\Phi(L))$ is weakly tighter than the previous one, $\Phi(\emptyset)$: since $\emptyset \subseteq \Phi(L)$, it follows that $\Phi(\Phi(L)) \subseteq \Phi(\emptyset)$. By the same logic, the new lower bounding set is weakly tighter than the previous one. We conclude that

$$\emptyset \subseteq \Phi(L) \subseteq \Phi(\Phi(\emptyset)) \subseteq \mathcal{L}^* \subseteq \Phi(\Phi(L)) \subseteq \Phi(\emptyset) \subseteq L.$$

Extending this logic inductively, each application of the squeezing step produces weakly tighter bounding sets. The procedure converges in at most $|L|$ steps, since each iteration either adds at least one location to the lower bounding set or removes at least one from the upper bounding set, and there are $|L|$ locations. An analogous argument holds when the objective function satisfies SCD-C from below. This completes our sketch of the proof of Theorem 1.

We refer to the iterative application of the squeezing step until convergence as the *squeezing procedure*, denoting the corresponding operator by S^K , and its output, $S^K([\emptyset, L])$, as the *reduced* bounding sets. If the reduced bounding sets coincide at convergence, the optimal decision set has been found. If the reduced bounding

sets do not coincide, they nevertheless define a weakly tighter subdomain of the original CDCP, on which we can identify \mathcal{L}^* either by brute-force evaluation of all remaining candidate decision sets or by a refinement procedure that we introduce next.

The Branching Procedure We introduce a recursive branching procedure to identify \mathcal{L}^* on a reduced domain $[\underline{\mathcal{L}}, \overline{\mathcal{L}}]$ with $\underline{\mathcal{L}} \subset \overline{\mathcal{L}}$. The branching step selects a location ℓ from $\overline{\mathcal{L}} \setminus \underline{\mathcal{L}}$, and creates two branches: the first imposes that ℓ is included in the firm's decision set, while the other imposes that ℓ is excluded. Each branch defines a *restricted* CDCP, specifically the problem of choosing $\mathcal{L} \subseteq (\overline{\mathcal{L}} \setminus \underline{\mathcal{L}}) \setminus \{\ell\}$ to maximize the modified objective function $\tilde{f} : \mathcal{P}((\overline{\mathcal{L}} \setminus \underline{\mathcal{L}}) \setminus \{\ell\}) \rightarrow \mathbb{R}$ where:

$$\begin{aligned} \tilde{f}(\mathcal{L}) &\equiv f(\mathcal{L} \cup \underline{\mathcal{L}} \cup \{\ell\}) && \text{on the branch that includes } \ell, \\ \tilde{f}(\mathcal{L}) &\equiv f(\mathcal{L} \cup \underline{\mathcal{L}}) && \text{on the branch that excludes } \ell. \end{aligned}$$

The branching step applies the squeezing procedure until convergence to each restricted CDCP, yielding a pair of reduced bounding sets for each branch. The branching step then translates these bounding sets for the restricted CDCP into bounding sets for the original CDCP by adding all items in $\underline{\mathcal{L}}$ to both bounding sets on each branch. On the branch that includes ℓ , the item ℓ is also added to both bounding sets.

The output of the branching step is two pairs of conditional bounding sets for the original CDCP, one on each branch. If the bounding sets on a branch coincide, that branch has converged to a *conditionally* optimal decision set of the original CDCP; otherwise, the branching step is applied recursively. The recursive application of the branching step creates a tree in which each terminal node corresponds to a conditionally optimal decision set of the original CDCP.

We refer to the recursive application of the branching step until convergence on all branches as the *branching procedure* and define $\Lambda([\underline{\mathcal{L}}, \overline{\mathcal{L}}])$ as the collection of conditionally optimal decision sets across the terminal nodes of all branches. The globally optimal decision set is $\mathcal{L}^* = \arg \max_{\mathcal{L} \in \Lambda([\underline{\mathcal{L}}, \overline{\mathcal{L}}])} f(\mathcal{L})$.⁷

Figure 2 illustrates the branching procedure: for a given pair of reduced bounding sets, two branches are created by applying the branching step with $\ell \in \overline{\mathcal{L}} \setminus \underline{\mathcal{L}}$. The branch on the left, which includes ℓ , converges to a pair of bounding sets with identical lower and upper bounds, identifying \mathcal{L}_1^* as the optimal decision set

⁷While $\Lambda([\underline{\mathcal{L}}, \overline{\mathcal{L}}])$ may depend on the locations selected for branching, we show in the Online Appendix that it always contains the global maximum, so the branching procedure correctly identifies \mathcal{L}^* regardless of the branching order.




FIGURE 2: THE BRANCHING PROCEDURE

Note: This figure shows a tree of subproblems resulting from applying the branching procedure recursively. Convergence on a single branch occurs when the squeezing procedure returns a conditionally optimal set, indicated by a terminal node. The final output of the full recursive procedure is the collection of all conditionally optimal sets, in this example $\{\mathcal{L}_1^*, \mathcal{L}_2^*, \mathcal{L}_3^*\}$.

conditional on including ℓ . The branch on the right, which excludes ℓ , converges to a pair of bounding sets which do not coincide, resulting in the bounding sets $[\overline{\mathcal{L}'}, \underline{\mathcal{L}'}]$ on the original domain, so branching recurs: a new location $\ell' \in \overline{\mathcal{L}'} \setminus \underline{\mathcal{L}'}$ is selected, creating an additional pair of branches. This recursive process builds a tree, terminating in a conditionally optimal decision set on each branch so that $\Lambda([\underline{\mathcal{L}}, \overline{\mathcal{L}}]) = \{\mathcal{L}_1^*, \mathcal{L}_2^*, \mathcal{L}_3^*\}$.

Branching is an alternative to the brute-force approach of evaluating the objective function at every decision set in the reduced domain. Intuitively, the branching procedure incorporates the brute-force logic one location at a time: it applies the squeezing procedure as much as possible, and when squeezing alone cannot identify the optimum, it fixes one location's status and repeats. In the worst case, when squeezing eliminates no decision set on any branch, the branching procedure evaluates the same decision sets as the brute-force approach.⁸

1.3. The Mathematics of Squeezing and Branching

This section explains the mathematical foundation of our squeezing method, formally connects it to Jia [2008], and identifies the new challenges that arise when the

⁸The wall-clock time advantage of branching over the brute-force approach depends on how effectively the squeezing procedure reduces the domain on each branch. When the squeezing procedure reduces the domain by little on each branch, the cost of repeated application may offset the savings from evaluating the objective function over fewer decision sets.

objective function satisfies SCD-C from above rather than below.

Jia [2008], Supermodularity, and SCD-C from Below Jia [2008] studies the location decisions of a retail chain, assuming individual store locations are complements so that the objective function is supermodular. In this context, Jia [2008] introduces the Φ mapping and recasts finding the optimal set of store locations, \mathcal{L}^* , as finding the fixed points of Φ , since $\Phi(\mathcal{L}^*) = \mathcal{L}^*$ by construction.⁹

Since the objective function is supermodular, Φ is order-preserving and Jia [2008] invokes the theorem of Tarski [1955] to guarantee the existence of well-defined smallest and largest fixed points, \mathcal{L}^{inf} and \mathcal{L}^{sup} . Together, these fixed points form a pair of bounding sets $[\mathcal{L}^{\text{inf}}, \mathcal{L}^{\text{sup}}]$, since necessarily $\mathcal{L}^{\text{inf}} \subseteq \mathcal{L}^* \subseteq \mathcal{L}^{\text{sup}}$.

Jia [2008] identifies the bounding sets $[\mathcal{L}^{\text{inf}}, \mathcal{L}^{\text{sup}}]$ by applying Φ to \emptyset and L separately until convergence. This approach is rooted in Kleene’s fixed point theorem, which states that iteratively applying an order-preserving map to \emptyset always converges to its smallest fixed point \mathcal{L}^{inf} , while applying it to L always converges to its largest fixed point \mathcal{L}^{sup} .¹⁰ In cases where the bounding sets $[\mathcal{L}^{\text{inf}}, \mathcal{L}^{\text{sup}}]$ do not identify \mathcal{L}^* , Jia [2008] applies the brute-force approach on the reduced domain.

In the case of SCD-C from below, our squeezing procedure implements the same method as Jia [2008]. With an order-preserving Φ , the squeezing step simplifies to $S([\underline{\mathcal{L}}, \overline{\mathcal{L}}]) = [\Phi(\underline{\mathcal{L}}), \Phi(\overline{\mathcal{L}})]$, which is equivalent to applying Φ separately to the lower and upper bounding sets. Consequently, the reduced bounding sets produced by the squeezing procedure coincide with the smallest and largest fixed points of Φ , so that $S^K([\emptyset, L]) = [\mathcal{L}^{\text{inf}}, \mathcal{L}^{\text{sup}}]$.

Our first contribution relative to Jia [2008] is to extend the applicability of the paper’s method beyond supermodular objective functions by showing that SCD-C from below is the necessary and sufficient condition for Φ to be order-preserving.

The Challenges of SCD-C from Above and Fixed Edges Our second contribution is to develop a solution method that applies when the objective function satisfies SCD-C from above. The key difficulty in this case is that Φ is order-reversing, so that the fixed point theorems of Tarski [1955] and Kleene do not apply; a smallest and largest fixed point bounding the optimal decision set may not exist.

⁹The mapping Φ has antecedents in the Operations Research literature on Boolean optimization. It presents a simple updating rule for decisions based on the discrete analogue of a first order condition [see, e.g., Boros and Hammer, 2002].

¹⁰See Stoltenberg-Hansen et al. [1994] for an exposition. The theorem derives from results in Kleene [1936] and Kleene [1938], though it was not stated explicitly there.

To see why, consider a firm choosing between two perfectly substitutable production locations, $L = \{\text{USA}, \text{CAN}\}$. Suppose that each location has positive marginal value when the other is inactive, but neither has positive marginal value when both are active, as each fully substitutes for the other. Applying the approach in [Jia \[2008\]](#) yields $\Phi(\emptyset) = \{\text{USA}, \text{CAN}\}$ and $\Phi(\{\text{USA}, \text{CAN}\}) = \emptyset$: iterating from either extreme oscillates and does not converge to a fixed point.

Although repeated application of Φ may not converge to a fixed point with SCD-C from above, it does converge to a pair of sets between which it alternates, as in the example above. We call a pair $[\mathcal{L}, \mathcal{L}']$ with $\mathcal{L} \subseteq \mathcal{L}'$ satisfying $\mathcal{L} = \Phi(\mathcal{L}')$ and $\mathcal{L}' = \Phi(\mathcal{L})$ a *fixed edge* of Φ . Fixed edges generalize the notion of fixed points: if \mathcal{L} is a fixed point of Φ , then $[\mathcal{L}, \mathcal{L}]$ is a fixed edge.

[Klimeš \[1981\]](#) shows that, although an order-reversing map need not have smallest and largest fixed points, it always has a fixed edge $[\mathcal{L}^{\text{inf}}, \mathcal{L}^{\text{sup}}]$ that is extreme in the sense that $\mathcal{L}^{\text{inf}} \subseteq \mathcal{L} \subseteq \mathcal{L}' \subseteq \mathcal{L}^{\text{sup}}$ for all other fixed edges $[\mathcal{L}, \mathcal{L}']$.¹¹ The extreme fixed edge therefore serves as a pair of bounding sets, just as the smallest and largest fixed points do in the order-preserving case. With SCD-C from above, the squeezing procedure applied to the domain bounding sets converges to the extreme fixed edge of Φ .

The insights in this section imply that the squeezing step S is itself an order-preserving mapping whenever the underlying objective function satisfies SCD-C. An alternative proof of [Theorem 1](#) therefore follows by applying the theorems of [Tarski \[1955\]](#) and Kleene to S directly.

The squeezing procedure converges to $[\mathcal{L}^*, \mathcal{L}^*]$ if and only if \mathcal{L}^* is the unique fixed point of Φ when f satisfies SCD-C from below, or the unique fixed edge of Φ when f satisfies SCD-C from above. When multiple fixed points or fixed edges exist, the branching procedure recovers all fixed points: every fixed point is a terminal node of the branching tree.

1.4. Solving CDCPs for Heterogeneous Agents

We now extend the framework to a setting with heterogeneous agents. We consider an objective function $f : \mathcal{P}(L) \times \mathbb{R} \rightarrow \mathbb{R}$ that maps a decision set $\mathcal{L} \in \mathcal{P}(L)$ and an agent type $z \in \mathbb{R}$ to a scalar payoff. We focus on single-dimensional heterogeneity

¹¹The intuition behind [Klimeš \[1981\]](#) is simple. If Φ is order-reversing, then $\Phi^2 \equiv \Phi \circ \Phi$ is order-preserving. By [Tarski \[1955\]](#), Φ^2 has a smallest and largest fixed point. Any fixed point of Φ^2 satisfies $\Phi(\Phi(\mathcal{L})) = \mathcal{L}$, so the pair $[\mathcal{L}, \Phi(\mathcal{L})]$ is a fixed edge of Φ .

and assume that the objective function is continuous in z for all $\mathcal{L} \in \mathcal{P}(L)$.¹² Solving CDCPs for heterogeneous agents amounts to recovering the policy function that maps each agent’s type to its optimal decision set.

Definition 5 (Policy function). The policy function $\mathcal{L}^* : \mathbb{R} \rightarrow \mathcal{P}(L)$ maps each type z to its optimal decision set, so that $\mathcal{L}^*(z) \equiv \arg \max_{\mathcal{L} \in \mathcal{P}(L)} f(\mathcal{L}, z)$.

In our multinational production example, the policy function maps each firm’s productivity to its optimal combination of production locations.¹³

We introduce a method to solve CDCPs with heterogeneous agents whenever the objective function satisfies SCD-C and an additional single crossing condition which we introduce next.

Single Crossing Differences in Type The single crossing differences in type (SCD-T) property restricts how the marginal value of a location varies with firm productivity.

Definition 6 (SCD-T). The objective function f exhibits single crossing differences in type if, for all $\ell \in L$, decision sets $\mathcal{L} \in \mathcal{P}(L)$, and types $z, z' \in \mathbb{R}$ such that $z < z'$:

$$D_{\ell} f(\mathcal{L}, z) \geq 0 \quad \Rightarrow \quad D_{\ell} f(\mathcal{L}, z') \geq 0.$$

SCD-T ensures that if a location ℓ has a non-negative marginal value as part of a decision set \mathcal{L} for a firm with productivity z , then it also has a non-negative marginal value for any firm with productivity $z' > z$. If instead the marginal value of ℓ crosses zero from above, the problem satisfies SCD-T with the transformed type $-z$.¹⁴

Just as SCD-C enables the squeezing procedure, which rules out decision sets without evaluating their payoff, SCD-T enables a *generalized* squeezing procedure that rules out decision sets for entire productivity ranges at once without evaluating their payoffs for any individual productivity value within those ranges.

¹²In an earlier working paper, we showed how to extend the methods in this section to settings with multidimensional heterogeneity [see [Arkolakis et al., 2023](#)].

¹³The policy function is well-defined when the optimal decision set is unique for each z except at cutoff productivity values. The Online Appendix provides a sufficient condition.

¹⁴This property is similar to, but weaker than, the single crossing property introduced by [Milgrom and Shannon \[1994\]](#), later referred to as the single crossing *differences* condition by [Milgrom \[2004\]](#). We formally establish the relationship between these conditions in the Online Appendix. The firm type z can represent any characteristic, endogenous or exogenous, that affects payoffs, and the policy function $\mathcal{L}^*(z)$ describes how optimal choices respond to changes in that characteristic. When z is endogenous, the policy function is often referred to as a best-response function.

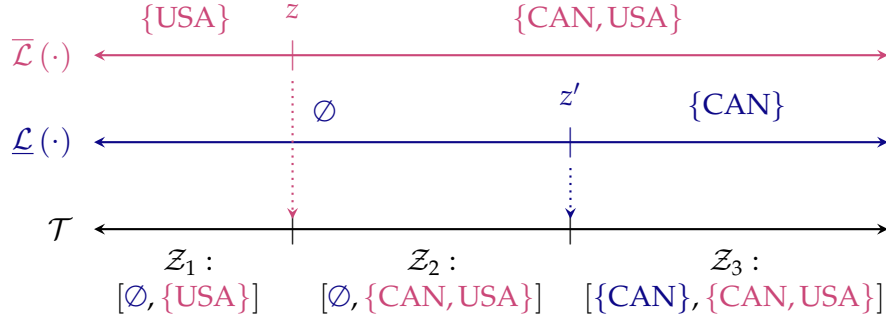


FIGURE 3: PARTITIONING THE PRODUCTIVITY SPACE BY COMMON BOUNDING SETS

Note: The middle line shows a lower bounding function, the top line an upper bounding function. Together, these two set-valued functions imply the partitioning \mathcal{T} , which creates intervals of productivity. In this figure, there are three intervals, so $\mathcal{T} = \{Z_1, Z_2, Z_3\}$. All productivity values within an interval Z_t share the same bounding sets.

The Generalized Squeezing Procedure We extend the notion of bounding sets to *bounding functions*: set-valued functions of firm productivity $\underline{\mathcal{L}}(\cdot)$ and $\overline{\mathcal{L}}(\cdot)$, such that $\underline{\mathcal{L}}(z) \subseteq \mathcal{L}^*(z) \subseteq \overline{\mathcal{L}}(z)$ for all $z \in \mathbb{R}$. We refer to $[\underline{\mathcal{L}}(\cdot), \overline{\mathcal{L}}(\cdot)] \equiv [\emptyset, L]$ as the domain bounding functions, since these constant functions nest the policy function for all z .

Any pair of bounding functions induces a partition of the firm productivity space into intervals over which the lower and upper bounding sets are constant. We define this partition as:

$$\begin{aligned} \mathcal{T}([\underline{\mathcal{L}}(\cdot), \overline{\mathcal{L}}(\cdot)]) &= \{Z_1, \dots, Z_t, \dots, Z_T\} \\ \text{such that } Z_t &= \{z \in \mathbb{R} \mid \underline{\mathcal{L}}(z) = \underline{\mathcal{L}}_t, \overline{\mathcal{L}}(z) = \overline{\mathcal{L}}_t\}, \end{aligned}$$

where t indexes the intervals over which neither bounding function changes its value. When the bounding functions are clear from context, we write \mathcal{T} for brevity.

Figure 3 illustrates how a pair of bounding functions partitions the productivity space into intervals. The middle and top lines represent the lower and upper bounding functions respectively, each changing value at different cutoffs, z and z' . The line at the bottom shows the resulting partition: the productivity intervals over which both bounding functions are constant. These intervals identify the cutoff productivity values at which either the lower or upper bounding function changes value.

We now define a generalized version of the squeezing step above.

Definition 7 (Generalized squeezing step). The generalized squeezing step is the mapping S^g defined for all $\mathcal{L}(\cdot), \mathcal{L}'(\cdot) : \mathbb{R} \rightarrow \mathcal{P}(L)$ as:

$$S^g([\mathcal{L}(\cdot), \mathcal{L}'(\cdot)]) \equiv [\inf\{\Phi^g(\mathcal{L}(\cdot), \cdot), \Phi^g(\mathcal{L}'(\cdot), \cdot)\}, \sup\{\Phi^g(\mathcal{L}(\cdot), \cdot), \Phi^g(\mathcal{L}'(\cdot), \cdot)\}]$$

where the mapping $\Phi^g : \mathcal{P}(L) \times \mathbb{R} \rightarrow \mathcal{P}(L)$ is defined as

$$\Phi^g(\mathcal{L}, z) \equiv \{\ell \in L \mid z \geq z_\ell^g(\mathcal{L})\}$$

and the functions $z_\ell^g : \mathcal{P}(L) \rightarrow \mathbb{R}$ are defined as $z_\ell^g(\mathcal{L}) \equiv \inf\{z \in \mathbb{R} \mid D_\ell(\mathcal{L}, z) = 0\}$ for each $\ell \in L$.

The generalized squeezing step extends the logic of Definition 4 to entire ranges of productivity values at a time when the objective function satisfies SCD-C and SCD-T. For all \mathcal{L} and z , $\Phi^g(\mathcal{L}, z)$ coincides with $\Phi(\mathcal{L})$ when the objective function is $f(\cdot, z)$. For all locations ℓ and decision sets \mathcal{L} , SCD-T implies that, rather than computing the marginal value of ℓ at each z , it suffices to check whether z lies above or below the cutoff $z_\ell^g(\mathcal{L})$ at which firms are indifferent to including ℓ in \mathcal{L} . Thus, the bounding functions update for entire productivity ranges at once, rather than point by point.¹⁵

When the objective function satisfies both SCD-C and SCD-T, iteratively applying the generalized squeezing step to the domain bounding functions returns a weakly tighter pair at each iteration. Iteration converges once $S^g([\underline{\mathcal{L}}(\cdot), \overline{\mathcal{L}}(\cdot)]) = [\underline{\mathcal{L}}(\cdot), \overline{\mathcal{L}}(\cdot)]$. Since each bounding function can tighten at most $|L|$ times, iteration converges in at most $|L|$ applications. The following theorem formalizes this result.

Theorem 2. *If the objective function f exhibits SCD-C and SCD-T, iteratively applying the generalized squeezing step to the domain bounding functions $[\emptyset, L]$ yields a sequence of bounding functions so that for all $z \in \mathbb{R}$,*

$$\emptyset \subseteq \dots \subseteq \underline{\mathcal{L}}^{(k)}(z) \subseteq \underline{\mathcal{L}}^{(k+1)}(z) \subseteq \mathcal{L}^*(z) \subseteq \overline{\mathcal{L}}^{(k+1)}(z) \subseteq \overline{\mathcal{L}}^{(k)}(z) \subseteq \dots \subseteq L,$$

where k denotes the iteration count. Iteration on the generalized squeezing step converges in $K \leq |L|$ applications.

We now illustrate how the generalized squeezing step proceeds for an objective function satisfying SCD-C from above and SCD-T. Consider the domain bounding functions, $[\emptyset, L]$, and the associated productivity partition \mathcal{T} . Since the domain

¹⁵If $D_\ell(\mathcal{L}, z) < 0$ for all z , we define $z_\ell^g(\mathcal{L}) \equiv \infty$; if $D_\ell(\mathcal{L}, z) > 0$ for all z , we define $z_\ell^g(\mathcal{L}) \equiv -\infty$. The range of each z_ℓ^g is thus the two-point compactification of the real line $\mathbb{R} \cup \{-\infty, \infty\}$.

bounding functions are constant for all z , the partition of the productivity space contains a single interval comprising the full productivity range. For each ℓ , we compute two cutoffs: $z_\ell^s(\emptyset)$, at which the marginal value of ℓ is zero in \emptyset , and $z_\ell^s(L)$ at which it is zero in L . The SCD-C and SCD-T conditions together imply $z_\ell^s(\emptyset) \leq z_\ell^s(L)$.

The indifference cutoffs $z_\ell^s(\emptyset)$ and $z_\ell^s(L)$ for each ℓ update the bounding functions as follows. For all firms with productivity $z < z_\ell^s(\emptyset)$, ℓ is not part of the optimal decision set, so the upper bounding function updates to $L \setminus \{\ell\}$. Conversely, for all firms with productivity $z > z_\ell^s(L)$, ℓ is part of the decision set, so the lower bounding function updates to $\{\ell\}$.

Figure 3 depicts the outcome of updating the domain bounding functions for $\ell = \text{CAN}$. The figure shows the CAN-specific cutoffs $z_{\text{CAN}}^s(\emptyset)$ as z and $z_{\text{CAN}}^s(L)$ as z' . The updated bounding functions now vary with productivity, changing value at $z_{\text{CAN}}^s(\emptyset)$ and $z_{\text{CAN}}^s(L)$, and hence partition the productivity space into three intervals. A full update of the bounding functions requires computing such cutoffs for all locations ℓ . In subsequent applications, the generalized squeezing step applies to each partition interval separately, returning a finer partition and thus tightening the bounding functions around the policy function.

We refer to the iterative application of the generalized squeezing step until convergence as the *generalized squeezing procedure* and the resulting tightened bounding functions as *reduced* bounding functions. The reduced bounding functions pin down the policy function on any interval for which they coincide.

For productivity intervals for which bounding functions differ, Appendix A defines two refinement methods that converge to the policy function on all intervals of the productivity space. The first, the *generalized branching procedure*, extends the branching procedure from Section 1.2 to the heterogeneous-agent case. The second, *iterative cutoff search*, requires the objective function to satisfy a stronger version of SCD-T, which we call *strong SCD-T*. Strong SCD-T requires that, for any pair of decision sets \mathcal{L} and \mathcal{L}' , the difference in payoffs $f(\mathcal{L}, z) - f(\mathcal{L}', z)$ crosses zero for at most one value of z . Strong SCD-T implies SCD-T, but not vice versa: SCD-T only requires such single crossing for pairs $(\mathcal{L} \setminus \{\ell\}, \mathcal{L} \cup \{\ell\})$ that differ by a single element.¹⁶ Strong SCD-T, which holds in the application of Section 2, imparts additional structure on the policy function that iterative cutoff search leverages.

We refer to the application of generalized squeezing and its refinement until convergence as the *policy function method*. In contrast to discretization or bisection methods, the policy function method uses the monotonicity afforded by SCD-T to

¹⁶SCD-T also allows the crossing at zero to occur for an interval of productivity values instead of at a single value; strong SCD-T rules out this possibility.

identify the exact cutoffs at which the policy function changes value. This approach eliminates unnecessary computation at non-cutoff productivity values and avoids the interpolation required by discretization methods. Section 4 shows that such interpolation can introduce substantial errors in aggregate variables in the context of a calibrated multinational production model.

2. A Quantitative Model of Multinational Production

In this section, we introduce a quantitative model of multinational production (MP) and derive conditions under which our solution method applies. In the model, a firm's problem of choosing multinational production locations maps directly into the CDCP framework of Section 1: the set of candidate locations L corresponds to the set of countries, the objective function f to the firm's operating profits, and the agent type z to firm productivity. The single crossing conditions translate into restrictions on the demand and cost elasticities. In the Online Appendix, we extend the framework to accommodate a broader class of demand and cost functions.

Environment The world economy consists of a discrete set of countries L . Firms are headquartered in one country, produce in a set of countries, and sell to consumers in all countries. We index headquarter locations by i , production locations by ℓ , and final consumption destinations by n . Each firm produces a differentiated final good variety with a firm-specific productivity z . Since firms with the same headquarter location and productivity behave identically, we index them by i and z alone. In every location ℓ , a mass of H_ℓ households inelastically supplies one unit of labor at wage w_ℓ . Labor markets are perfectly competitive and output markets are monopolistically competitive.

Demand System Consumers in all destinations have identical CES preferences with elasticity $\sigma > 1$ over the set of available final goods. The demand function for an individual variety as a function of its destination-specific price p_n is

$$q_n(p_n) = Q_n \left(\frac{p_n}{P_n} \right)^{-\sigma},$$

where Q_n is the CES consumption basket and P_n the associated ideal price index. Firms and consumers take the aggregate objects Q_n and P_n as given.

Production Technology Consider a firm headquartered in location i with productivity z that operates a set of production locations $\mathcal{L} \subseteq L$. The firm delivers its final good to destination n by combining production from its locations ℓ . The firm's marginal cost of serving destination n is given by a CES aggregator of the marginal cost at each of its production locations:

$$(1) \quad c_{in}(\mathcal{L}, z) = \frac{1}{z} \left[\sum_{\ell \in \mathcal{L}} \zeta_{i\ell n}^{1-\varepsilon} \right]^{\frac{1}{1-\varepsilon}}, \quad \zeta_{i\ell n} \equiv \gamma_{i\ell} \frac{w_\ell}{A_\ell} \tau_{\ell n}.$$

The parameter $\varepsilon > 1$ denotes the elasticity of substitution among production locations. For each production location, w_ℓ is the equilibrium wage rate and A_ℓ is an exogenous location-specific productivity shifter. Firms also face a bilateral MP cost $\gamma_{i\ell}$, which captures factors such as communication or language costs, as well as a bilateral cost of trade $\tau_{\ell n}$. We summarize all cost shifters in a trilateral resistance term denoted $\zeta_{i\ell n}$.

Equation (1) shows that, all else equal, expanding the production location set lowers a firm's marginal cost. However, the cost savings from each additional location fall as the production location set grows. As $\varepsilon \rightarrow \infty$, firms only use their lowest-cost location.

The marginal cost aggregator in equation (1) captures, in reduced form, that firms use output from all their production locations to assemble their final good for any destination. One microfoundation is that, for each destination market, a firm's locations compete to be the lowest-cost supplier for each of a continuum of inputs into the final product. Adding a production location expands the set of potential suppliers and lowers the firm's marginal cost, but at a diminishing rate: as the location set grows, the probability that any given location is the lowest-cost supplier for a specific input declines. In the Online Appendix, we detail this microfoundation with Fréchet-distributed cost shocks, as in [Antràs et al. \[2017\]](#) and [Tintelnot \[2017\]](#).¹⁷ The equilibrium allocations in our model depend only on the aggregator in equation (1) and not on the underlying microfoundation.

Profit Maximization Consider a firm headquartered in location i with productivity z . The firm's profit maximization problem has two parts: choosing a set of production locations \mathcal{L} and setting the price of its final good in each destination

¹⁷ [Antràs et al. \[2017\]](#) model firms that pay fixed costs to add sourcing partners subject to Fréchet cost shocks, while [Tintelnot \[2017\]](#) studies firms that establish plants across locations with plant-specific Fréchet shocks. [Tintelnot \[2017\]](#) restricts the Fréchet shape parameter to exceed a common demand and cost elasticity, which implies SCD-C from above. Our framework allows the two elasticities to differ. Alternative foundations include multivariate Pareto shocks [[Arkolakis et al., 2018](#)] or intermediate goods that are imperfect substitutes across locations [[Antràs et al., 2024b](#)].

market. For a given \mathcal{L} , the firm sets its price in destination n as a constant markup over marginal cost, $c_{in}(\mathcal{L}, z)$, so that

$$p_{in}(\mathcal{L}, z) = \frac{\sigma}{\sigma - 1} c_{in}(\mathcal{L}, z).$$

Given its optimal pricing policy, the firm chooses the optimal location decision set $\mathcal{L}_i^*(z)$ to maximize its operating profits:

$$(2) \quad \max_{\mathcal{L} \in \mathcal{P}(L)} \pi_i(\mathcal{L}, z) \equiv \max_{\mathcal{L} \in \mathcal{P}(L)} \left\{ \sum_{n \in L} \frac{1}{\sigma} q_n(p_{in}(\mathcal{L}, z)) p_{in}(\mathcal{L}, z) - \sum_{\ell \in L} w_\ell f_{i\ell} \right\},$$

where $f_{i\ell}$ is the fixed cost of producing in location ℓ , denominated in location ℓ production labor. Note that, with CES preferences, variable profits are a share $1/\sigma$ of revenue.

Finally, the firm must first pay a labor-denominated entry cost f_i^e to draw a productivity z from an exogenous distribution $G_i(z)$. The firm produces positive quantities if it has non-negative operating profits; otherwise, it exits right after paying f_i^e .

Aggregation and the Equilibrium System We now turn to aggregation and the equilibrium determination of aggregate variables.

The first equilibrium condition is a zero profit condition that pins down the cutoff productivity \tilde{z}_i below which firms would have negative operating profits and thus exit instead:

$$(3) \quad \pi_i(\mathcal{L}_i^*(\tilde{z}_i), \tilde{z}_i) = 0.$$

The second equilibrium condition is a free entry condition that pins down the total mass of entrants M_i in each origin country i by equalizing the entry cost on the left with the expected operating profits on the right:

$$(4) \quad w_i f_i^e = \int_{\tilde{z}_i}^{\infty} \pi_i(\mathcal{L}_i^*(z), z) dG_i(z).$$

Price indices in each destination market n aggregate the prices of all goods offered:

$$(5) \quad P_n^{1-\sigma} = \sum_{i \in L} M_i \int_{\tilde{z}_i}^{\infty} p_{in}(\mathcal{L}_i^*(z), z)^{1-\sigma} dG_i(z).$$

In addition, the labor market must clear in each production location ℓ . There are three sources of labor demand: variable labor requirements from all the production

sites operating in country ℓ , the entry costs of all the firms headquartered in ℓ , and the fixed costs incurred by foreign and domestic firms to set up production in location ℓ . Labor market clearing equates the total labor supply in location ℓ to total labor demand as follows:

$$(6) \quad w_\ell H_\ell = \frac{\sigma - 1}{\sigma} \sum_{i \in L} \sum_{n \in L} M_i X_n \int_{\tilde{z}_i}^{\infty} \frac{\mathbb{1}_{i\ell}(z) \tilde{\zeta}_{i\ell n}^{1-\varepsilon}}{\sum_{k \in \mathcal{L}_i^*(z)} \tilde{\zeta}_{ikn}^{1-\varepsilon}} \left[\frac{p_{in}(\mathcal{L}_i^*(z), z)}{P_n} \right]^{1-\sigma} dG_i(z) \\ + M_\ell w_\ell f_\ell^\varepsilon + \sum_{i \in L} M_i w_\ell f_{i\ell} \int_{\tilde{z}_i}^{\infty} \mathbb{1}_{i\ell}(z) dG_i(z),$$

where X_n denotes total expenditure on final goods in destination n and the indicator $\mathbb{1}_{i\ell}(z) \equiv \mathbb{1}\{\ell \in \mathcal{L}_i^*(z)\}$ is 1 if a firm headquartered in i with productivity z opens a production location in country ℓ .¹⁸

Lastly, balance of payments requires that total expenditure from consumers in a market n equals their total income:

$$(7) \quad X_n = w_n H_n.$$

The general equilibrium in our model is defined as follows.

Definition. General equilibrium is a set of policy functions $\{\mathcal{L}_i(\cdot)\}_i$ and aggregate variables $\{w_i, \tilde{z}_i, M_i, P_i, X_i\}_i$ so that

1. given the aggregate variables, the policy functions solve the firm optimization problem in equation (2); and
2. given the policy functions, the aggregate variables satisfy equations (3)–(7).

Establishing the SCD-C and SCD-T Conditions We conclude this section by establishing the SCD-C and SCD-T properties in the model.

To build intuition, consider a symmetric version of our economy where $\tilde{\zeta}_{i\ell n} = \tilde{\zeta}$, so that the marginal value of an additional production location depends only on the number of active locations, not on their identity:

$$(8) \quad D_\ell \pi_i(\mathcal{L}, z) = \mathcal{G} z^{\sigma-1} \left(|\mathcal{L} \cup \{\ell\}|^{\frac{\sigma-1}{\varepsilon-1}} - |\mathcal{L} \setminus \{\ell\}|^{\frac{\sigma-1}{\varepsilon-1}} \right) - w_\ell f_{i\ell},$$

where \mathcal{G} is a composite general equilibrium constant. A single sufficient statistic, the ratio $(\sigma - 1)/(\varepsilon - 1)$, governs whether locations are complements or substitutes

¹⁸Note that, with CES demand, a fraction $(\sigma - 1)/\sigma$ of a firm's total sales are variable production costs; location $\ell \in \mathcal{L}_i^*(z)$ accounts for a share $\mathbb{1}_{i\ell}(z) \tilde{\zeta}_{i\ell n}^{1-\varepsilon} / \sum_{k \in \mathcal{L}_i^*(z)} \tilde{\zeta}_{ikn}^{1-\varepsilon}$ of these costs.

and determines the type of SCD-C the profit function satisfies. When $(\sigma - 1)/(\varepsilon - 1) > 1$, the marginal value of any location is increasing in the number of other locations, so that supermodularity and hence SCD-C from below holds. When $(\sigma - 1)/(\varepsilon - 1) < 1$, the marginal value of any location is decreasing in the number of other locations, so that submodularity and hence SCD-C from above holds. When the ratio is one, the marginal value is independent of the firm's decision set \mathcal{L} .¹⁹

The ratio $(\sigma - 1)/(\varepsilon - 1)$ summarizes the net effect of two opposing forces that generate interdependencies among production locations. On the demand side, the elasticity of substitution $\sigma > 1$ governs the firm's ability to scale: since demand is price-sensitive, the same price reduction translates to a larger sales volume if the firm is already low-price, so that each additional production location raises the value of the next. On the cost side, since $\varepsilon > 1$, production locations are substitutes, as they effectively compete with one another to supply each destination. When $(\sigma - 1)/(\varepsilon - 1) > 1$, the demand side complementarity dominates the supply side substitutability.

Equation (8) also shows that SCD-T is always satisfied in our model since the marginal value of location ℓ is a strictly increasing function of the firm's productivity type given $\sigma > 1$. Moreover, this parameter restriction ensures strong SCD-T: for any two decision sets, there is at most one productivity at which the firm is indifferent between them.

As we show formally in the Online Appendix, these parameter conditions remain the sufficient conditions for SCD-C and strong SCD-T in the full model, so that the policy function method with iterative cutoff search is applicable.²⁰

¹⁹Equation (8) highlights that the firm's CDCP arises from the combination of two ingredients: interdependencies among locations and the fixed costs of setting up production locations. Without interdependencies, that is, when $(\sigma - 1)/(\varepsilon - 1) = 1$, the marginal value of any production location is independent of all others and hence the firm faces $|L|$ independent decisions. Without fixed costs, that is, when $f_{i\ell} = 0 \forall i, \ell$, the marginal value of all locations is always non-negative, so firms establish production locations in all countries, regardless of interdependencies.

²⁰In the Online Appendix, we also show how to verify SCD-C and SCD-T for more general demand functions $q_n(p_n)$ and marginal cost functions $c_{in}(\mathcal{L}, z)$ of which the CES functions in this section are special cases. Our generalized framework nests demand systems as in [Arkolakis et al. \[2019\]](#) and production structures as in [Ramondo \[2014\]](#), [Arkolakis et al. \[2018\]](#), and [Lind and Ramondo \[2023\]](#). In general, the strength of the complementarities or substitutabilities across locations depends on how changes in marginal costs translate into differences in variable profits, which is determined by the elasticity of demand and the pass-through elasticity of costs to price. In the present model, markups are constant so that the pass-through elasticity is always 1; consequently, only the elasticity of demand matters.

3. Quantification

We calibrate the model to data on trade and MP across 32 countries. As part of the calibration routine, we solve the full heterogeneous-agent general equilibrium model for each parameter guess, a task our policy function method makes computationally feasible. This section provides an overview of the calibration strategy, with further details on how it uses our solution methods in Appendix C. The Online Appendix discusses data treatment and model fit.

3.1. Data

We use information on manufacturing trade and MP for 32 OECD and European countries from [Alviarez \[2019\]](#). For each host country, the dataset contains the number of foreign affiliates in manufacturing and their total sales by origin country. For example, the dataset contains the number of Canadian affiliates of German firms and their total sales. The dataset also contains bilateral trade flows for all country pairs. For all variables, the data reflect average values over the period from 2003 to 2012.²¹

While the dataset contains foreign affiliate counts for each country, it lacks total firm counts and information on firm entry and survival. To supplement the dataset, we obtain total firm counts for each country from the UNIDO Industrial Statistics Database, and data on the one-year survival rates of firms in each country from the OECD Structural and Demographic Business Statistics.

To construct standard bilateral gravity controls, we use the CEPII database [see [Conte et al., 2023](#)], which provides measures of geographic distance, a common border indicator, a shared colonial past indicator, and a shared language indicator. In addition, we obtain tariff data from the Global Tariff database [see [Teti, 2024](#)] and data on real GDP per capita and total employment from the Penn World Table [see [Feenstra et al., 2015](#)].

3.2. Calibration Strategy

Table 1 summarizes our calibration strategy, listing each parameter and the moment it targets. We set σ and ε externally and then use country-level aggregate outcomes together with bilateral trade, MP, and affiliate count flows to discipline the remaining parameters.

²¹To construct the dataset, [Alviarez \[2019\]](#) combined data from the OECD, the Eurostat Foreign Affiliate Statistics database, the Bureau of Economic Analysis, and Bureau van Dijk's Orbis dataset.

TABLE 1: PARAMETERS AND TARGET MOMENTS

PARAMETER	DESCRIPTION	VALUE/TARGET MOMENT
<i>Set Externally</i>		
σ	Demand elasticity	Set to 4 [Arkolakis et al., 2018]
$\frac{\sigma-1}{\varepsilon-1}$	Interdependencies among locations	Set from 0.25 to 3.9 with benchmark values $\frac{3}{2}$ and $\frac{2}{3}$
<i>Calibrated Internally</i>		
ζ	Firm Pareto shape	Sales right-tail [Arkolakis, 2010]
\underline{z}_i	Firm Pareto minimum	Share of global MP conducted by firms HQed in i [Alviarez, 2019]
A_ℓ	Location productivity	GDP per capita (PWT)
f_i	Fixed cost (origin)	One-year firm survival rate (OECD)
f_i^e	Entry cost	Total firm counts (UNIDO)
H_ℓ	Labor supply	Employment (PWT)
$\tau_{\ell n}, \gamma_{i\ell}, \nu_{i\ell}$	Bilateral trade, MP, fixed costs	Gravity coefs. on trade flows, MP flows, affiliate counts [Alviarez, 2019]

Note: This table summarizes the calibration strategy of the model. Two parameters, σ and ε , are set externally; the remaining are calibrated simultaneously in a single iterative routine that targets the moments described in the last column, subject to general equilibrium conditions. For each empirical moment, we list the data source. PWT refers to the Penn World Table, OECD to the OECD Structural and Demographic Business Statistics, and UNIDO to the UNIDO Industrial Statistics Database.

Demand Elasticity and Location Substitution Elasticity: σ and ε The ratio $(\sigma - 1)/(\varepsilon - 1)$ determines whether locations are complements or substitutes, as well as the strength of these forces. We set $\sigma = 4$, following Arkolakis et al. [2018].²² To cover a wide range of complementarity and substitutability, we calibrate the model for 37 values of $(\sigma - 1)/(\varepsilon - 1)$, evenly spaced between 0.25 and 3.9, which implicitly defines a range of values for ε . We choose two values of this ratio as the benchmarks for our quantitative exercises below, one for the complements case and one for the substitutes case. The substitutes benchmark is $(\sigma - 1)/(\varepsilon - 1) = 2/3$, corresponding to $\varepsilon = 5.5$ as in Arkolakis et al. [2018], while the complements benchmark is $(\sigma - 1)/(\varepsilon - 1) = 3/2$, chosen for symmetry.

²²Our choice of σ falls within the range of estimates of Broda and Weinstein [2006]. Arkolakis et al. [2018] show that it is also consistent with markup estimates from the manufacturing sector. Moreover, our value is similar to the estimate of $\sigma = 3.89$ from Head and Mayer [2019]. While Head and Mayer [2019] focuses on MP in the car industry, their estimation strategy is consistent with our theoretical setup and also generates markups in line with the microeconomic estimates from the car industry in Goldberg [1995] and Berry et al. [1995].

Country Productivity Parameters and Fixed Costs We assume that the firm productivity distribution $G_i(z)$ is Pareto with shape parameter ζ and minimum \underline{z}_i , which is specific to each headquarter country i . The model generates a firm sales distribution with a Pareto tail of $\zeta/(\sigma - 1)$, and we set $\zeta = 1.65 \times (\sigma - 1)$ to match the tail estimate in [Arkolakis \[2010\]](#).

The model features two vectors of country-specific productivity terms: the minimum of the firm Pareto distribution, \underline{z}_i , which acts as a headquarter-location productivity shifter; and the production-location productivity shifter, A_ℓ . We choose the first to match the share of global foreign production attributable to firms headquartered in i , and the second to match the observed GDP per capita for each country.

For the calibration, we decompose the fixed cost of setting up a production location in ℓ for firms headquartered in location i into an origin-specific shifter and a bilateral term, $f_{i\ell} \equiv f_i v_{i\ell}$. We choose the base component of fixed costs f_i to match the observed one-year firm survival rate and the entry cost f_i^e to match the mass of entering firms. In the data, we compute the total mass of entrants as a country's total firm count divided by its one-year firm survival rate. We set total labor supply in each location, H_ℓ , to match total employment in the country.

Trade Costs, MP Costs, Fixed Costs In the theory, three bilateral cost matrices shape the patterns of trade flows, foreign affiliate sales, and foreign affiliate counts across country pairs: the matrix of trade costs $\{\tau_{\ell n}\}_{\ell n}$, the matrix of MP costs $\{\gamma_{i\ell}\}_{i\ell}$, and the matrix of the bilateral components of fixed costs $\{v_{i\ell}\}_{i\ell}$. Following [Tintelnot \[2017\]](#), we parameterize these matrices as constant elasticity functions of the standard bilateral gravity variables: geographic distance, a shared colonial past indicator, a shared language indicator, and a common border indicator. Each matrix has a separate elasticity for every gravity variable, totaling twelve elasticities. We additionally include bilateral tariffs in the trade cost matrix, with an elasticity of one. We present the expressions for all three bilateral cost matrices in the Online Appendix.

We identify the elasticities using an indirect inference approach, since the model does not deliver closed-form aggregate gravity equations. Using trade flows, MP flows, and bilateral affiliate counts from the model, we estimate separate gravity equations for each and choose the elasticities to match the corresponding regression coefficients in the data.²³

²³For any bilateral pair without MP activity in the data, we set $\gamma_{i\ell} = \infty$.

4. Quantitative Exercises and Counterfactuals

This section applies our solution method to the calibrated model. We first demonstrate the method’s computational advantages over incumbent approaches, then examine how the distribution of welfare effects from MP depends on whether production locations act as complements or substitutes. For each exercise, we use either our benchmark complements and substitutes calibrations, or the full range of calibrated economies differing in the degree to which locations are complements or substitutes.

4.1. Computational Performance

We conduct three numerical experiments, demonstrating the advantages of our solution method in terms of speed, precision, and breadth of applicability. For experiments that vary the number of countries, we sample from the location-specific parameters of our calibrated model to generate synthetic countries.²⁴

Speed We compare our policy function method (“Policy”), which combines generalized squeezing with the iterative cutoff search refinement, against two alternative approaches that discretize firm productivity on a grid of $2^{14} \approx 16000$ points. The “Naive” approach solves the CDCP at each grid point using the brute-force approach of computing the profits for all possible location combinations. The “Squeezing” approach applies the squeezing procedure and then applies the brute-force approach to the reduced domain. Both alternatives use a productivity grid and therefore produce interpolation error in aggregation. Our policy function method is exact and hence avoids a grid entirely.

Table 2 reports the time (in seconds) required to compute the policy function across varying numbers of synthetic countries. We compute $\mathcal{L}_i^*(\cdot)$ for firms from each origin i and report the average time across origins. With just 16 countries, the naive method requires more than an hour, while the policy method solves the same problem in less than a tenth of a second. Even with 128 countries, the policy method recovers the average policy function in under 15 seconds, and with 256 countries in under 10 minutes.

²⁴A synthetic country is described by three types of parameters: fundamentals $\{A_\ell, f_i, H_\ell\}$, aggregates $\{w_\ell, P_\ell\}$, and bilateral costs $\{\tau_{\ell n}, \gamma_{i\ell}, v_{i\ell}\}$. For each, we non-parametrically fit a distribution to the estimated values and sample from it. We draw 256 synthetic countries this way, then order them by their market size, $w_\ell H_\ell$. In Table 2 and Figure 4, each numerical exercise subsets from this collection of synthetic countries, from the largest to smallest.

TABLE 2: RUNTIMES (SECONDS) FOR DIFFERENT SOLUTION METHODS

Countries	Complements			Substitutes		
	Naive (1)	Squeezing (2)	Policy (3)	Naive (1)	Squeezing (2)	Policy (3)
8	6	0.482	0.034	11	0.400	0.009
16	5215	2.699	0.087	8444	2.142	0.018
32	–	16.968	0.186	–	10.735	0.118
64	–	116	1.293	–	76.540	1.324
128	–	895	14.702	–	874	14.167
256	–	7375	374	–	4797	565
Grid points	2^{14}	2^{14}	–	2^{14}	2^{14}	–

Note: This table reports the time (in seconds) to compute the firm policy function $\mathcal{L}_i^*(\cdot)$ using three methods. The “Naive” method discretizes firm productivity on a grid and evaluates the profit function for all possible production location combinations at each grid point. The “Squeezing” method applies the squeezing procedure at each grid point until convergence, then uses the brute-force approach on the reduced domain. Our “Policy” method solves for the exact policy function without discretization. Results are shown for both the complements and substitutes calibrations. Synthetic countries are generated by sampling from the distribution of parameter estimates obtained from the calibrated model. We report the average time to solve for $\mathcal{L}_i^*(\cdot)$ across all origins i . All computations were performed on an Apple M1 (2020) CPU.

In the complements case, the squeezing method corresponds to the incumbent approach pioneered by [Antràs et al. \[2017\]](#), and our policy function method is one to two orders of magnitude faster across all country counts. In the substitutes case, the naive brute-force approach represents the incumbent approach, and our method improves on it by four to five orders of magnitude. Approximately half of the speed gains come from extending the squeezing method of [Jia \[2008\]](#) to the substitutes case, and the remaining half from introducing our policy function method, which avoids solving the CDCP at every grid point.

Table 2 does not account for the discretization error of the alternative approaches; we quantify the size of this error next.

Precision Our policy function method solves for the *exact* policy function, allowing us to quantify the discretization error of other methods, which arises because decision sets must be interpolated between grid points.

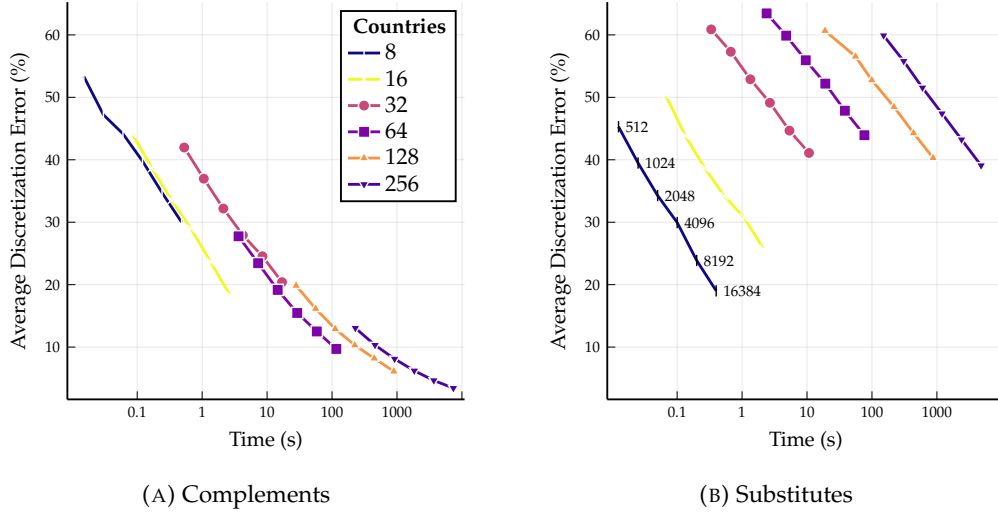


FIGURE 4: THE PRECISION-TIME FRONTIER WITH DISCRETIZATION METHODS
Note: This figure shows the trade-off between computation time and discretization error for the “Squeezing” method, for varying numbers of countries. Discretization error is measured as the average percentage deviation of trilateral flows computed using the discretized policy function from those computed using the exact policy function obtained by the “Policy” method. Specifically, we compute $N^{-3} \sum_{i,\ell,n \in L} |\tilde{X}_{i\ell n} / X_{i\ell n} - 1| \times 100\%$, where $X_{i\ell n}$ denotes sales to destination n by production sites in ℓ whose headquarters are in i , computed using the exact policy function, and $\tilde{X}_{i\ell n}$ is the corresponding object computed using the “Squeezing” method with interpolation between grid points. All computations were performed on an Apple M1 (2020) CPU.

To quantify the trade-off between computational time and discretization error inherent in grid-based approaches, we first compute the total value of trade flows from origin i to destination n via production in country ℓ , denoted $X_{i\ell n}$, by aggregating over the exact policy function. We then compute the corresponding approximation $\tilde{X}_{i\ell n}$ by aggregating over the discretized policy function of the “Squeezing” approach, interpolating between grid points. We measure discretization error as the average absolute percentage deviation between flows computed from the discretized policy function and the exact policy function, that is, $N^{-3} \sum_{i,\ell,n \in L} |\tilde{X}_{i\ell n} / X_{i\ell n} - 1| \times 100\%$.²⁵

Discretization error is large even with many grid points in both the complements and substitutes case, regardless of the number of countries. Figure 4 graphs the discretization error against computational time for varying grid densities in both the complements and substitutes cases. Each line shows the precision-time frontier

²⁵Since the support of productivity is unbounded, the only flows $X_{i\ell n}$ with zero value are those where $\gamma_{i\ell} = \infty$. In this case, the approximation $\tilde{X}_{i\ell n}$ is also zero, so we set the error to 0%.

for a different number of countries. With 512 grid points and 8 countries, computation is fast, but the average discretization error exceeds 40% regardless of whether locations are complements or substitutes. Even with over 16000 grid points, the error remains near or above 20% in both cases. Doubling the number of grid points reduces the error by 5 to 10 percentage points but doubles the computation time.

Discretization error decreases with the number of countries in the complements case but shows no systematic pattern in the substitutes case, a difference that may partly reflect the different structure of the policy function in each. In the complements case, the policy function has a nesting structure: more productive firms produce in all locations chosen by less productive firms and possibly more.²⁶ Grid-based methods thus induce less error since they only miss marginal locations added between grid points, and this error tends to shrink as the number of countries grows. In the substitutes case, the policy function has no nesting structure: more productive firms may operate entirely different location sets, so grid-based methods may miss policy function switches between grid points entirely. Unlike the complements case, adding more countries does not systematically reduce this error.

Wide Applicability In a final exercise, we examine how computational time varies across the full range of our calibrations indexed by the degree of complementarity or substitutability among production locations. For each calibration, the left panel of Figure 5 shows the average time required to solve for the policy function across origin countries, using our policy function method. The figure also separates the computational time spent on squeezing versus the iterative cutoff search refinement. We mark the two benchmark calibrations with diamond symbols on the x-axis.

Across all degrees of complementarity and substitutability, our policy function method takes under 0.1 seconds on average to solve for the policy function. Computation is fastest near the vertical line where production locations are independent. As locations become more complementary or substitutable, computation time increases.

²⁶Milgrom and Shannon [1994] show that the policy function exhibits a nesting structure when the objective function satisfies their single crossing and quasi-supermodularity conditions (the Online Appendix specifies the relationship between their conditions and SCD-C and SCD-T). Both conditions hold in our application, so that higher-productivity firms choose all locations selected by lower-productivity firms and possibly more: $z < z'$ implies $\mathcal{L}^*(z) \subseteq \mathcal{L}^*(z')$. The nesting result has been used in applied theory and empirical work, including by Antràs et al. [2017] in their analysis of firm-level sourcing decisions, and is reminiscent of the positive assortative patterns studied in Costinot [2009].

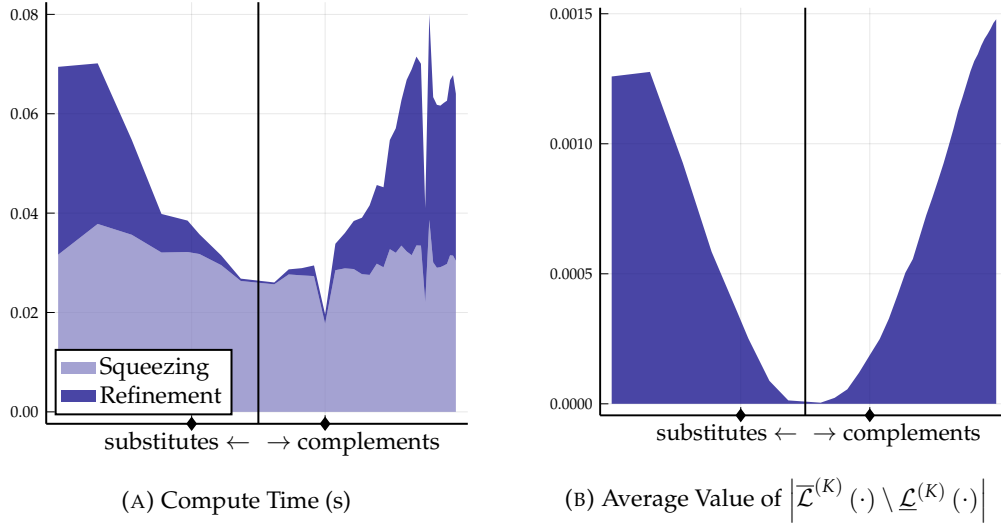


FIGURE 5: DEGREE OF COMPLEMENTARITY OR SUBSTITUTABILITY AND COMPUTATIONAL PERFORMANCE

Note: This figure shows how computation time and the effectiveness of the policy function method vary with the degree of complementarity or substitutability across locations as measured by $\frac{\sigma-1}{\varepsilon-1} / \left(1 + \frac{\sigma-1}{\varepsilon-1}\right)$. The x-axis varies ε while holding $\sigma = 4$ fixed. The vertical line marks the case where locations are neither complements nor substitutes, with $\frac{\sigma-1}{\varepsilon-1} = 1$. We indicate the two benchmark calibrations with diamond symbols on the x-axis. The left panel decomposes the total computation time (in seconds) spent on generalized squeezing relative to the iterative cutoff search refinement. The right panel shows the average number of locations separating the upper and lower bounding set across all firms and countries after convergence of the generalized squeezing procedure. All computations were performed on an Apple M1 (2020) CPU.

Theorem 2 helps explain the patterns in the left panel. It establishes that the generalized squeezing method never takes more than $|L|$ applications, in line with the relatively flat squeezing time across all parameter combinations.

The right panel shows a proxy measure for the effectiveness of the squeezing procedure: the average number of locations separating the upper and lower bounding sets across all firms and countries after convergence of the generalized squeezing procedure. Few locations separate the bounding sets after convergence, though this number increases with the strength of complementarity or substitutability, consistent with the widening gap between squeezing and total time in the left panel.

4.2. Welfare Gains from Multinational Production

We examine how the welfare gains from MP vary with the degree of complementarity or substitutability among locations by comparing each calibrated equilibrium to its corresponding MP-autarky counterfactual in which firms are restricted to produce only in their headquarter country. Formally, the MP-autarky counterfactuals set bilateral MP costs to infinity, so that $\gamma_{i\ell} = \infty$ for all $i \neq \ell$.

Figure 6 reports the resulting welfare gains from MP across the full range of complementarity and substitutability. The light gray band shows the full cross-country range of effects and the dark gray band the interquartile range; the figure also highlights the welfare effects in selected countries. Two patterns stand out. First, at a given level of complementarity or substitutability, welfare gains from MP are uneven across countries. Second, this cross-country dispersion widens as locations become more complementary.

For a given degree of complementarity or substitutability, two main factors shape the differences in welfare gains across countries. The first is the cost of conducting MP, which differs across countries in the calibration: moving from MP autarky to the calibrated equilibrium thus represents a larger reduction in MP costs for some countries than for others. The Netherlands, for example, has its costs of doing MP lowered by more than most countries and therefore experiences larger welfare gains. The second is productivity. Countries can be productive both as headquarter locations—because of high headquarter productivity—and as production locations—because of low productivity-adjusted wages, low fixed costs, and good market access. Countries with these advantages benefit more from MP through the foreign expansion of their firms or by attracting foreign affiliates.

Moving from substitutability to complementarity in Figure 6, the distribution of welfare gains across countries fans out. As cross-location complementarity rises, welfare gains increase most in countries that serve as headquarter locations for highly productive firms, particularly those with low costs of conducting MP, while gains in economies without productive headquarters are dampened and can even turn negative. The welfare gains in the United States, for example, triple moving from strong substitutability to strong complementarity, whereas Lithuania and Romania experience losses with strong complementarity. With substitutability, by contrast, gains are more evenly distributed across countries.

This pattern arises because complementarity strengthens multinational scale economies, amplifying the marginal cost advantages of the most productive firms. As scale economies intensify, the most productive firms can leverage their affiliates to achieve even lower marginal costs and expand sales relative to less productive firms, so welfare gains concentrate in the countries where those firms are headquartered.

Countries that are less productive as headquarters and more productive as production locations see welfare losses, as foreign multinational set up local affiliates that crowd out domestic headquarters. When locations are substitutes, by contrast, production in one location cannibalizes production in others, limiting these scale effects and distributing gains more evenly across countries.

This concentration effect is reflected in the difference between the GDP-weighted average welfare gain across countries and the median welfare gain. The GDP-weighted average welfare gain is largest when complementarity is strongest, because scale economies at the most productive firms raise global output the most. The median welfare gain, however, is largest when substitutability is strongest, because gains are distributed more evenly and a broader set of countries benefits from MP.²⁷

To further understand the sources of variation in welfare gains across countries, we decompose welfare gains following [Arkolakis et al. \[2012\]](#), adapted to our setting.²⁸ The welfare change in country i can be written as the sum of an openness term, a variety term, and an average productivity term:

$$(9) \quad \ln \frac{\hat{w}_i}{\hat{P}_i} = \underbrace{\ln \hat{\phi}_{iii}^{-\frac{1}{\sigma-1}}}_{\text{openness}} + \underbrace{\ln \hat{z}_i^{-\frac{\zeta}{\sigma-1}} + \ln \hat{M}_i^{\frac{1}{\sigma-1}}}_{\text{variety}} + \underbrace{\ln \hat{z}_i + \ln \left[\sum_{\mathcal{Z}_{i,t} \in \mathcal{T}_i} \lambda_{iii,t} s_{iii,t}^{1-\frac{\sigma-1}{\varepsilon-1}} \right]^{-\frac{1}{\sigma-1}}}_{\text{average productivity}}.$$

In this expression, $\hat{x} \equiv x/x'$ denotes the ratio of a variable in the baseline equilibrium to its value under MP autarky. The term ϕ_{iii} is the share of expenditure in country i on goods produced domestically by firms headquartered in i ; it serves as an inverse measure of openness to trade and MP. The mass of entrants M_i and the productivity cutoff \tilde{z}_i capture entry and selection effects in country i .

The policy function $\mathcal{L}_i^*(\cdot)$ partitions firms into productivity intervals $\{\mathcal{Z}_{i,t}\}_t$, within which firms choose the same set of production locations. For firms in interval t , the term $s_{i\ell i,t}$ denotes the share of their domestic sales produced in country ℓ , so $s_{i\ell i,t}$ sums to 1 across production locations ℓ . The term $\lambda_{iii,t}$ denotes the share of total domestic sales, produced domestically, that are accounted for by firms in that

²⁷The largest GDP-weighted average effect is 7.64%, occurring under the strongest complementarity, while the largest median welfare effect is 6.70%, occurring under the strongest substitutability.

²⁸We derive this welfare formula in Appendix B under the assumption that, in the initial equilibrium, all active firms include the headquarter location in their optimal set of production locations. This assumption holds in all numerical exercises.

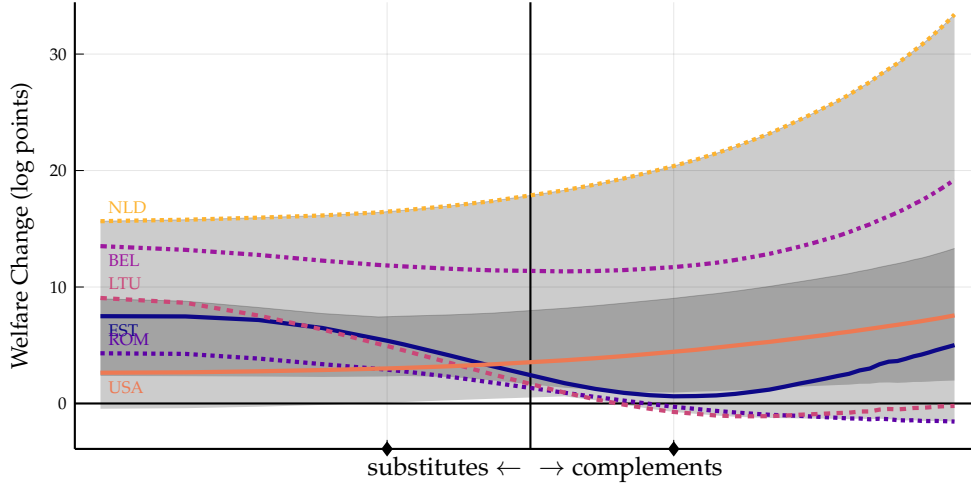


FIGURE 6: WELFARE GAINS FROM MULTINATIONAL PRODUCTION

Note: This figure shows the log point welfare change ($100 \times \ln(\hat{w}_i/\hat{P}_i)$) from moving from MP autarky to the calibrated economy, across model calibrations with different degrees of complementarity or substitutability among locations as measured by the ratio $\frac{\sigma-1}{\varepsilon-1} / \left(1 + \frac{\sigma-1}{\varepsilon-1}\right)$. In this expression, $\hat{x} \equiv x/x'$ denotes the ratio of a variable in the baseline equilibrium to its value in MP autarky. The x-axis varies ε while holding $\sigma = 4$ fixed (its value in the baseline calibration). The vertical line indicates the case where locations are neither complements nor substitutes, with $\frac{\sigma-1}{\varepsilon-1} = 1$. The light gray band depicts the full range of welfare changes across all 32 countries in our calibration, while the dark gray band shows the interquartile range. The lines trace the welfare impact for select countries across the different degrees of complementarity or substitutability.

interval. The shares $\lambda_{iii,t}$ sum to 1 across intervals t . Together, these terms capture how MP reshapes the distribution of sales across firms and locations.

Figure 7 shows the welfare gains from MP for every country in our sample in pink, and their decomposition into the openness, variety, and average productivity channels in shades of blue. For readability, we show the welfare effects only for our benchmark degrees of complementarity and substitutability.

The openness channel is positive for all countries and reflects access to cheaper foreign production. Countries spend less on locally produced goods by domestic firms as they open to MP, effectively shifting their production possibility frontier outward. This effect is largest for two types of countries: economies that import large volumes and hence benefit from a global decrease in marginal costs through MP, and low-productivity economies that benefit from new, productive foreign affiliates producing locally, giving them access to foreign varieties at lower prices.

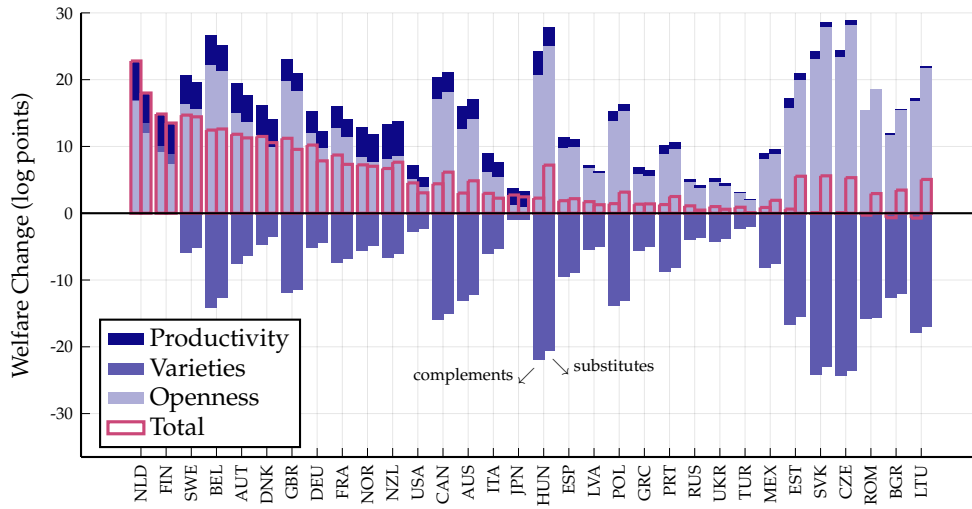


FIGURE 7: DECOMPOSING THE WELFARE GAINS FROM MULTINATIONAL PRODUCTION

Note: This figure shows the log point welfare change, $100 \times \ln(\hat{w}_i / \hat{P}_i)$, from moving from MP autarky to the calibrated economy as pink outlines. In addition, the figure decomposes the welfare changes into the contributions from changes in openness, the number of available varieties, and average productivity from equation (9). The left bars are for the complements calibration, the right bars for the substitutes calibration. The countries are ordered by the size of the total welfare effect in the complements calibration.

The variety channel reflects changes in the mass of domestic varieties. With complementarity, productive firms expand MP across locations, raising wages and survival cutoffs in host countries. Countries that mostly serve as hosts to the affiliates of foreign firms see their own domestic entrants crowded out, generating a negative variety effect. By contrast, in the Netherlands and Finland, countries where many multinationals are headquartered, global expansion raises profits, entry, and the mass of available local varieties. With substitutability, both effects are muted: multinationals operate fewer locations, earn lower profits, and crowd out fewer domestic firms in host countries.

The productivity channel is positive and captures changes in sales-weighted average productivity. In countries where multinationals are headquartered, the domestic market share of this class of productive firms expands relative to firms only producing domestically; in countries with few multinationals, this reallocation effect is small.²⁹ With complementarity, scale economies amplify this expansion and thus

²⁹At the extreme, if no domestically-headquartered firm engages in MP, the distribution of domestic

the variation across countries of the productivity effect. Substitutability dampens the effect and its variation, since the competition among locations prevents multinational firms from expanding domestic production as much.

Together, these channels explain the variation in welfare gains from MP across countries and calibrations. Complementarity magnifies the advantages of highly productive and open economies and intensifies reallocation toward globally dominant firms. Substitutability limits this concentration, so gains are more even across countries. Thus, the strength and direction of interdependencies among production locations shape not only the aggregate gains from MP, but also their distribution across firms and countries.³⁰

5. Conclusion

This paper studies combinatorial discrete choice problems, in which agents select items whose values depend on which others are chosen. Without structure on the objective function, finding the optimal decision set requires evaluating the payoffs of all possible combinations, a number that grows exponentially in the number of items, limiting the scope of quantitative models with such decision problems. We develop an exact solution method for such problems when the objective function satisfies a single crossing condition, which imparts structure on how each item's value changes as the chosen set expands. The method recovers the global optimum without exhaustive evaluation and extends to heterogeneous agents.

We apply the method to a general equilibrium model of multinational production in which heterogeneous firms choose international production locations subject to fixed costs. When locations are complements, scale economies concentrate production in the most productive firms, amplifying the welfare gains from multinational production in their headquarter countries while leading to small gains or even losses for countries that primarily host foreign affiliates. On the other hand, substitutability among locations limits this concentrating effect, so that gains are positive and more even across countries.

firms can be described by a representative firm as in the closed economy of Melitz [2003]. In this case, as a consequence of CES demand together with Pareto productivity, changes in MP costs cause no reallocations among firms, so that the productivity effect is zero. Competition for local labor instead lowers entry.

³⁰In the Online Appendix, we examine how the welfare losses from MP autarky depend on the presence of fixed costs. Our benchmark calibration features fixed costs; without them, all firms produce in all countries and the CDCP disappears. Removing fixed costs lowers the concentration effect of MP. When we recalibrate the model without fixed costs to match the same data, the average welfare gains from MP are larger in both the complements and substitutes cases, in line with Tintelnot [2017].

Beyond multinational production, combinatorial discrete choice problems arise in a range of economic settings, such as sourcing decisions, network formation, technology adoption, and store location choices, yet are often simplified to preserve tractability at the cost of abstracting from economically relevant interdependencies. Our approach allows researchers to model such interdependencies and rich heterogeneity without sacrificing exactness or computational feasibility. Future research drawing on our tools could estimate the sign and strength of cross-choice interactions from micro data, analyze policies that alter complementarities or fixed costs, and quantify models with variable markups that respond to complementarity-driven scale economies.

References

- Alonso Alfaro-Ureña, Juanma Castro-Vincenzi, Sebastián Fanelli, and Eduardo Morales. Firm export dynamics in interdependent markets. Technical report, National Bureau of Economic Research, 2024.
- Vanessa Alviarez. Multinational production and comparative advantage. *Journal of International Economics*, 119:1–54, 2019.
- Pol Antràs, Teresa C. Fort, and Felix Tintelnot. The margins of global sourcing: Theory and evidence from us firms. *American Economic Review*, 107(9):2514–64, 2017.
- Pol Antràs, Evgenii Fadeev, Teresa C. Fort, and Felix Tintelnot. Exporting, global sourcing, and multinational activity: Theory and evidence from the united states. *The Review of Economics and Statistics*, pages 1–48, 04 2024a. ISSN 0034-6535. doi: 10.1162/rest_a_01450.
- Pol Antràs, Evgenii Fadeev, Teresa C. Fort, and Felix Tintelnot. Export-platform fdi: Cannibalization or complementarity? *AEA Papers and Proceedings*, 114:130–35, May 2024b. doi: 10.1257/pandp.20241021.
- Costas Arkolakis. Market penetration costs and the new consumers margin in international trade. *Journal of Political Economy*, 118(6):1151–1199, 2010.
- Costas Arkolakis, Arnaud Costinot, and Andrés Rodríguez-Clare. New trade models, same old gains? *American Economic Review*, 102(1):94–130, 2012.
- Costas Arkolakis, Natalia Ramondo, Andrés Rodríguez-Clare, and Stephen Yeaple. Innovation and production in the global economy. *American Economic Review*, 108(8):2128–73, 2018.

- Costas Arkolakis, Arnaud Costinot, Dave Donaldson, and Andrés Rodríguez-Clare. The elusive pro-competitive effects of trade. *The Review of Economic Studies*, 86(1):46–80, 01 2019. ISSN 0034-6527. doi: 10.1093/restud/rdx075. URL <https://doi.org/10.1093/restud/rdx075>.
- Costas Arkolakis, Fabian Eckert, and Rowan Shi. Combinatorial discrete choice: A quantitative model of multinational location decisions. Technical report, National Bureau of Economic Research, 2023.
- Kristian Behrens, Giordano Mion, Yasusada Murata, and Jens Suedekum. Quantifying the Gap Between Equilibrium and Optimum under Monopolistic Competition. *The Quarterly Journal of Economics*, 135(4):2299–2360, 2020.
- Steven T. Berry, James A. Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.
- Endre Boros and Peter L. Hammer. Pseudo-boolean optimization. *Discrete Applied Mathematics*, 123(1):155–225, 2002. ISSN 0166-218X. doi: [https://doi.org/10.1016/S0166-218X\(01\)00341-9](https://doi.org/10.1016/S0166-218X(01)00341-9).
- Christian Broda and David E. Weinstein. Globalization and the gains from variety. *The Quarterly Journal of Economics*, 121(2):541–585, 2006.
- Maddalena Conte, Pierre Cotterlaz, Thierry Mayer, et al. *The CEPII Gravity Database*. 2023.
- Arnaud Costinot. An elementary theory of comparative advantage. *Econometrica*, 77(4):1165–1192, 2009. doi: <https://doi.org/10.3982/ECTA7636>.
- Sebastian Dyrda, Guangbin Hong, and Joseph B. Steinberg. A macroeconomic perspective on taxing multinational enterprises. *Journal of International Economics*, 152:104022, 2024. ISSN 0022-1996. doi: <https://doi.org/10.1016/j.jinteco.2024.104022>.
- Jonathan Eaton and Samuel Kortum. Technology, geography, and trade. *Econometrica*, 70(5):1741–1779, 2002.
- Pablo D Fajgelbaum, Eduardo Morales, Juan Carlos Suárez Serrato, and Owen Zidar. State taxes and spatial misallocation. *The Review of Economic Studies*, 86(1):333–376, 2019.
- Ying Fan and Chenyu Yang. Competition, product proliferation, and welfare: A study of the us smartphone market. *American Economic Journal: Microeconomics*, 12(2):99–134, 2020.
- Robert C. Feenstra, Robert Inklaar, and Marcel P. Timmer. The next generation of the penn world table. *American Economic Review*, 105(10):3150–3182, 2015.

- Matthew Gentzkow. Valuing new goods in a model with complementarity: Online newspapers. *American Economic Review*, 97(3):713–744, June 2007. doi: 10.1257/aer.97.3.713.
- Pinelopi Koujianou Goldberg. Product differentiation and oligopoly in international markets: The case of the u.s. automobile industry. *Econometrica*, 63(4): 891–951, 1995.
- Keith Head and Thierry Mayer. Brands in motion: How frictions shape multinational production. *American Economic Review*, 109(9):3073–3124, 2019.
- Keith Head, Thierry Mayer, Marc Melitz, and Chenying Yang. Industrial policies for multi-stage production: The battle for battery-powered vehicles. Working Paper 34884, National Bureau of Economic Research, February 2026. URL <http://www.nber.org/papers/w34884>.
- Igal Hendel. Estimating multiple-discrete choice models: An application to computerization returns. *The Review of Economic Studies*, 66(2):423–446, 1999.
- Jonas Hjort, Hannes Malmberg, and Todd Schoellman. The missing middle managers: Labor costs, firm structure, and development. Technical report, National Bureau of Economic Research, 2022.
- Panle Jia. What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica*, 76(6):1263–1316, 2008.
- Stephen C. Kleene. General recursive functions of natural numbers. *Mathematische Annalen*, 112:727–742, 12 1936.
- Stephen C. Kleene. On notation for ordinal numbers. *Journal of Symbolic Logic*, 3(4): 150–155, 1938.
- Peter J Klenow and Jonathan L Willis. Real rigidities and nominal price changes. *Economica*, 83(331):443–472, 2016.
- Jiří Klimeš. Fixed edge theorems for complete lattices. *Archivum Mathematicum*, 17(4):227–234, 1981.
- Adrian Kulesza. Combinatorial discrete choice with deep reinforcement learning, 2024.
- Nelson Lind and Natalia Ramondo. Trade with correlation. *American Economic Review*, 113(2):317–53, February 2023.
- Daniel McFadden. Conditional logit analysis of qualitative choice behavior. *Frontiers in Econometrics*, 1974.

- Marc J. Melitz. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, 71(6):1695–1725, 2003.
- Paul Milgrom. *Putting Auction Theory to Work*. Cambridge University Press, 2004.
- Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica*, 62(1):157–180, 1994.
- Marc-Andreas Muendler and Sascha O. Becker. Margins of multinational labor substitution. *American Economic Review*, 100(5):1999–2030, 2010.
- Robert A Pollak. Additive utility functions and linear engel curves. *The Review of Economic Studies*, 38(4):401–414, 1971.
- Natalia Ramondo. A quantitative approach to multinational production. *Journal of International Economics*, 93(1):108–122, 2014.
- Natalia Ramondo and Andrés Rodríguez-Clare. Trade, multinational production, and the gains from openness. *Journal of Political Economy*, 121(2):273–322, 2013.
- Natalia Ramondo, Andrés Rodríguez-Clare, and Felix Tintelnot. Multinational production: Data and stylized facts. *American Economic Review*, 105(5):530–536, 2015.
- JMC Santos Silva and Silvana Tenreyro. The log of gravity. *The Review of Economics and Statistics*, 88(4):641–658, 2006.
- Ina Simonovska. Income differences and prices of tradables: Insights from an online retailer. *The Review of Economic Studies*, 82(4):1612–1656, 2015.
- Sebastian Sotelo. Practical aspects of implementing the multinomial pml estimator. *Ann Arbor: University of Michigan, mimeo*, 2019.
- V. Stoltenberg-Hansen, I. Lindström, and E. R. Griffor. *Mathematical Theory of Domains*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1994.
- Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955.
- Feodora A. Teti. Missing tariffs. Technical Report 11590, CESifo Working Papers, 2024. Feodora Teti’s Global Tariff Database (v_beta1-2024-12).
- Felix Tintelnot. Global production with export platforms. *The Quarterly Journal of Economics*, 132(1):157–209, 2017.
- Kenneth E. Train, Daniel McFadden, and Moshe Ben-Akiva. The demand for local telephone service: A fully discrete model of residential calling patterns and service choices. *The RAND Journal of Economics*, 18(1), 1987.

A. Appendix A: Methods for Policy Function Refinement

In this section, we introduce two methods of refining the policy function after the generalized squeezing procedure has converged. Generalized squeezing produces two reduced bounding functions that partition the type space into intervals, on each of which the lower and upper bounding functions take constant values $[\underline{\mathcal{L}}_t, \overline{\mathcal{L}}_t]$. Where the bounding functions coincide, the optimal decision set equals their common value. Where they do not, further refinement is required. We focus on solving for the policy function over one such interval \mathcal{Z}_t with $\underline{\mathcal{L}}_t \subset \overline{\mathcal{L}}_t$; the full policy function is recovered by applying a refinement method to each such interval in turn.

The first refinement method, generalized branching, requires SCD-C and SCD-T. The second, iterative cutoff search, requires the stronger condition of strong SCD-T and is the method we use in our quantitative application. The Online Appendix contains all proofs.

A.1. The Generalized Branching Procedure

We extend the branching logic from Section 1 to heterogeneous agent settings. On an interval \mathcal{Z}_t where $\underline{\mathcal{L}}_t \subset \overline{\mathcal{L}}_t$, the generalized branching step selects one location $\ell \in \overline{\mathcal{L}}_t \setminus \underline{\mathcal{L}}_t$ and creates two *branches*: one that imposes ℓ is included in the decision sets and the other that imposes ℓ is excluded.

Each branch corresponds to the problem of solving for the policy function that maps each type $z \in \mathcal{Z}_t$ to the conditionally optimal decision set \mathcal{L} maximizing the modified objective function $\tilde{f} : \mathcal{P}((\overline{\mathcal{L}}_t \setminus \underline{\mathcal{L}}_t) \setminus \{\ell\}) \times \mathcal{Z}_t \rightarrow \mathbb{R}$ where:

$$\begin{aligned} \tilde{f}(\mathcal{L}, z) &\equiv f(\mathcal{L} \cup \underline{\mathcal{L}}_t \cup \{\ell\}, z) && \text{on the branch that includes } \ell, \\ \tilde{f}(\mathcal{L}, z) &\equiv f(\mathcal{L} \cup \underline{\mathcal{L}}_t, z) && \text{on the branch that excludes } \ell. \end{aligned}$$

The generalized branching step applies the generalized squeezing procedure to each restricted CDCP, yielding a pair of reduced bounding functions for each branch. It then translates these bounding functions for the restricted CDCPs into bounding functions for the original CDCP by adding all items in $\underline{\mathcal{L}}_t$ to both bounding functions on each branch. On the branch that includes ℓ , the item ℓ is also added to both bounding functions. The output of the generalized branching step is thus a pair of conditional bounding functions for the original CDCP on each branch which each induce a new partition on \mathcal{Z}_t .

On type intervals in this new partition for which the bounding functions on a branch coincide, the branch terminates, having identified a conditionally optimal




FIGURE 8: THE GENERALIZED BRANCHING PROCEDURE

Note: The figure illustrates an example outcome of the generalized branching procedure, in an example where $\underline{\mathcal{L}}_t = \emptyset$ and $\bar{\mathcal{L}}_t = \{\text{USA}, \text{CAN}\}$. The branching step selects $\ell = \text{USA}$ and creates two branches: one which presumes USA is in the optimal set and the other which presumes the opposite. Convergence on a single branch occurs when the generalized squeezing procedure returns a conditionally optimal policy function. The final output of the full recursive generalized branching procedure is the collection of all conditionally optimal policy functions.

policy function; on other intervals, the generalized branching step is applied recursively to each sub-interval where the bounding functions differ. We refer to the recursive application of the generalized branching step until all branches terminate on all intervals as the “generalized branching procedure.” Because each step strictly reduces the number of locations separating the upper and lower bounding functions, the procedure branches a finite number of times.

To obtain the global policy function on \mathcal{Z}_t , we combine the results across all terminal branches. Each terminal branch yields a conditionally optimal policy function, which is piecewise constant but may change value at different productivity types across branches. We intersect these policy functions to partition \mathcal{Z}_t into subintervals of types on which none of the conditionally optimal policy functions change value. On each such subinterval, we collect the candidate decision sets from the conditionally optimal policy functions. For each type $z \in \mathcal{Z}_t$, the global policy function chooses the decision set that yields the highest value among these candidate decision sets.

Figure 8 illustrates this refinement on an interval \mathcal{Z}_t where the converged bounding functions are $\underline{\mathcal{L}}_t = \emptyset$ and $\bar{\mathcal{L}}_t = \{\text{USA}, \text{CAN}\}$. Branching on the USA location creates two branches: one that includes the USA in the optimal set, in pink, and one that excludes it, in blue. In the example, applying generalized squeezing on

each branch produces a conditionally optimal policy function that maps each type to an optimal decision set, shown in the two middle lines. The bottom line collects the candidate optimal sets for each element of the type space partition. The optimal policy function is found by choosing, for each type $z \in \mathcal{Z}_t$, the profit-maximizing decision set among the conditionally optimal decision sets.

A.2. Iterative Cutoff Search

Iterative cutoff search is an alternative refinement method that relies on the additional structure that a stronger version of SCD-T induces on the policy function.

Definition (Strong SCD-T). For any pair of decision sets $(\mathcal{L}, \mathcal{L}')$ and types $z < z'$,

$$f(\mathcal{L}, z) - f(\mathcal{L}', z) \geq 0 \quad \Rightarrow \quad f(\mathcal{L}, z') - f(\mathcal{L}', z') > 0$$

Strong SCD-T requires that, for any pair of decision sets, there is at most one type that receives the same payoff from the two sets. When the objective function satisfies strong SCD-T, the policy function does not change value and then revert to a previous value as the type increases: if two types share the same optimal decision set, all intermediate types share it as well. The policy function is thus a sequence of sub-intervals, each with a distinct optimal decision set.

Consider an interval \mathcal{Z}_t from generalized squeezing with end points $z < z'$. If the optimal decision sets on the two end points coincide, so that $\mathcal{L}^*(z) = \mathcal{L}^*(z')$, then this decision set must also be optimal for all types within the interval.

If the optimal decision sets on the two end points do not coincide, we apply iterative cutoff search. It pins down the policy function two adjacent sub-intervals at a time by narrowing the interval until there are only two unique values of the policy function. It then finds the cutoff type where the optimal decision set switches, using the fact that both decision sets must be optimal and deliver the same value at the cutoff.

Given an interval where $\mathcal{L}^*(z) \neq \mathcal{L}^*(z')$, strong SCD-T guarantees a unique type $\bar{z} \in [z, z']$ where the two optimal decision sets of the end points have the same value. If either of the sets is also optimal at \bar{z} , then the policy function must switch from $\mathcal{L}^*(z)$ to $\mathcal{L}^*(z')$ at \bar{z} . The interval is therefore resolved into two sub-intervals, each with constant optimal decision sets.³¹

If instead $\mathcal{L}^*(\bar{z})$ differs from both end point sets, then there are at least three sub-intervals within $[z, z']$, since three distinct values of the policy function have been

³¹As we discuss in Appendix B, the model of our quantitative application admits a closed-form expression for the indifferent type.

identified. Iterative cutoff search narrows its search to $[z, \bar{z}]$ and iterates again. Once it succeeds in finding two adjacent sub-intervals of the policy function, as described, it returns to iterate on the remainder of the interval.

Definition (Iterative cutoff search). Consider the interval $[z^{(0)}, z'^{(0)}]$. Iterate as follows.

1. Identify the type $\bar{z}^{(k)}$ that is indifferent between $\mathcal{L}^*(z^{(k)})$ and $\mathcal{L}^*(z'^{(k)})$. Solve for the optimal decision set at this type, $\mathcal{L}^*(\bar{z}^{(k)})$, over the reduced domain implied by $[\underline{\mathcal{L}}_t, \bar{\mathcal{L}}_t]$.
2. Update as follows.
 - a) If $\mathcal{L}^*(\bar{z}^{(k)})$ coincides with either $\mathcal{L}^*(z^{(k)})$ or $\mathcal{L}^*(z'^{(k)})$, then update the policy function

$$\mathcal{L}^*(z) = \begin{cases} \mathcal{L}^*(z^{(k)}) & \text{for } z \in [z^{(k)}, \bar{z}^{(k)}] \\ \mathcal{L}^*(z'^{(k)}) & \text{for } z \in [\bar{z}^{(k)}, z'^{(k)}] \end{cases}.$$

If $z'^{(k)} = z'^{(0)}$, the policy function has been solved for the full interval. Terminate iteration. Otherwise, set $z^{(k+1)} = z'^{(k)}$ and $z'^{(k+1)} = z'^{(0)}$.

- b) Otherwise, set $z^{(k+1)} = z^{(k)}$ and $z'^{(k+1)} = \bar{z}^{(k)}$.

Computing $\mathcal{L}^*(\bar{z})$ at the candidate cutoff types requires solving a CDCP. We implement this step using the squeezing and branching methods from the main paper. Thus, while SCD-C is not required for iterative cutoff search, we use it to compute the optimal decision sets in practice.

B. Appendix B: Model Derivations

In this section, we establish the SCD-C and SCD-T conditions in the firm problem of Section 2 and derive the welfare expression in equation (9). In the Online Appendix, we establish the single crossing differences conditions in a broader class of frameworks.

B.1. Establishing SCD-C and SCD-T

Let $\mathcal{L} \subset \mathcal{L}'$. The marginal value of location ℓ to the set \mathcal{L}' in terms of its marginal value to \mathcal{L} is:

$$\begin{aligned} D_\ell \pi_i(\mathcal{L}', z) &= D_\ell \pi_i(\mathcal{L}, z) \\ &+ \left(\frac{\sigma - 1}{\varepsilon - 1} - 1 \right) \frac{\sigma - 1}{\varepsilon - 1} \sum_{n \in \mathcal{L}} \frac{1}{\sigma} X_n \left(\frac{\sigma - 1}{\sigma} P_n z \right)^{\sigma - 1} \\ &\times \int_0^{\xi_{in}^{1-\varepsilon}} \int_{\Theta_{in}(\mathcal{L} \setminus \{\ell\})}^{\Theta_{in}(\mathcal{L}' \setminus \{\ell\})} [t + u]^{\frac{\sigma-1}{\varepsilon-1} - 2} du dt \end{aligned}$$

where we denote $\Theta_{in}(\mathcal{L}) \equiv \sum_{\ell \in \mathcal{L}} \xi_{i\ell n}^{1-\varepsilon}$. Since \mathcal{L}' nests \mathcal{L} , $\Theta_{in}(\mathcal{L}') \geq \Theta_{in}(\mathcal{L})$ for all markets n and the integrals are positive.

As the marginal value shows, the profit function satisfies the SCD-C condition. When the ratio $(\sigma - 1)/(\varepsilon - 1)$ is above 1, the marginal value of location ℓ in \mathcal{L}' is positive whenever its marginal value in \mathcal{L} is positive. In this case, the profit function satisfies SCD-C from below. On the other hand, when the ratio is below 1, the marginal value of ℓ is negative in \mathcal{L}' whenever it is negative in \mathcal{L} , so that the profit function satisfies SCD-C from above.

Now, let $\{\mathcal{L}, \mathcal{L}'\}$ be any pair of production location sets. The difference in profits associated with each of these decision sets can be written:

$$\begin{aligned} \pi_i(\mathcal{L}', z) - \pi_i(\mathcal{L}, z) &= z^{\sigma-1} \sum_{n \in \mathcal{L}} \frac{1}{\sigma} X_n \left(\frac{\sigma - 1}{\sigma} P_n \right)^{\sigma-1} \left[\Theta_{in}(\mathcal{L}')^{\frac{\sigma-1}{\varepsilon-1}} - \Theta_{in}(\mathcal{L})^{\frac{\sigma-1}{\varepsilon-1}} \right] \\ &- \sum_{\ell \in \mathcal{L}} w_\ell f_{i\ell} (\mathbb{1}[\ell \in \mathcal{L}'] - \mathbb{1}[\ell \in \mathcal{L}]) \end{aligned}$$

Since $\sigma > 1$, this difference is strictly monotonic in the productivity z , so there is at most one productivity at which the firm is indifferent between the two production location sets.³² Then, the profit function satisfies strong SCD-T, and, in turn, SCD-T. The type indifferent between the two production sets can be expressed in closed form.

B.2. Counterfactual Welfare

The share of expenditure in n that is captured by goods produced in ℓ by firms originating from i , denoted $\phi_{i\ell n}$, integrates over the behavior of all firms. It simplifies

³²Strong SCD-T is violated only in the knife-edge case where two decision sets exist such that the total variable profit improvement (summed across markets n) is zero *and* the difference in total fixed cost payments (summed across locations ℓ) is also zero.

to:

$$\begin{aligned} \phi_{i\ell n} &= M_i \left(\frac{\sigma-1}{\sigma} P_n \right)^{\sigma-1} \\ &\times \sum_{t=1}^T \mathbb{1}_{i\ell,t} \bar{\zeta}_{i\ell n}^{1-\varepsilon} \Theta_{in}(\mathcal{L}_{i,t})^{\frac{\sigma-1}{\varepsilon-1}-1} \left[(z_{i,t+1})^{\sigma-1-\zeta} - (z_{i,t})^{\sigma-1-\zeta} \right] \bar{\zeta} \end{aligned}$$

where $\bar{\zeta}$ is a constant of integration. We split the productivity range of active firms into the intervals along which the policy function is constant, indexed by $t \in \{1, \dots, T\}$, denoting with $z_{i,t}$ the left end point of each interval and setting $z_{i,T+1} = \infty$. Within each interval, let $\mathcal{L}_{i,t}$ be the optimal set and $\mathbb{1}_{i\ell,t} \equiv \mathbb{1}[\ell \in \mathcal{L}_{i,t}]$ indicate whether location ℓ is in the optimal set.

In what follows, we assume that fundamentals are such that all active firms establish production in their country of origin i , so that $\mathbb{1}_{ii,t} = 1$ for all t . This assumption holds in all our calibrations.

For MP autarky, we set $\gamma'_{i\ell} = \infty$ for all $i \neq \ell$, so that active firms only produce domestically. Then, there is only one interval, $[\bar{z}'_i, \infty)$, and the optimal location set along this interval is $\{i\}$. The counterfactual change in ϕ_{iii} simplifies to:

$$\frac{\phi'_{iii}}{\phi_{iii}} = \left(\frac{M'_i}{M_i} \right) \left(\frac{P'_i}{P_i} \right)^{\sigma-1} \left(\frac{\bar{\zeta}'_{iii}}{\bar{\zeta}_{iii}} \right)^{1-\sigma} \left(\frac{\bar{z}'_i}{\bar{z}_i} \right)^{\sigma-1-\zeta} \sum_t \lambda_{iii,t} s_{iii,t}^{1-\frac{\sigma-1}{\varepsilon-1}},$$

where

$$\begin{aligned} s_{i\ell i,t} &\equiv \frac{\bar{\zeta}_{i\ell i}^{1-\varepsilon}}{\sum_{\ell' \in \mathcal{L}_{i,t}} \bar{\zeta}_{i\ell' i}^{1-\varepsilon}}, \\ \lambda_{iii,t} &\equiv \frac{s_{iii,t} \Theta_{ii}(\mathcal{L}_{i,t})^{\frac{\sigma-1}{\varepsilon-1}} \left[(z_{i,t})^{\sigma-1-\zeta} - (z_{i,t+1})^{\sigma-1-\zeta} \right]}{\sum_{t'=1}^T s_{iii,t'} \Theta_{ii}(\mathcal{L}_{i,t'})^{\frac{\sigma-1}{\varepsilon-1}} \left[(z_{i,t'})^{\sigma-1-\zeta} - (z_{i,t'+1})^{\sigma-1-\zeta} \right]}. \end{aligned}$$

Using $\bar{\zeta}'_{iii}/\bar{\zeta}_{iii} = w'_i/w_i$, we arrive at equation (9).

C. Appendix C: Computational Implementation of Our Solution Methods

We implement all solution methods necessary for the quantitative exercises in the Julia package accompanying this paper, available at <https://github.com/rowanxshi/CDCP.jl>. The package contains the squeezing and branching procedures to solve

the CDCP at a specific firm productivity. It also contains the generalized squeezing procedure and iterative cutoff search to solve for the policy function across a range of firm productivity values. Our replication package, which calls the Julia solution package, implements all quantitative exercises with the multinational model. This section discusses how our methods enter the quantitative exercises of Sections 3 and 4.

C.1. Calibration of GE Model

The first application of our solution methods is to calibrate the model in general equilibrium. Calibration requires solving for the policy functions for each origin country, $\{\mathcal{L}_i^*(\cdot)\}_i$, at every new parameter guess. We thus embed the policy function method within the calibration routine.

As Table 1 lays out, we first set the parameters $\{\sigma, \varepsilon, \zeta\}$. Our calibration finds model parameters such that the data values of $\{w_\ell, M_\ell, H_\ell\}_\ell$ are an equilibrium.

We start with an initial guess for the market size aggregate $\{P_n^{(0)}\}_n$, the fundamentals $\{A_\ell^{(0)}, \underline{z}_\ell^{(0)}, f_\ell^{(0)}\}_\ell$, and the bilateral cost parameterization $\{\bar{\tau}_\ell^{(0)}, \bar{\gamma}_\ell^{(0)}, \bar{v}_\ell^{(0)}\}_\ell$ with $\{\kappa_\tau^{v(0)}, \kappa_\gamma^{v(0)}, \kappa_v^{v(0)}\}_v$. We then iterate as follows.

1. Solve the policy functions $\{\mathcal{L}_i^*(\cdot)\}_i$ of the firm's CDCP from equation (2) using the policy function method. The policy function also determines \tilde{z}_i as the lowest type to operate at least one location, satisfying equation (3).
2. Given $\{\mathcal{L}_i^*(\cdot)\}_i$, compute trade, MP, and affiliate flows in the model; estimate the PPML gravity equation for each (see the Online Appendix for details).
3. Use deviations from aggregate conditions and moments to update the guess:

- update $\{P_n^{(k)}\}_n$ using equation (5);
- update $\{A_\ell^{(k)}, \underline{z}_\ell^{(k)}, f_\ell^{(k)}\}_\ell$ using respectively deviations from equation (6), the gap between model and data share of total foreign MP that is attributable to firms headquartered in each country, and the gap between model and data survival rate in each country;
- update $\{\bar{\tau}_\ell^{(k)}, \bar{\gamma}_\ell^{(k)}, \bar{v}_\ell^{(k)}\}_\ell$ using the gap between model and data own shares of trade (domestic sales), MP (domestic production), and affiliates (domestic production locations); and

- update $\{\kappa_\tau^{v(k)}, \kappa_\gamma^{v(k)}, \kappa_v^{v(k)}\}_v$ using the gap between the model and data coefficients of the PPML regression.

The iteration converges to (near) zero on all deviations. The routine thus calibrates the model while enforcing GE conditions at the same time. After convergence, we invert equation (4) to recover f_i^e .

C.2. Time Trials and GE Counterfactuals

Section 4 begins with a subsection benchmarking the computational performance of our methods using the calibrated models. In Table 2 and Figure 4, we compare the squeezing method, discretized on a grid over the productivity range, against the policy function method. To do so, we use each approach to solve the CDCP across the full range of firm productivity values, holding fixed all model parameters, fundamentals, and aggregates. In Figure 5, for each value of the ratio $(\sigma - 1)/(\varepsilon - 1)$, we calibrate the model as described above and apply the policy function method. None of these figures requires computing the general equilibrium of the model.

To compute the welfare effects from multinational production requires solving the model in general equilibrium. To do so, given a set of parameters and fundamentals, we use an iterative routine that solves for the policy functions, $\{\mathcal{L}_i^*(\cdot)\}_i$, at each new guess for aggregates.

We start with an initial guess for the aggregates $\{P_\ell^{(0)}, w_\ell^{(0)}, M_\ell^{(0)}\}_{\ell'}$, then iterate as follows.

1. Given aggregates $\{P_\ell^{(k)}, w_\ell^{(k)}, M_\ell^{(k)}\}_{\ell'}$, set $X_\ell = w_\ell H_\ell$ for all ℓ to satisfy equation (7).
2. Solve the policy functions $\{\mathcal{L}_i^*(\cdot)\}_i$ of the firm's CDCP from equation (2) using the policy function method. The policy function also determines \tilde{z}_i as the lowest type to operate at least one location, satisfying equation (3).
3. Given the policy functions $\{\mathcal{L}_i^*(\cdot)\}_i$, use deviations from aggregate conditions to update the guess:
 - update $\{P_\ell^{(k)}\}_\ell$ using equation (5);
 - update $\{w_\ell^{(k)}\}_\ell$ using equation (6); and
 - update $\{M_\ell^{(k)}\}_\ell$ using equation (4).

In the counterfactual of Section 4, solving for the policy function becomes trivial. Firms open no foreign locations, so they compare $\mathcal{L} = \emptyset$, which represents being inactive, against $\mathcal{L} = \{i\}$, which represents producing domestically. Solving for the policy function then simplifies to identifying the cutoffs $\{\bar{z}_i\}_i$, where the fixed cost of setting up the domestic location perfectly offsets the variable profits that the firm would earn.