Industrial policies for multi-stage production: The battle for battery-powered vehicles

Keith Head[1], Thierry Mayer[2], Marc Melitz[3], and Chenying Yang[4]

[1]University of British Columbia
[2]Sciences Po
[3]Harvard University
[4]Singapore Management University

April 13, 2024

Preliminary and Incomplete Draft.

Abstract

Many countries have set ambitious targets for transitioning away from fossil fuels. The plans generally involve switching from combustion engines to electric vehicles (EVs). As batteries constitute around 40% of the cost of EVs, firms need to establish low-cost battery supply chains in order to make EVs attractive to consumers. At the same time, governments increasingly use tax and subsidy schemes to induce firms to locate more stages of the supply chain within their jurisdictions. We specify a multi-stage supply chain for EVs from battery cell production to vehicle distribution. Each car producer selects where to open facilities at each stage considering production costs, transport costs, tariffs and subsidies. This is a difficult combinatorial choice problem, but we leverage a mixed integer linear program formulation which can be solved in under a minute. We estimate the parameters of our model—which include the variable production costs and fixed plant/model activation costs—using observed sourcing decisions for all production stages over the period 2015 to 2022. The next step is a set of counterfactuals that compute how policy interventions affect the final pattern of production and trade in this sector. Ultimately, we plan to use the model to quantify the impact of competing industrial policies on global CO2 emissions.

1 Introduction

The rise of global supply chains has dramatically changed the landscape of the international organization of production. Firms slice up their production process, retaining only a subset of the stages within their domestic economies. Yet, this international organization of production does not come with international cooperation in government policies, especially in growing industries, like electric vehicles (EV). As more countries accelerate their transition to green economies and announce their ban on internal combustion engine vehicle sales by 2035, countries compete to get EVs made in their countries by providing domestic subsidies.

The 2022 US Inflation Reduction Act (IRA) awards consumer and production subsidies that do not extend to European-made cars. European governments are considering measures to counter those policies and to stem the rising tide of Chinese imports. Canada, meanwhile has promised to match US IRA production subsidies for plants established in Canada. Deputy prime minister Freeland defended the Canadian government's decision to spend roughly $30 billion on incentives to induce Volkswagen and Stellantis to build factories in Ontario, saying "Our government is absolutely determined that Canada gets its fair share of those green jobs."

The value chain for EVs is economically interesting because of countervailing forces at work. First, there are China's cost advantages in all stages of the EV value chain from refining of minerals to final assembly. The IRA contains a provision that prevents subsidies from being applied to vehicles containing Chinese minerals and other components. Second, there are large transport costs because batteries are heavy, bulky once arranged in packs, and the EVs themselves are challenging to transport because of their weight and fire risk. Finally, there are the subsidies and protectionist rules designed to pull EV and battery production into the consuming countries. With large fixed costs for each new facility, these forces interact in complex ways to determine the equilibrium locations of production and the distribution of sales.

This paper studies how firms endogenously form supply chains over space and how these chains shape the spatial distribution of the EV industry. Firms decide where to build plants and from whom to source inputs at every stage along a supply chain. Characterizing such allocation of production stages to countries is a challenging computational problem for even moderate number of alternative locations. Even with just one stage of production, the facility location problem (as it is referred to in the operations research literature) is already an NP-hard problem. Loosely speaking, this means that there are no algorithms guaranteed to solve it in polynomial time, or put more simply the problem-solving time "explodes" as the number of locations increases. The difficulty is compounded under multi-stage production because the optimal location of a given stage is not only a function of the marginal cost and fixed cost at that production stage, but is also shaped by the proximity of that location to the desired locations of production of the upstream and downstream stages. One contribution of this paper is to adapt techniques from operations research (OR) to solve for the geography of global supply chains with fixed costs. We extend the multi-stage production cost minimization problem considered in OR to allow for endogenous demand and market entry.

We show in simulations that the way firms respond to government policies depends in complex ways on the geographic structure and parameter choices. Various types of industrial policy can be effective, ineffective, or even counterproductive. Therefore, one cannot make policy recommendations without bringing in detailed data on the costs of distance, borders, and trade agreements. Variable and fixed costs of production need to be estimated. We propose a methodology to do so that uses different decisions the firms make to extract the implied parameters.

Once the model is quantified through data and estimation, we perform counterfactual exercises to determine how various industry policies currently in use are predicted to affect the extent to which various countries participate in domestic, regional, or global supply chains. We also quantify the welfare consequences of these policies. Specifically, we examine the effects of the US Inflation Reduction Act (IRA) which subsidize the US-assembled EVs with batteries sourced from its trade partners. The conditions to qualify for the EV tax credit have raised concerns from major battery makers, including those in China, Japan, South Korea and the EU. In response, the European Union is also considering loosening its rules to allow governments to provide more subsidies for EV manufacturers, leading to a subsidy war that may potentially be wasteful. In addition, the restrictions on cheap input sources also increase the cost of US car manufacturers, leading to a potential welfare loss.

However, one may argue that the transformation of green technology and the global urgency of decarbonization make the more subsidies—on both sides of the Atlantic—the better. Analyzing the welfare effects of US IRA alone is not sufficient unless the policy is compared to an alternative with international cooperation. To answer this question, we plan on deriving the global optimal subsidy and predicting the optimal spatial distribution of the EV supply chains.

The paper is organized as follows. Section 2 positions this paper within the literature on multi-stage production and the role of fixed costs in location choices and sourcing along the value chain. Section 3 describes three key facts about the EV industry that in-form core assumptions in the model. Section 4 presents the model of multi-stage production market entry, and equilibrium. In section 5, we verify the computational feasibility and illustrate how policies work within a stripped down model with a single decision maker. Section 6 develops the estimation framework and presents the estimated and calibrated parameters needed to solve the model using real world data. Section 7 carries out counterfactual policy exercises. Section 8 concludes.

2 Literature

Table 1: Positioning our model in the literature

		number of production* stages	
		Single ($K = 1$)	Multiple ($K > 1$)
	yes	ARRY 2018	Antras & de Gortari 2020
Constant returns? (CRS)		Head & Mayer 2019	Tyazhelnikov 2022
	no	Tintelnot 2017	de Gortari 2020 note
		AFT 2017, AES 2023	AFFT 2024 ($K = 2$)
		Castro-Vincenzi (2022) †	This paper ($K = 3$)

Note: * excludes market entry (stage $K + 1$ here). † features capacity constraints, other non-CRS have fixed costs

Plant location with fixed costs is a hard combinatorial problem as the number of potential locations grows. Even with just a single production stage, brute force requires evaluating 2^N alternative solutions. Jia (2008) pioneered using supermodularity of the profit function to reduce computational difficulty. Antràs et al. (2017) pointed out that the global sourcing problem can be supermodular or submodular depending on the relative magnitudes of the key demand and supply parameters, σ and θ. With σ denoting the elasticity of substitution between varieties in the demand function, and θ denoting the Frechet shape parameter on the supply side, supermodularity obtains if and only if $\sigma > 1 + \theta$. This restriction holds in Antràs et al. (2017): $\sigma = 3.9$, $\theta = 1.8$. However, Arkolakis et al. (2023) find $\sigma = 4$, $\theta = 4.5$ implying submodularity. They extend the Jia (2008) method to exploit submodularity with a "generalized squeezing" algorithm.

Our structure requires a different approach because one mechanism for submodularity and two mechanisms pushing for supermodularity. The submodularity mechanism is that plants at the same stage k substitute for each other. The supermodularity mechanism is that active upstream or downstream production facilities at stages $k \leq K$ complement each other. Furthermore, activated distribution ($K + 1$) facilities increase the profitability of adding production plants. We do not think oarametric restrictions to ensure global sub- or super- exist for for this setup. One attractive feature of our approach is that it is not necessary to rely on them.

3 Facts

This section describes our data and presents some useful facts about the battery EV industry.

3.1 Data on batteries and electric vehicles

We use three main sources of data in this project.

1. The first one is a dataset on the value chain of batteries for electric vehicles produced by the consultancy firm IHS-Markit (now part of S&P). For each vehicle produced in a given assembly plant, the dataset provides the source of the cells, modules, pack with associated respective plants and suppliers (LG, CATL, Panasonic etc.). Additional details are provided regarding the shape of the cell (prismatic, pouch or cylinder), the detailed chemistry (various forms of two main types: NMC and LFP), and the battery capacity in KWh. This is production-based data (years 2015 to 2022), with total volumes of cars assembled associated battery elements, but no information on where the cars are sold.

2. The second source is the updated version of the sales data (also from IHS-Markit) used in Head and Mayer (2019). It provides bilateral volumes at the car model level for each of 75 markets, and providing information on the plant of assembly. It also provides information on the manufacturing firm at that assembly plant (which can differ from the brand name, i.e. GM vs Chevrolet). importantly, the car model information does not include fuel type: when a flow of Hyundai Konas appears in a dataset, it sums all fuel types, Gas, Diesel, Hybrid, Plug-in hybrid and BEVs in that case. This is not an issue for models that are only produced as BEVs (the Nissan Leaf for instance), but there is no way to identify those in this source.

3. The third source, still produced by IHS-Markit is called New Registration. This is a destination-based dataset available for 24 markets, where sales / registration are reported, with more detail on the car model (most importantly for us; fuel type), but no information on the assembly plant. We use this last source to identify i) models

that are only BEVs, which means that the allocation of production in data source 1) can be safely done with data source 2), ii) models that have mixed fuel types, for which we compute a share of BEVs in total sales over the markets in order to allocate production of the BEV version.

The combination of those 3 sources help us construct a novel type of micro-level GVC dataset where we can track where the final product is sold (and how much), where it is assembled, and where do the elements of its core value chain (the battery pack) are themselves produced. All plants along the value chain have been geo-coded with latitude and longitude. The data is available for the quasi universe of the nascent BEV industry (none of the major destinations or countries of production are missing), over the period 2015-2022. Frictions that will be used in estimation stages use the usual sources in this literature (CEPII gravity data for distance and RTAs), WITS for tariffs.

3.2 Main markets of BEVs

Table 2: The 15 top markets of BEVs in 2022

Rank	Country	# Models	# Firms	Sales (000')	EV share (%)
1	China	150	46	3875.8	16.1
2	United States	40	15	819.7	5.9
3	Germany	61	24	334.6	11.6
4	United Kingdom	49	16	215.1	11.3
5	Norway	59	22	131.2	64.0
6	South Korea	30	8	117.2	7.1
7	France	49	17	106.7	5.7
8	Canada	35	14	94.8	6.2
9	Sweden	55	21	73.4	22.7
10	Japan	33	12	71.8	1.7
11	Netherlands	57	22	51.6	13.8
12	Switzerland	46	16	35.8	14.0
13	Belgium	52	18	32.8	7.7
14	Italy	50	17	32.0	2.2
15	Australia	26	11	31.7	3.0
16	Rest Of World	23	10	305.2	4.0

Note: Rest of world row reports averages across 59 countries for numbers of models and EV share and the sum for Sales.

Table 2 shows the top 15 markets for BEVs in 2022. China is by far the largest market in all dimensions except the share of EVs. Relative to the US it has three times more active firms and almost four times as many models. The majority of the large ICE vehicle-making countries (US, France, Korea, Japan, Canada, and Italy) have EV shares under 10%. In the model it will be important to capture both market size in terms of car buyers and pro-EV policies (The three countries with the highest EV shares—Norway, Sweden, and China—all had early aggressive subsidies to induce buyers to switch to EVs).

3.3 Actors along the value chain

Figure 1: Battery cell production is highly concentrated

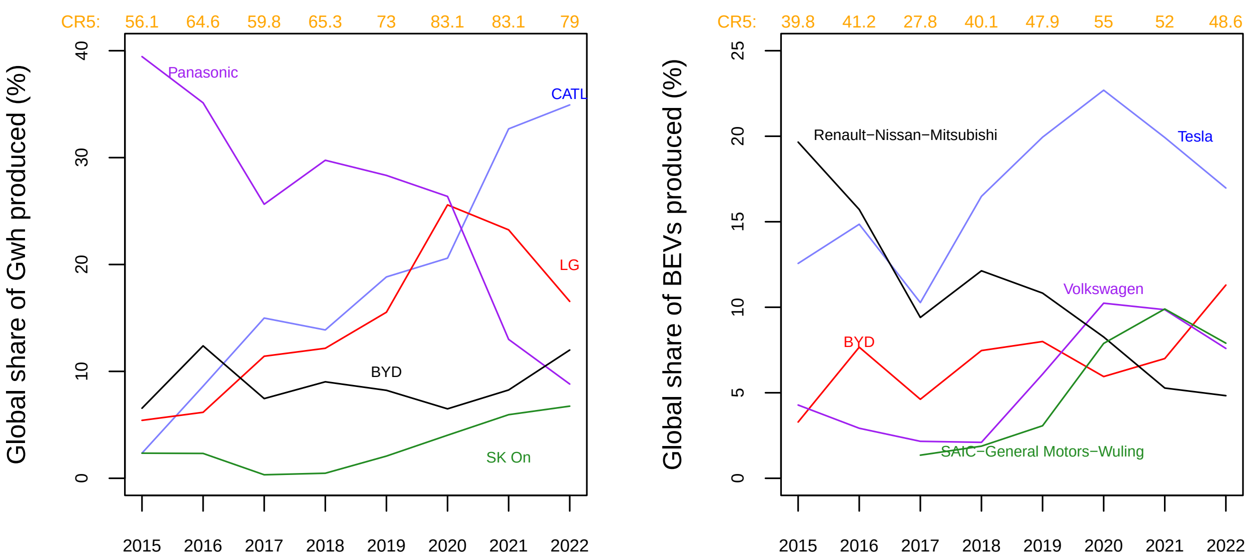

(a) Cells

(b) Electric Vehicles

Figure 1 shows that a handful of firms dominate the industry. The top five firms account for nearly 80% of all cells produced (measured in terms of gigawatt-hours to distinguish capacity differences). Vehicle assembly is less concentrated, but the top five still account for almost half of all EVs sold.

In our data there are 63 firms that assemble EVs in 2022. Of those, 14 sell exclusively in China. Table 3 shows that the top 13 firms—ranked by number of sales destinations—account for 96% of all EVs sold outside China. Because the Chinese market is so large, these firms account for much smaller share (62.5%) of world production. Even among the top firms we see that most firms sell in a minority of the markets.

Table 4 shows the major players at the first stage of production considered in this paper: cell manufacturing. CATL of China equips twice as many cars as its nearest rival, LG of Korea, but it focuses on a different shape of cell. It is also active in the less expensive

Table 3: 13 top firms, their 139 models in 2022

No.	Manufacturer	# Markets	# Models	Production Cum. Share (%)	Sales-exCHN Cum. Shr (%)
1	Tesla	26	4	17.2	40.2
2	Volkswagen	19	18	24.8	56.6
3	BMW	15	7	28.1	61.9
4	Hyundai	14	5	32.7	70.9
5	Geely	13	13	37.9	74.0
6	Mercedes-Benz	11	6	40.1	78.3
7	Ford	10	2	41.8	82.5
8	Nissan-Mitsubishi	9	4	43.3	86.9
9	BYD	8	10	54.8	87.6
10	Renault	8	6	58.2	91.4
11	SAIC	7	9	60.9	93.2
12	Toyota	7	4	61.4	94.1
13	General Motors	4	4	62.5	96.1

Table 4: Cell suppliers are specialized (2022 data)

Supplier	Vol (000s)	# clients	Shape			Material	
			Prismatic	Pouch	Cylinder	LFP	NMC
CATL	2142	47	99	1	0	42	58
LG Energy Solution	1074	15	0	70	30	0	100
BYD	964	6	100	0	0	99	1
Phylion	553	2	100	0	0	8	92
Panasonic	516	2	0	0	100	0	100
China Aviation Lithium	500	12	100	0	0	29	71
SK On	449	6	1	99	0	0	100
Gotion	321	16	53	19	27	91	9
CATL-SAIC	246	3	100	0	0	2	98
Samsung SDI	182	8	85	0	15	0	100

iron-based battery material called LFP. BYD, a vertically integrated EV maker from China is devoted almost entirely to prismatic LFP cells and its main client is its own downstream vehicle plants. The Japanese (Panasonic) and Korean (LG, SK On, Samsung) cell makers exclusively use the Nickel based NMC material. Individual car models use particular shapes and materials for their cells (no cars mix them) so these choices constrain the sourcing options for cell sourcing.

3.4 Location of plants and sourcing along the value chain

Figure 2: Cell plants in the major regions

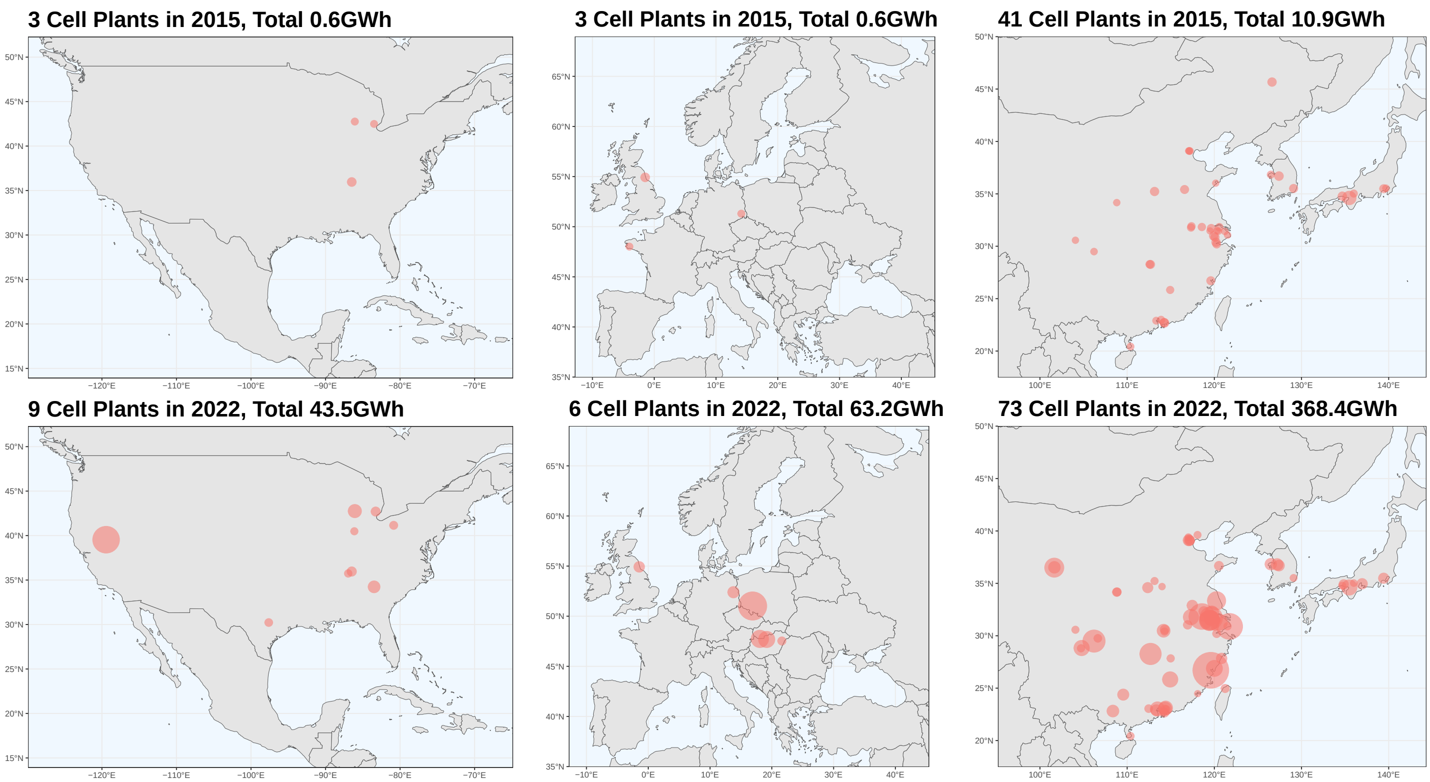

Figures 2, 3, and 4 show the impressive growth in the numbers of plants for cells, packs, and vehicles from 2015 to 2022—in all the main regions. The size of plant symbols corresponds to total output in gigawatt-hours for cells and packs and total sales for vehicles.

East Asia remains dominant in cell production but the number of plants in the US and Europe grow by factors of three and two respectively and the capacity in GWH rises by two orders of magnitude. While the growth of cells is impressive, it is also worth noting that many of countries in Figure 4 with multiple vehicle plants (Italy, Spain, Portugal, and Turkey) lack local cell production.

Figure 3: Pack plants in the major regions

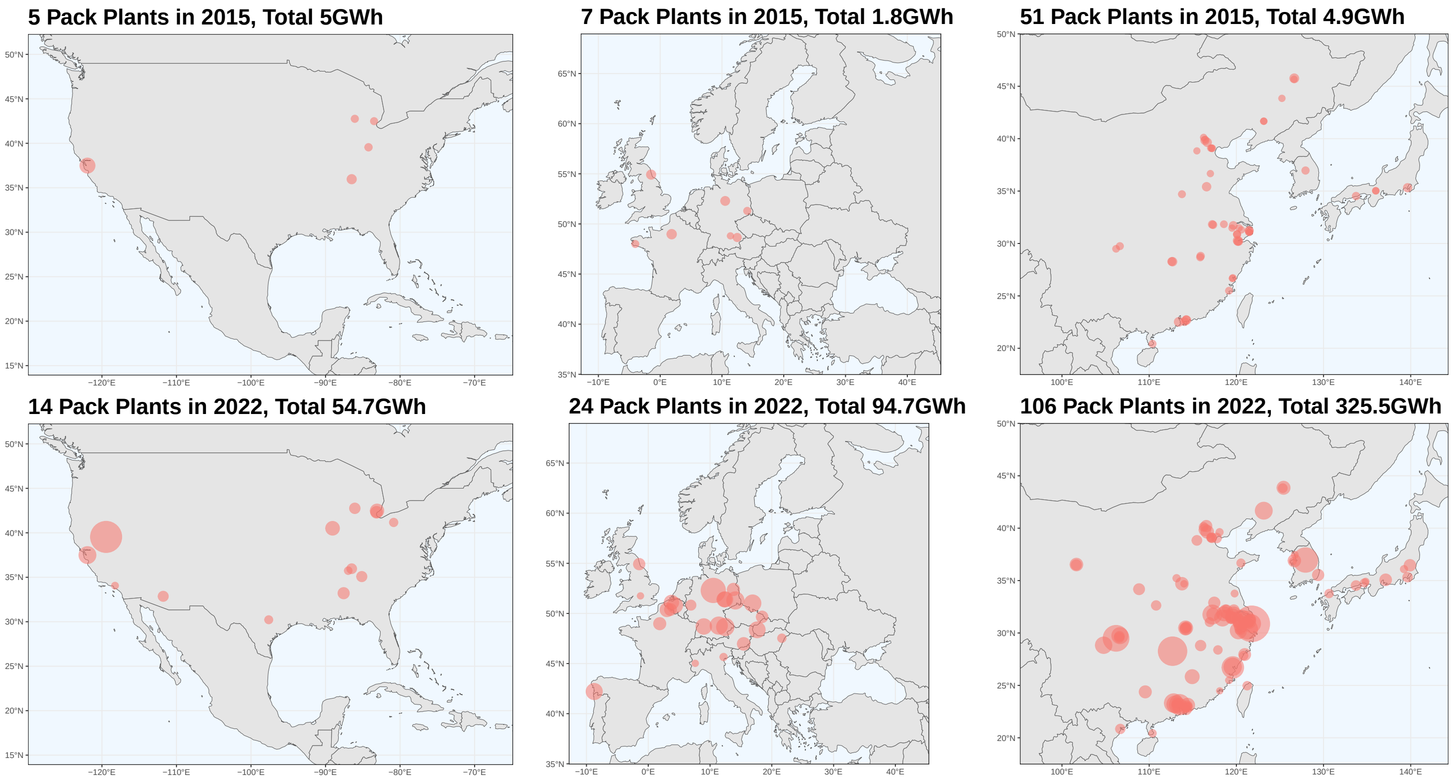

Figure 3 shows packs are intermediate: there are more pack plants than cell plants in every region in both years but fewer pack plants than vehicle plants. We will see that this "fanning out" is a natural feature of the multi-stage production model and does not rely on differences in either transport costs or fixed costs across stages.

The maps show the geographic dispersion of plants but they do not tell us which plants trade with which. Table 5 provides some insight into the distances that components travel. By showing medians and means we see the huge skewness in the distribution of distances that the output of a stage travels before becoming an input. For comparison purposes we also show distance for two of the most important parts of the ICE vehicle: the engine and the transmission. The finding that over 50% of modules travel 1km reflects our decision to code intra-plant distances as 1 to make them easy to recognize. Of the modules that travel long distances a very large share of them are put together in the Cell factory. Also many cells skip the module stage and go straight into packs. For this reason, we will omit modules from our representation of the value chain.

Table 6 shows that when we disaggregate as much as the data allow, a component in the value chain is almost always single-sourced. This is not the case at more aggregate levels of models and firms. The reason is that there are different types of battery com-

Figure 4: EV plants in the major regions

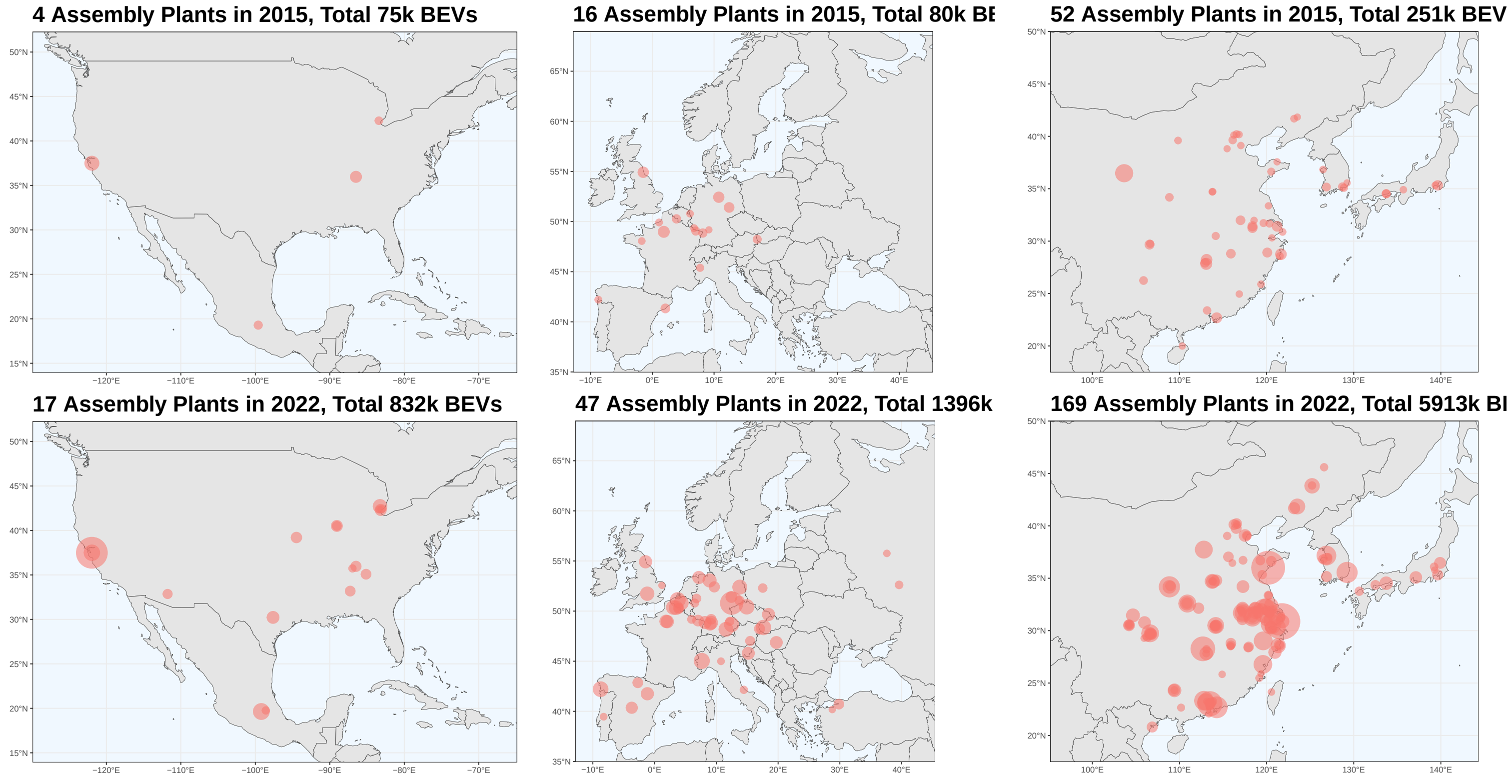

Table 5: How much do components travel?

Link	Year	Distance in km median	Distance in km mean
Battery Electric Vehicle			
Pack to Assembly	2015	299	819
	2022	215	641
Module to Pack	2015	1	994
	2022	1	830
Cell to Module	2015	13	1782
	2022	1	477
Internal Combustion Engine Vehicle			
Engine to Assembly	2018	133	1034
Transmission to Assembly	2018	681	2184

Table 6: Multi-sourcing is rare

Stage	Percent multi-sourcing by firms	models	specs
Cells	10	0.63	0.17
Modules	7.14	3.17	0.00
Packs	8.57	2.86	0.00
Battery chains	20	6.35	0.17
Assembly	17.46	2.44	2.44

Note: The numbers are percentages of models or firms that source from more than one country at each stage, conditional on the set of countries used in subsequent stages.

ponents and a purchaser may obtain one type of cell from one plant and a different type from another plant. Even at the final assembly stage we see that market countries tend to single source, with a narrowly defined vehicle only coming from two different plants 2.44% of the cases. In this respect, EVs do not differ notably from traditional vehicles. Head and Mayer (2019) find multi-sourcing (measured in terms of countries of origin rather than plants) for 2.3% of all model-market years from 2000 to 2018.

This fact is very important for thinking about how to formalize the firm's profit maximization problem. It is common in the trade literature that is applied to more aggregate data to think of each source country as supplying a distinct variety (Armington assumption) with firms using a CES aggregator to combine them. This is not appropriate at the detailed micro level. Instead, our setup realistically assumes that firms choose the lowest cost source for a narrowly defined component or final vehicle.[1]

There is another explanation for multi-sourcing which is also important to consider in setting up the firm's profit maximization problem: capacity constraints. We might see multi-sourcing not because of the love of variety common in trade models but simply because the preferred source is unable to meet the demand. However, the numbers in Table 6 suggest that capacity constraints are not a major concern. In line with the data patterns exhibited in the data, we employ an uncapacitated model in this paper.

[1]Thus, our component specifications are sufficiently detailed to correspond to the goods in the Eaton and Kortum (2002) continuum.

4 Multi-product Multi-stage Production

4.1 Demand

There are N consumer countries indexed $\{n : n \in N\}$. Each country has B_n buyers who select their preferred EV model m, among those offered in the market—or buy the outside good (internal combustion engine vehicle). Each household in market n has a common indirect utility $A_n + \xi_m - \alpha p_{nm}$, where A_n is the country-specific appeal of EVs (determined by *inter alia* charging stations, gas prices), α is the price sensitivity parameter, and p_{nm} is the price of model m in country n inclusive of any tariffs or consumer subsidy. The outside good is normalized to have indirect utility of 1. Household heterogeneity is distributed Type 1 Extreme Value, leading to the familiar expression for Logit demand after aggregating over households:

$$q_{nm} = \frac{\exp(A_n - \alpha p_{nm})}{1 + \sum_{m'} \exp(A_n - \alpha p_{nm'})} B_n. \tag{1}$$

Each firm f carries distinct varieties m. We assume the set of models for a firm M_{fn} and a measure of firms F_n depending on how many models and firms sell at destination location n.

4.2 The firm's problem

Figure 5 depicts a simplified version of our model to clarify the decisions we are modelling. As in the empirical model, we have $K = 3$, that is three stages of production. In addition, there is the decision to enter a consumer market. The schematic portrays a single firm that makes three models (identified by the colors Cyan, Magenta, and Orange). There are three potential locations: East, Central, and West.

The problem of the firm consists of three stages: At stage $K + 1$ the firm selects the set of car models to supply to each location n. In figure 5 Cyan is offered in West and Central; Magenta and Orange are only offered in Central. No model was profitable to offer in East. The firm simultaneously decides the locations of facilities along the value chain. In this example, only one location (Central) has an active cell plant; there are two Pack plants (Central is not activated at $k = 2$). Constrained by the active facilities, the firm chooses *paths* for each car model from cells to final consumers. We refer to this as an *assignment* problem. In our model once a firm activates a production, facility can be used by more than one model, as we see for E_2. The lines shown in gray are edge that are available but not selected. The decisions of which plants and markets to activate depend on quantities,

Figure 5: Schematic of a supply chain with $K = 3$

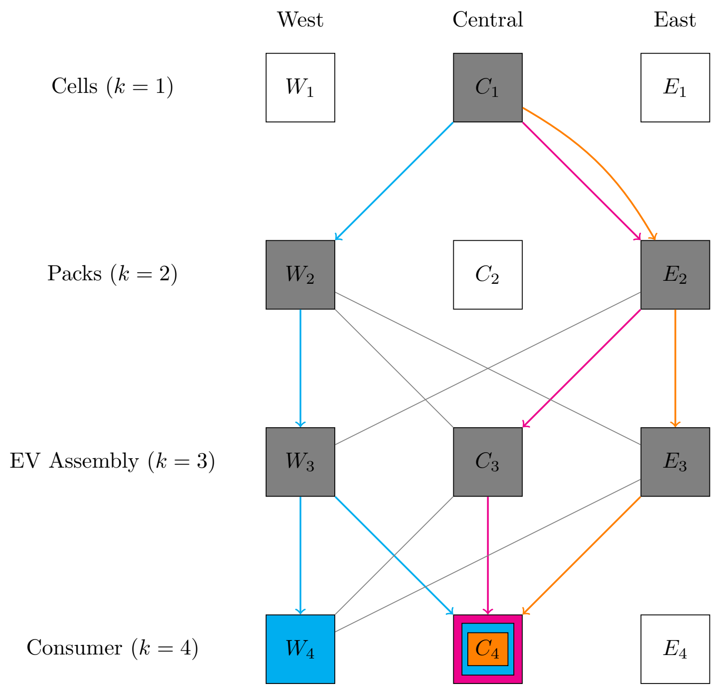

which are in turn determined (via a simple markup equation) by the sum of the marginal costs incurred at each node and along each edge.

A global value chain has levels of production $\{1, ..., K\}$, where level 1 is the furthest production stage from consumers and level K the closest stage. The set of potential locations where firms can open facilities is partitioned into different levels, $\{L_1, ..., L_K\}$, and each set L_k does not necessarily overlap with one another.[2] A particular production location at level k is denoted as ℓ_k. One can integrate the final consumption as level $K+1$ such that the location of final consumers $n \equiv \ell_{K+1} \in N = L_{K+1}$. We use these two interchangeably in the paper. Therefore, a path of value chain $\boldsymbol{\ell}_{nm} \equiv \{\ell_1(m), ..., \ell_K(m), \ell_{K+1}(m)\}$. A firm chooses sets of facilities to produce at each level and the set of destinations to sell, $\boldsymbol{\mathcal{L}}_f = \{\mathcal{L}_1(f), ..., \mathcal{L}_K(f), \mathcal{L}_{K+1}(f)\}$, and $\boldsymbol{\ell}_{nm} \in \boldsymbol{\mathcal{L}}_f$.

A firm is characterized by its country of origin, $h(f)$, the set of models it offers, M_f. Fixed costs incurred at the firm-stage are $\{\phi_{\ell k}(f)\}$ for $k = 1, ..., K$ Fixed cost of market access ϕ_{mn} are incurred for a single model in a market. They represent the cost of marketing, sales training, space for display, and so forth.

The production function at each level k is Leontief with γ_k upstream units needed to

[2]The data shows that there are 9 countries producing cells, 18 producing packs, 27 hosting EV assembly.

make one stage k unit along with β_k bundles of primary factors and other locally supplied inputs. For brevity we will refer to them as labor. Therefore, cost per unit on an edge from ℓ_k to ℓ_{k+1} are

$$c^{km}_{\ell_{k+1}\ell_k} = \beta_k(w_{\ell_k} + \tau_{\ell_{k+1}\ell_k} + u^m_{\ell_{k+1}\ell_k}), \tag{2}$$

where w_{ℓ_k} is the per unit labor cost at location ℓ_k, $\tau_{\ell_{k+1}\ell_k}$ are per unit trade cost including border and distance costs plus tariffs on the edge from ℓ_k to ℓ_{k+1}, and the $u^m_{\ell_{k+1}\ell_k}$ are the unobserved edge cost shocks. The $\beta_k \equiv \prod_{k+1}^{K} \gamma_k$, are the number of units required for the completed vehicle, where γ_k are input requirements at stage k.

Given the Leontief technology, the final cost of the assembled vehicle delivering to consumer n is

$$c(\boldsymbol{\ell}_{nm}) = \sum_{k=1}^{K} c^{km}_{\ell_{k+1}\ell_k} = \sum_{k=1}^{K} \beta_k(w_{\ell_k} + \tau_{\ell_{k+1}\ell_k} + u^m_{\ell_{k+1}\ell_k}). \tag{3}$$

The cost for a particular chain depends on three key aspects: (i) wage w_{ℓ_k} at locations of each level, (ii) the amount of cost saving associated with productivity for each model attributable to the bilateral locations $u^m_{\ell_{k+1}\ell_k}$, and (iii) the geography, as captured by a $N \times L_K$ matrix of trade cost $\tau_{n\ell_K}$ between consumers and assembly and a $L_{k+1} \times L_k$ matrix of trade cost $\tau_{\ell_{k+1}\ell_k}$ between two adjacent levels of production.

For a chosen path $\boldsymbol{\ell}_{nm}$, the firm also incurs fixed costs $\phi_{\ell k}$ in location ℓ for each stage of production k (presumably, a substantial portion of these are sunk once a facility is open). Note that these fixed costs are not indexed by model: Once a facility is open in ℓ for production stage k, we assume that it can accommodate any of the firm's models. We denote this facility location choice with a binary variable $y_{\ell k} \in \{0, 1\}$.

We start with an assumption of monopolistic competition, whereby the impact of a firm's prices across its models sold in market n on other firms' pricing decisions is negligible. Given the Logit demand (1), firms add a constant markup $1/\alpha$ to their delivered marginal cost $c(\boldsymbol{\ell}_{nm})$ for all their models. The resulting equilibrium quantity for each model then only depends on that model's delivered marginal cost $c(\boldsymbol{\ell}_{nm})$ and an aggregate cost index $\tilde{c}_n$:

$$q(c(\boldsymbol{\ell}_{nm}), \tilde{c}_n) = B_n \frac{\exp(A_n - 1)\exp[-\alpha c(\boldsymbol{\ell}_{nm})]}{1 + \exp(A_n - 1)\exp(-\alpha \tilde{c}_n)}, \qquad \tilde{c}_n \equiv -\frac{1}{\alpha} \ln\left\{\sum_m \exp[-\alpha c(\boldsymbol{\ell}_{nm})]\right\}. \tag{4}$$

In this case, the firm's variable profits earned from sales of model m to market n then also only depend on the delivered marginal cost and the cost index: $\pi(c(\boldsymbol{\ell}_{nm}), \tilde{c}_n) = q(c(\boldsymbol{\ell}_{nm}), \tilde{c}_n)/\alpha$.[3]

[3]We can extend our model to oligopolistic Bertrand pricing. In this case, the scalar cost index is replaced

Let ϕ_{nm} be the fixed market entry cost of selling model m in market n. The firm will then sell all models for which the variable profit $\pi(c(\boldsymbol{\ell}_{nm}), \tilde{c}_n)$ exceeds this cost. We denote those market entry choices with binary variables $y_{nm} \in \{0, 1\}$. Consider a firm's chosen set of paths $\boldsymbol{\ell}_{nm}$ for every model m and market n such that $y_{nm} = 1$ (this involves a set of choices for all the open facilities $y_{\ell k} = 1$ along those paths). The firm's total profits for that chosen set of paths is then:

$$\Pi = \sum_n \sum_m [\pi(c(\boldsymbol{\ell}_{nm}), \tilde{c}_n) - y_{nm}\phi_{nm}] - y_{\ell k} \sum_{k=1}^{K} \phi_{\ell_k} \tag{5}$$

In the following section, we describe how we solve for a firm's optimal set of paths $\boldsymbol{\ell}_{nm}$, including the determination of all the binary $y_{\ell k}$ facility location choices and y_{nm} market entry choices.

4.3 Multi-level uncapacitated facility location problem

We adapt the multi-level uncapacitated facility location problem (MUFLP) in the operations research literature described by Ortiz-Astorquiza et al. (2018) to our framework with multi-product and endogenous choice of market access. This results in a linear optimization problem with integer constraints. In a typical MUFLP, there are two types of decisions involved. One is the facility location decision $y_{\ell k}$ at each level of production. The other is the assignment of chosen paths across those locations and stages of production for a given product to a given destination. A chosen path $\boldsymbol{\ell} = \{\ell(1), \ell(2), \ldots, \ell(K)\}$ is represented by a set of indicator variables $x_{\boldsymbol{\ell} nm}$. If there are L_k potential location with 3 production stages $k = 1...K$, then there are a total of $L_1 \times ... \times L_K$ indicator variables $x_{\boldsymbol{\ell} nm}$ for a given product m and location n. Further multiplying by the number of products and number of locations yields the total number of x variables.[4]We further adapt this optimization problem to incorporate the market entry decisions y_{nm}. Maximizing total firm profits (5) can then be written as the following linear optimization problem with integer

by the entire vector of delivered marginal costs for all models sold in market n.

[4]The operations research literature refers to this MUFLP formulation as a path-based optimization problem. There is also an alternative edge-based formulation, which defines the x variables for "edges" from one stage k to stage $k+1$. We use the path-based formulation as this delivers a linear optimization problem in the x and y variables, even when the variable profit is non-linear in delivered marginal cost $c(\boldsymbol{\ell}_{nm})$ and cost index $\tilde{c}_n$.

constraints:

$$\max_{\mathbf{x},\mathbf{y}} \sum_n \sum_m \sum_{\boldsymbol{\ell}} \pi(c(\boldsymbol{\ell}_{nm}), \tilde{c}_n) x_{\boldsymbol{\ell}nm} - \sum_k \sum_{\ell \in L_k} y_{\ell k}\phi_{\ell k} - \sum_n \sum_m y_{nm}\phi_{nm},$$

$$\begin{aligned} \text{subject to } & \sum_{\ell \in L} x_{\boldsymbol{\ell}nm} \le y_{nm}, \forall n, m \\ & x_{\boldsymbol{\ell}nm} \le y_{\boldsymbol{\ell}_k k}, \forall n, m, k \\ & x_{\boldsymbol{\ell}nm} \ge 0, \\ & y_{\ell k}, y_{nm} \in \{0, 1\}, \forall \ell, n, m, k. \end{aligned}$$

In this formulation, we only need to specify the facility location and market entry decisions as binary variables. Given the constraints, the optimal path indicator variables $x_{\boldsymbol{\ell}nm}$ will always be either 0 or 1.[5]

5 Single-model Policy Simulation

We simulate a production chain of a single firm with one car model. There are three stages of production (cell, pack, vehicle assembly) in a world of 40 locations[6] (production and consumption sites). The coordinates of site are randomly drawn from the unit interval. Demand is assumed to rise with the square root of latitude: Northern locations demand more EVs than Middle locations, which themselves consume more electric cars than the South. To make costs observable in the maps we draw, we assume the marginal cost of production is determined by the Euclidean distance between countries, $c_{mn\boldsymbol{\ell}} = \tilde{d}_{\ell_1,\ell_2} + \tilde{d}_{\ell_2,\ell_3} + \tilde{d}_{\ell_3,n}$, where $\tilde{d}_{ab} = [(\text{lat}_a - \text{lat}_b)^2 + (\text{lat}_a - \text{lat}_b)^2]^{1/2}$. Fixed costs of establishing plants at any stage are set to be the same for all locations and stages. Fixed costs of market access are drawn randomly from a uniform distribution.

The low and high fixed costs to revenue ratios implied by the parameter settings in the simulation are 6.2% and 16%. Using newspaper reports on investment costs, we see that on average EV assembly plants involve a $0.85 billion investment and battery plants require $2.45 billion. We assume the battery investment costs refer to packs and include the fixed costs of the cell facilities. There are 205 EV assembly plants and 101 pack plants, so total EV investment is approximately $4.15 billion. Annualizing these up-front invest-

[5]In addition, we formulate the first constrain as an inequality, even though it will always hold with equality at the optimum. This is the version of the problem that we implement computationally (the solvers work more efficiently with constraints specified as inequalities versus equalities).

[6]As of 2021, 38 countries are involved in one or multiple stages of the production of EVs.

ments at rates or 5% or 10%, we compute fixed cost to revenue ratios for the EV industry as 7% to 14%. Our revenue calculation assumes an average price of $50,000 per EV for the 6.3 million EVs assembled worldwide in 2022. These "back of the envelope" calculations suggest that the fixed cost to revenue ratios in the simulation are in the right ballpark.

Each of the mapped simulation below takes about 20–40 seconds[7] to solve the firm's location problem, depending on the computer and the levels of fixed costs and subsidies. This establishes that the mixed integer programming solver can handle realistically sized problems in a reasonable amount of time. The code, which will be made available online, is written in Julia and uses the JuMP package to interface with the solver.

Figure 6: Simulation: no interventions

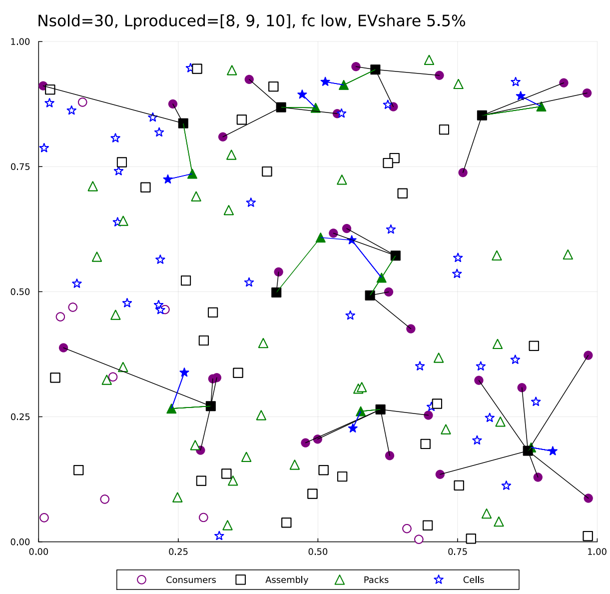

(a) Low fixed costs

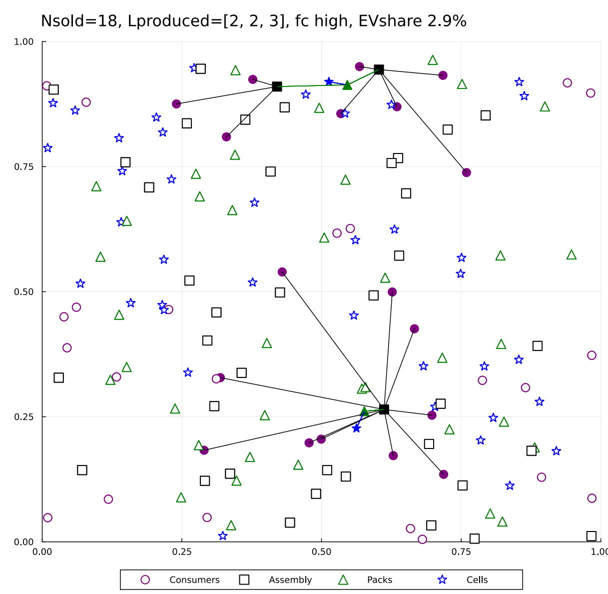

(b) High fixed costs

We begin the investigation of the MUFLP by simulating the model without any subsidies for EVs. The left panel of Figure 6 shows the case of low fixed costs. We see the three key features of the model: (i) there are more plants in the North, reflecting the higher demand, (ii) for any given latitude, plants tend to be central to be closer to more consumers, (iii) peripheral and southern consumers are less likely to be served (one quarter of the consumers do not have sufficient demand to justify the fixed costs of serving them). In the right panel, we multiply fixed costs by five. This reduces the number of cell plants by

[7]The program can be solved an order of magnitude more quickly in the arc-based formulation (ABF) described in Appendix A.2. Unfortunately while this method minimizes cost given quanity and market entry decisions, it is not guaranteed to maximize profits.

a factor of four and the number of assembly plants by a factor of three. The greater distances from assembly to consumer lower the number markets served from 30 to 18. The combination of higher marginal costs and reduced number of markets lowers the share of EVs in the market from 5.5% to 2.9%.

Figure 7: Simulation: $5,000 subsidies in Middle

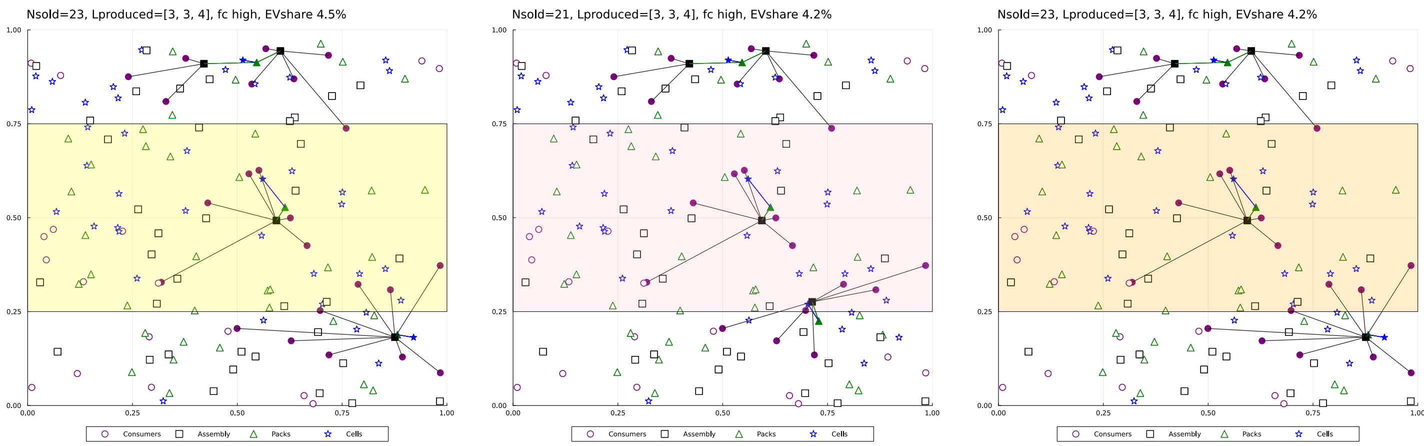

(a) Unconditional subsidy

(b) Final-assembly

(c) Full-path

Figure 7 brings consumer subsidies into the picture, continuing the high fixed cost case from Figure 6. The three frames correspond to three different types of rules for which vehicles qualify for the subsidy. In panel (a) any consumer in the designated region—shown in yellow—qualifies for the $5,000 dollar subsidy. This resembles the way credits are allocated in Canada and many European countries. In panel (b), the subsidy is only available for vehicles assembled in the region. This is similar to the situations in China and the US where there are consumer subsidies coupled with large tariffs on imported EVs. Panel (c) takes its inspiration from the IRA provisions requiring the whole value chain of batteries to come from countries with whom the US has a free trade agreement.

The unconditional subsidy in Middle (panel a) leads to a new assembly plant being opened in the South, and the existing plant in the Middle is replaced by a more centrally located plant. Five new consumer sites are activated, and the EV share rises from 2.9% to 4.5%. Making the subsidy eligible only to regionally produced vehicles leads assembly site serving Southern consumers to be relocated to the lower part of Middle. Two markets in the South are closed due to higher marginal costs. The EV share falls to 4.2%. Panel (c) makes qualification rules even stricter. The pack plant used in panel (b) would disqualify vehicle sales from the plant in the southern part of Middle. It is replaced by a new plant in the South, which leads to reactivating two Southern consumers. There are conflicting effects on affordabilty that leave the overall market share of EVs unchanged. This case illustrates the Laffer curve effect identified by Head et al. (2024): stricter rules can be counterproductive in terms of domestic production.

Figure 8: Simulation: $7,500 subsidies in Middle

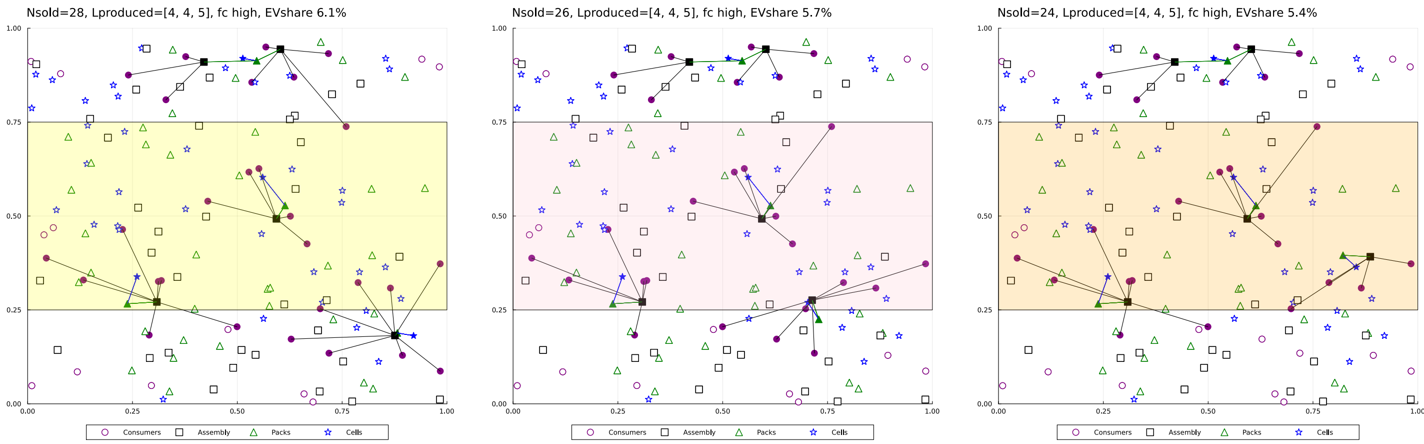

(a) Unconditional subsidy (b) Final-assembly (c) Full-path (IRA)

The simulation underlying Figure 8 keeps the same cost parameters as in Figure 7 but raises the subsidy to $7,500 to mimic the IRA. The key points are that the higher subsidy increases the number of activated consumers, and adds an extra plant at each stage. It also increases the EV share under all three policy rules. The interesting point is that the higher subsidy eliminates the outcome in panel (c) where the firm responds to the more stringent content rule by decamping to the South. However, the higher costs of complying with the rule in panel (c) imply that now the EV share falls relative to panel (b).

Figure 9: Simulation: $5,000 subsidies in North

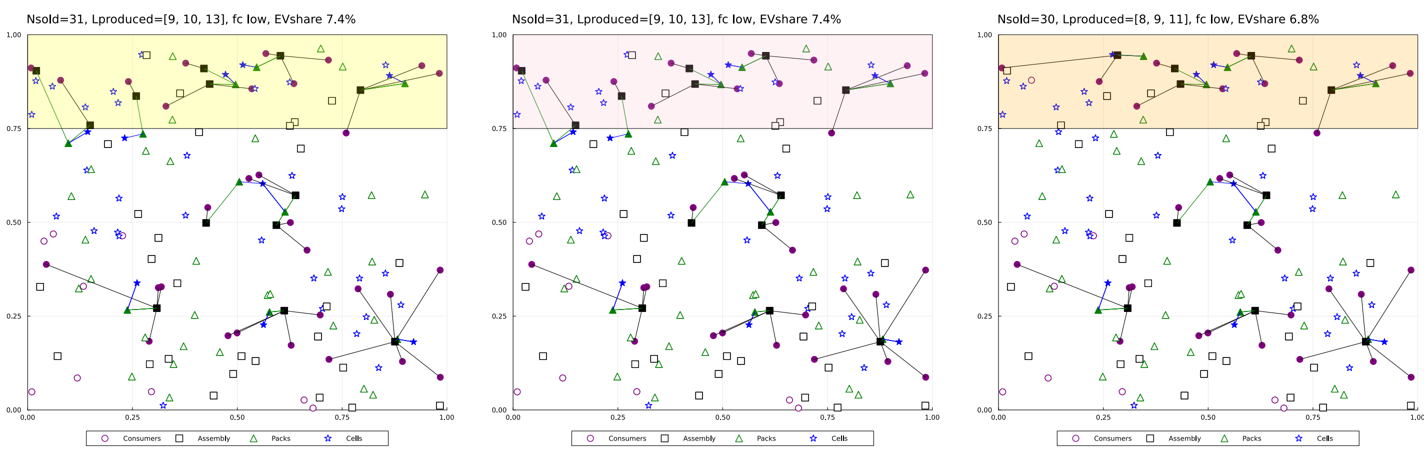

(a) Unconditional subsidy (b) Final-assembly (c) Full-path

Figure 9 departs from the previous case in three ways. First we revert back to $5,000 subsidies, second we lower fixed costs, and third we place the subsidy in the North. The key points to note are that in panels (a) and (b) two Northern plants are served by cell-pack value chains in Middle. The full-path rule in panel (c) eliminates the possibility of qualifying for subsidies with such paths. This lowers EV shares and reduces the number of consumers served. It also eliminates four factories. The EV share is still higher than under *laisses-faire* in Figure 6 (a) but the policy is unsuccessful from both the point of view of job creation (four fewer plants) and displacing ICE vehicles (lower EV share).

The takeaway from this series of simulations is that in a world where fixed costs are important, subsidies can reshape the geography of production but do so in ways that are not easy to predict, either quantitatively or even qualitatively. To obtain a better idea of the effects of the gamut of recent policies, we are going to need parameter estimates grounded in the data.

6 Estimation of the Theoretical Model

Applying a system of sourcing regressions, we can get the estimates of $w_{\ell_k} + \tau_{\ell_{k+1}\ell_k}$ for all locations at each level of production, and estimates of the Gumbel scale parameter θ_k driving the distribution of the random cost term $u_{\ell_{k+1}\ell_k}(m)$ (up to a normalization). These estimates for $k = 1, ..., K$ will be used to construct equation (A.1) together with the simulated Gumbel draws, and then solve the GVC for each firm-model-consumer, given guesses of fixed costs. Lastly, we plan to estimate the fixed costs via constrained maximum likelihood as in Tintelnot (2017) or via moment inequality.

6.1 Conditional choice estimation of variable costs

A nested logit model is used to estimate parameters driving variable costs along the value chain. Recall that edge costs are given by

$$c^{km}_{\ell_{k+1}\ell_k} = \beta_k(w_{\ell_k} + \tau_{\ell_{k+1}\ell_k} + u^m_{\ell_{k+1}\ell_k}),$$

with $\beta_k \equiv \prod_{k+1}^{K} \gamma_k$, where γ_k are input requirements at stage k. The per unit production costs in b are denoted with w_{ℓ_k}, and $\tau_{\ell_{k+1}\ell_k}$ is the per unit trade cost including border and distance costs plus tariffs on the edge from ℓ_k to ℓ_{k+1}. Finally $u^m_{\ell_{k+1}\ell_k}$ is the unobserved cost shocks on edge $\ell_{k+1}\ell_k$ for car model m. We assume that the sourcing decisions can be estimated through nested logit, with the following Ben-Akiva and Lerman (1985) conditions regarding $u^m_{\ell_{k+1}\ell_k}$:

1. All disturbance terms u are mutually independent.

2. The terms $u_{\ell_2\ell_1}(m)$ are independent and identically Gumbel distributed with scale parameter θ_1.

3. The terms $u_{\ell_{k+1}\ell_k}(m)$ is distributed so that $\sum_{k=1}^{r}\beta_k u_{\ell_{k+1}\ell_k}(m)$ is Gumbel distributed with scale parameter θ_r, for $r = 2, ..., K$.

The scale parameters $\boldsymbol{\theta}$ vary across levels of production. We there let each stage's input be heterogeneous in various degrees. An alternative approach is followed by Antràs and De Gortari (2020) where the stage-specific Fréchet distribution is assumed to have the same shape parameter.

Conditional on the set of level 1 facilities established by a firm $\mathcal{L}_1(f)$, it chooses for each model the location with the lowest cost to serve ℓ_1 due to constant returns to scale in the variable costs of production. We can then derive that the conditional probability of ℓ_1 being the cost-minimizing location of supplying to ℓ_2 as

$$\Pr\left(\ell_1|\ell_2;\mathcal{L}_1(f)\right) = \Pr\left(c^{1m}_{\ell_2\ell_1} \leq \min_{\ell' \neq \ell_1, \ell' \in \mathcal{L}_1(f)}\{c^{1m}_{\ell_2\ell'}\}\right) = \frac{\exp\left[-\theta_1\left(w_{\ell_1} + \tau_{\ell_2\ell_1}\right)\right]}{\sum_{\ell\in\mathcal{L}_1(f)}\exp\left[-\theta_1\left(w_\ell + \tau_{\ell_2\ell}\right)\right]}. \tag{6}$$

An important component of the nesting approach to estimating production costs along the value chain is the integration of the cost of upstream parts in downstream decisions. The natural approach that follows from the assumptions of Ben-Akiva and Lerman (1985) is to include in the choice of a location for the battery pack the expected cost of cells that will be used in that pack. With extreme value distributions, this expected cost for pack plant ℓ_2 is

$$\mathcal{E}_{\ell_2}(f) = -\frac{1}{\theta_1}\ln\sum_{\ell\in\mathcal{L}_1(f)}\exp\left[-\theta_1\left(w_\ell + \tau_{\ell_2\ell}\right)\right]. \tag{7}$$

We refer to this term as the *inclusive cost* which depends on the sourcing potential of ℓ_2 and differs across firms due to various choice sets $\mathcal{L}_1(f)$. For example, Tesla and Audi may have different inclusive costs for its pack plants in China because Tesla has contracted with cell suppliers in China, whereas Audi has only contracted with cell suppliers in South Korea and Europe. The inclusive cost for Tesla's pack plants in China also differs from those in the US due to different trade costs from its set of cell plants.

Moving down the supply chain and applying condition 3, the probability that ℓ_2 is selected conditional on ℓ_3 being chosen is

$$\Pr\left(\ell_2|\ell_3;\mathcal{L}_2(f)\right) = \frac{\exp\left[-\theta_2\left(w_{\ell_2} + \tau_{\ell_3\ell_2} + \gamma_2\mathcal{E}_{\ell_2}(f)\right)\right]}{\sum_{\ell\in\mathcal{L}_2(f)}\exp\left[-\theta_2\left(w_\ell + \tau_{\ell_3\ell} + \gamma_2\mathcal{E}_\ell(f)\right)\right]}, \tag{8}$$

where γ_2 is the number of stage 1 parts (cells), that enter stage 2 (packs). Generalizing to the k-level, the conditional sourcing probability is

$$\Pr\left(\ell_k|\ell_{k+1};\mathcal{L}_k(f)\right)=\frac{\exp\left[-\theta_k\left(w_{\ell_k}+\tau_{\ell_{k+1}\ell_k}+\gamma_k\mathcal{E}_{\ell_k}(f)\right)\right]}{\sum_{\ell\in\mathcal{L}_k(f)}\exp\left[-\theta_k\left(w_\ell+\tau_{\ell_{k+1}\ell}+\gamma_k\mathcal{E}_\ell(f)\right)\right]},\tag{9}$$

where the inclusive cost of k-level producers is

$$\mathcal{E}_{\ell_k}(f)=-\frac{1}{\theta_{k-1}}\ln\sum_{\ell\in\mathcal{L}_{k-1}(f)}\exp\left[-\theta_{k-1}\left(w_\ell+\tau_{\ell_k\ell}+\gamma_{k-1}\mathcal{E}_\ell(f)\right)\right].\tag{10}$$

At each stage, the conditional sourcing probability can be estimated empirically with PPML. The dependent variable is the discrete choice (taken by the plant in stage $k+1$) of sourcing stage k from ℓ_k rather than any other available location in the set $\mathcal{L}_k(f)$:[8]

$$\mathbb{1}\left(\ell_k=1|\ell_{k+1};\mathcal{L}_k(f)\right)=\exp\Big(\mathrm{FE}_{\ell_k}+\mathrm{FE}_{fk}-\theta_k[\boldsymbol{X}'_{\ell_{k+1}\ell_k}\delta_k+\epsilon_{\ell_{k+1}\ell_k}+\mathbb{1}_{k>1}\gamma_k\mathcal{E}_{\ell_k}(f)]\Big),\tag{11}$$

where the origin fixed effects $\mathrm{FE}_{\ell_k}=-\theta_k\left(w_{\ell_k}\right)$, captures local production costs, and $\mathrm{FE}_{fk}=-\ln\sum_{\ell\in\mathcal{L}_k(f)}\exp\left[-\theta_k\left(w_\ell+\tau_{\ell_{k+1}\ell}+\gamma_k\mathcal{E}_\ell(f)\right)\right]$ accounts for the attractiveness of the set of potential choices available to firm f at stage k. The bilateral trade costs are parameterized as $\tau_{\ell_{k+1}\ell_k}=\boldsymbol{X}'_{\ell_{k+1}\ell_k}\delta_k+\epsilon_{\ell_{k+1}\ell_k}$, with $\boldsymbol{X}_{\ell_{k+1}}$ including bilateral frictions such as bilateral distance, border, RTA and tariffs, and $\epsilon_{\ell_{k+1}\ell_k}$ being the unobservable frictions. We estimate this nested logit model sequentially from the most upstream stage where inclusive cost is zero, and continue downward through the supply chain using the inclusive cost constructed from the $k-1$ level in estimating the k-level parameters.

6.2 Empirical results of variable costs estimation

Table 7 combines estimation results for the three stages considered in our analysis of the BEV industry: battery cells, battery packs ans assembly. The first and last three columns use different specifications for $\boldsymbol{X}_{\ell_{k+1}}$. Columns (1) to (3) use national borders, distance and RTAs. The last three replace the RTA dummy with per unit tariffs between the two locations/stages; which is obviously very correlated with RTAs, since all RTAs approved under GATT/WTO have to achieve zero tariffs after a maximum of ten years on manufacturing goods.

[8]The estimation of those conditional logit problems through PPML with appropriate sets of fixed effects uses the equivalence known in the literature since (Guimaraes et al., 2003) and has been used in a similar context by Head and Mayer (2019).

Table 7: Combined regression results

	(1) Cells	(2) Packs	(3) Vehicles	(4) Cells	(5) Packs	(6) Vehicles
border	-1.64[a]	-3.74[a]	-1.72[a]	-0.713	-3.21[a]	-0.797[a]
	(0.471)	(0.898)	(0.460)	(0.472)	(0.462)	(0.305)
log distance	-0.565[a]	-0.231[a]	-0.069	-0.566[a]	-0.234[a]	-0.093[c]
	(0.011)	(0.018)	(0.054)	(0.012)	(0.019)	(0.052)
RTA	1.03[b]	0.965[b]	0.833[a]			
	(0.407)	(0.475)	(0.177)			
EC		-0.067	-0.485[a]		-0.033	-0.021[a]
		(0.052)	(0.093)		(0.026)	(0.004)
tariff ($)				-2.71[a]	-0.197	-0.137[a]
				(0.740)	(1.00)	(0.045)
θ	5.83	2.44	0.34	2.71	0.11	0.08
Cost elasticity	-26.82	-8.43	-17.19	-12.48	-0.40	-3.94
Observations	17,280	16,005	54,933	17,280	16,005	54,933
Squared Correlation	0.511	0.320	0.244	0.513	0.319	0.247

Note: EC stands for expected costs, with definition being given in equation (10).

The effect of distance on sourcing decisions decline in absolute value as we move downstream. The effect is quite strong (around -0.57) and significant for cells, but is reduced by more than half for packs, and is finally quite small for the distance separating assembly from final destination. This last result contrasts with the ones obtained with vehicles of all fuel types (and therefore mostly ICEs) from Head and Mayer (2019), where distance effects were stronger. BEVs therefore seem to be less sensitive to physical distance once assembled.[9] However, the RTA effect in column (3) is notably stronger than for ICEs, signaling that intra-continental / intra-regional trade is more dominant than for traditional vehicles. The RTA effects are uniformly strong along the value chain, approximately tripling the sourcing probablility ($\exp(1.03) = 2.80$ for cells). The effect of national borders is also very strong, specially for pack sourcing (with a coefficient of -3.74, there is very little room for packs being sourced outside of national borders). The coefficients on expected costs of upstream stages are overall small, and largely insignificant for packs, when the effect is more robust for assembly.

We have two ways to estimate θ_k, an important parameter when converting estimates of frictions into costs. The easiest is to include directly per unit bilateral tariffs on the

[9] An important distinction between the two studies is that the present one is much more granular, being able to geo-locate all plants, rather than to simply assign them to countries.

relevant stage.[10] Since the coefficient on tariffs directly provides an estimate of θ_1 (cells), column(4) estimate is the tariff elasticity at 2.71. For packs and vehicles, we average coefficients of tariffs and expected costs, resulting in lower and noisy values. The second approach uses the fact that Border and RTA dummies also reflect the impact of tariffs. We can compute preferential margins, i.e. the MFN tariff for border and the difference between MFN and applied rates for RTAs, and use those margins to infer θ_k.

- For stage $k = 1$, we proceed as following:
    1. Estimate cell sourcing with RTA and Border (w/o tariffs, column 1).
    2. Across locations ℓ considered as a source for cells, compute $\theta_1^{\text{RTA}} = \frac{\hat{\beta}^{\text{RTA}}}{v_2 \text{mean}(\text{tar}^{\text{MFN}}_{\ell_2\ell} - \text{tar}^{\text{RTA}}_{\ell_2\ell})}$, and $\theta_1^{\text{Border}} = \frac{\hat{\beta}^{\text{Border}}}{v_2 \text{mean}(\text{tar}^{\text{MFN}}_{\ell_2\ell})}$
    3. Our final estimate is $\hat{\theta}_1 = \text{mean}(\theta_1^{\text{RTA}}, \theta_1^{\text{Border}})$
    4. This is an upper bound because of the existence of NTBs in the impact of borders and RTAs.
- For $k > 1$ (in columns 2 and 3), there is one additional estimate of θ_k coming from the expected cost term, $-\gamma_k \mathcal{E}_{\ell_k}(f)$.

The resulting θ_k are 5.83, 2.44 and 0.34 for cells, packs and vehicles respectively. Recalling that θ is a measure of how homogeneous are the random terms across locations in the variable costs of a car model for a given stage, it seems intuitive that this measure of homogeneity is declining as one gets closer to the final good. There is a correspondence between θ_k and the usual cost elasticity, obtained by multiplying θ_k by the unit price of cells/packs/vehicles. Those are shown under the θ_k line in Table 7. There is wide variation in those elasticities, and now obvious comparison with the literature possible regarding cells and packs. Averaging the ones for vehicles from columns (3) and (6), one obtains a cost elasticity of 10.5, larger but not out of range from than the one obtained by Head and Mayer (2019) for all vehicles.

Figure 10 uses the estimates of country fixed effects to reveal relative costs for each stage. This is obtained as $w_{\ell_k} = -\text{FE}_{\ell_k}/\theta_k$, and normalized with respect to the USA in each k. The competitiveness of China, South Korea and Japan is visible for cells. Poland

[10] We use WITS database to identify preferential and MFN tariffs: For cells, we use HS 850650, for packs HS 850760, and for BEVs, HS 870380 and 870490 for passenger cars and light commercial vehicles respectively. WITS provides ad valorem tarif rates. We transform those into unitary taxation (in USD per units) using values inferred from figures relevant for the Tesla Model 3 in 2021. Those are 43800 euros for the car, 12100 euros per battery pack (both converted in USD). Since there are about 3000 cells in this pack, the unit price of a cell is 4.6 USD.

Figure 10: Relative costs in the value chain of BEVs

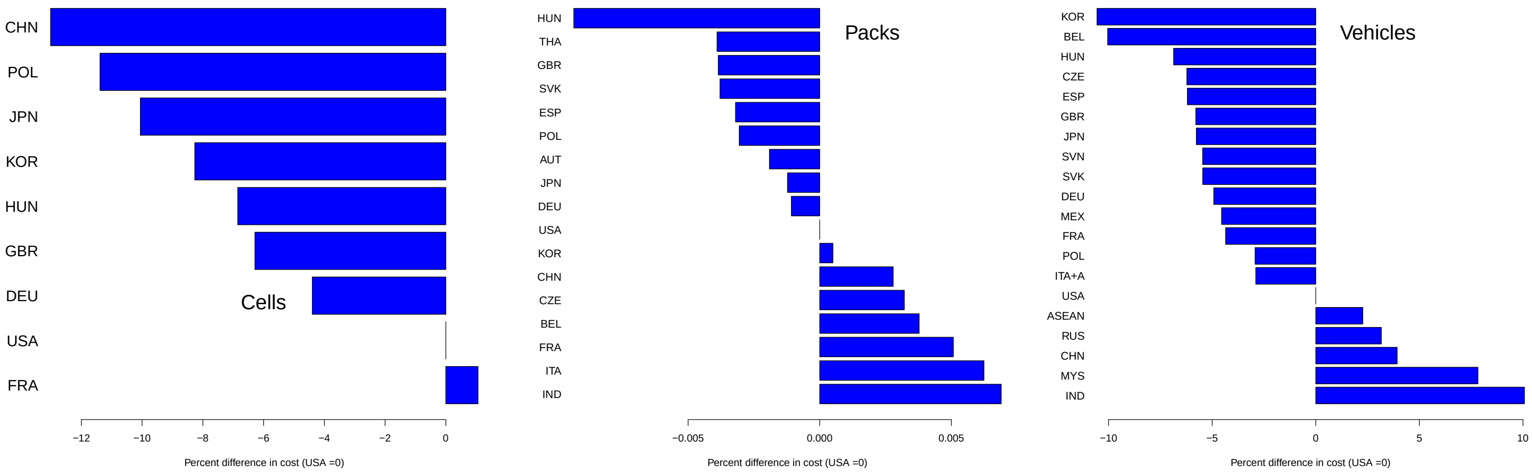

(a) Cells (b) Packs (c) Vehicles

and Hungary are also estimated to have low costs of cells, based on their capacity to serve pack factories in many places in Europe. Hungary is also (with Slovakia and Poland) a performing place in terms of pack production. Korea and Belgium are identified as very good places for EV assembly as opposed India, Malaysia or even China. This is due to the fact that the electric Volvos and Audis assembled in Belgium sell in many places in the world (not only EU), when most Indian, Malaysian and Chinese-assembled EVs are sold locally. To take precise examples: The Hyundai electric Kona produced in India is sold only in India, when the Korean-made one is sold in 59 destinations, and the figure for the Czech-made one in 39. The case of Malaysia is also quite telling since it assembles electric Volvos (the same as in Belgium) since 2022 that are sold only in Malaysia, and a few in Thailand.[11] In 2022, the Volswagen electric SUV ID.4 is produced in 4 plants: 2 in Germany, 1 in the USA, 2 in China. Each of the German plants serves at least 29 markets, the US plant 2 markets, and the 2 Chinese plants serve only China. Those FE estimates reflect those patterns: while the situation is evolving rapidly for China, in our 2015-2022 sample most of the Chinese plants still serve only China. Other consumers in the world prefer to source their BEVs from alternative plants, which translates into high estimated costs.

[11] https://paultan.org/2023/02/03/volvo-car-malaysia-to-begin-exporting-cars-made-in-shah-alam-to-the-philippines-and-vietnam-this-year/.

6.3 Completing the parameter set for the model

Calibration of demand: Using car price data at each market, and sales of EV vs. ICE cars,

$$q_{mn} = \frac{\exp(A_n - \alpha p_{mn})}{1 + \sum_j y_{jn} \exp(A_n - \alpha p_{jn})} B_n$$
$$Q_n^{\mathrm{EV}} = \sum_m y_{mn} q_{mn} = \frac{\exp(A_n) \sum_m y_{mn} \exp(-\alpha p_{mn})}{1 + \sum_j y_{jn} \exp(A_n - \alpha p_{jn})} B_n$$
$$Q_n^{\mathrm{ICE}} = \frac{1}{1 + \sum_j y_{jn} \exp(A_n - \alpha p_{jn})} B_n,$$

we can rearrange those equations and solve for the demand shifter A_n in each destination as

$$A_n = \ln\left(\frac{Q_n^{\mathrm{EV}}}{Q_n^{\mathrm{ICE}}}\right) - \ln\left(\sum_m y_{mn} \exp(-\alpha p_{mn})\right).$$

In order to proceed with that inversion, we need a value for α. First note that own price elasticity of a car model j with logit demand is αp_j. We therefore alibrate α as the average own price elasticity divided by the average price of all car models for a given year. Table 8 provides a number of average elasticities from a set of papers that estimate demand for cars. The left column gives figures for papers focusing on EVs, while the right column does not distinguish between EVs and traditional gas-powered vehicles, and therefore mostly covers ICEs. In table 8, we see that several recent papers find own price elasticities centered around four (the mean all studies is 5, but EVs appear to have lower elasticities, averaging 3.25 across the EV-specific papers).

Table 8: Price responsiveness of car demand from recent literature

EV-only	**elas.**	**Mainly ICE**	**elas.**
Barwick et al. (2024)	4.2	Beresteanu and Li (2011)	8.4
Kwon (2023)	4.4	Castro-Vincenzi (2024)	4.3
Li et al. (2022)	3.7	Colon and Gortmaker (2020)	3.9
Li (2023)	3.7	Coşar et al. (2018)	14.9
Li et al. (2017)	1.3	Goldberg (1995)	3.3
Linn (2022)	5.3	Goldberg and Verboven (2001)	5.2
Muehlegger and Rapson (2022)	2.1	Head and Mayer (2019)	3.9
Springel (2021)	1.8	Li (2018)	9.5
Xing et al. (2021)	2.8		

Figure 11 provides the result for the year 2022, expressed in relative terms with respect the the United States. The top position of Norway in A_n ranking reflects very generous policies, high income, and probably preferences for clean vehicles. On the other side of the spectrum, China's low A_n reflects low EV adoption *relative to the low prices and high EV variety*.

Figure 11: Relative EV appeal (A_n) for selected markets

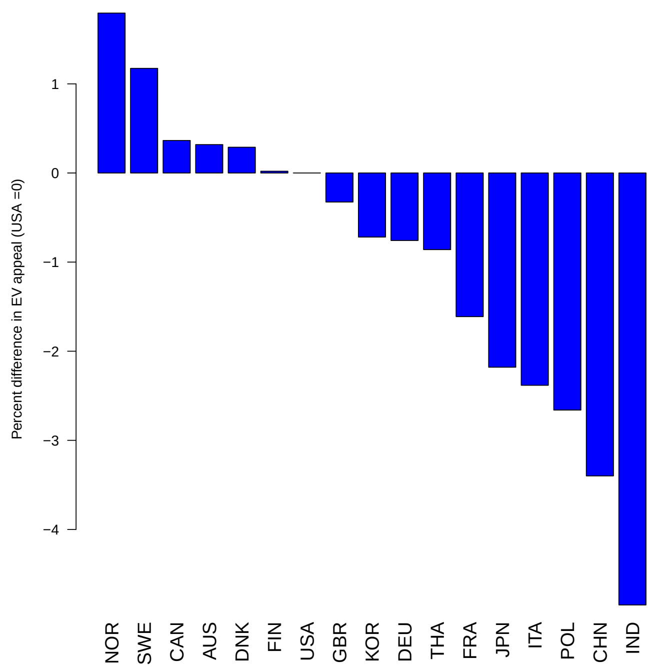

7 Proposed Counterfactual

7.1 Unilateral subsidy

The first policy we will consider is the unconditional subsidy shown in the left panel of figures 7–9. Outsourcing inputs from foreign countries could reduce the government's willingness to let subsidy rents flow to domestic firms. If the government internalizes the surplus of domestic consumers and domestic producers only, transfers flowing to foreign firms will not be credited to the government's "surplus ledger." This will reduce the government's incentives given to domestic firms, which is the subsidy-leakage effect documented in Bown et al. (2022).

In this counterfactual exercise, we will predict the spatial distribution of electric vehicle supply chains after the implementation of EV tax credits that to consumers from a given country, but that are available for all vehicles regardless of the locations of assembly and component procurement.

7.2 Unilateral subsidy conditional on sourcing locations

One way to mitigate the subsidy leakage is through regulating the input sources of cars to be eligible for the subsidy. The US IRA is one example. Beginning in 2023, 40% of an EV battery's minerals and 50% of the components will need to come from the US or Free Trade Agreement (FTA) partners. In 2027 and 2029, this requirement will increase to 80% for minerals and 100% for components.

Based on the subsidized input source locations, we will again predict the geography of global supply chain of EVs and the associate welfare effects to compare to the results from the first counterfactual.

7.3 Subsidy with international cooperation

Since outsourcing inputs introduces distortion due to the domestic government's desire to limit leakage of the subsidy to foreign firms, there is thus scope for international cooperation to improve global surplus. Suppose that cooperation comes in the form of an agreement among countries to maximize expected global surplus. The country will then increase the subsidy from the equilibrium levels absent international cooperation.

Then the question is whether the subsidy provided by the US IRA is too low due to the subsidy-leakage effect, or too high considering the recent competition in EV adoption by different countries, or coincidentally achieve the global optimum? In order to answer this question, we need to derive the analytical form of equilibrium subsidy at the country level and global level as a function of facilities locations and paths of sourcing. The international treaty on EV subsidy presumably involves the US, EU, and China. Since the US and China have already had relevant subsidies in place, we will examine the welfare effects of an EV subsidy implemented uniformly by the EU countries at the US-equivalent level (US$7,500 with input requirements from EU countries) and the model-predicted equilibrium level.

8 Conclusion

To be added when the estimation and counterfactuals are completed.

References

Antràs, P. and A. De Gortari (2020). On the geography of global value chains. *Econometrica 88*(4), 1553–1598.

Antràs, P., T. C. Fort, and F. Tintelnot (2017). The margins of global sourcing: Theory and evidence from US firms. *American Economic Review 107*(9), 2514–64.

Arkolakis, C., F. Eckert, and R. Shi (2023, November). Combinatorial discrete choice: A quantitative model of multinational location decisions. Working Paper 31877, National Bureau of Economic Research.

Ben-Akiva, M. E. and S. R. Lerman (1985). *Discrete choice analysis: theory and application to travel demand*, Volume 9. MIT press.

Bown, C. P., C. M. Snyder, and R. W. Staiger (2022). Vaccine supply chain resilience and international cooperation. Dartmouth College Economics Department working paper.

Eaton, J. and S. Kortum (2002). Technology, geography, and trade. *Econometrica 70*(5), 1741–1779.

Guimaraes, P., O. Figueirdo, and D. Woodward (2003). A tractable approach to the firm location decision problem. *The Review of Economics and Statistics 85*(1), 201–204.

Head, K. and T. Mayer (2019). Brands in motion: How frictions shape multinational production. *American Economic Review 109*(9), 3073–3124.

Head, K., T. Mayer, and M. Melitz (2024). The Laffer curve for rules of origin. *Journal of International Economics 150*, 103911.

Jia, P. (2008). What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry. *Econometrica 76*(6), 1263–1316.

Ortiz-Astorquiza, C., I. Contreras, and G. Laporte (2018). Multi-level facility location problems. *European Journal of Operational Research 267*(3), 791–805.

Tintelnot, F. (2017). Global production with export platforms. *The Quarterly Journal of Economics 132*(1), 157–209.

A Multilevel Uncapacitated Facility Location Problems

A.1 Solving the combinatorial choice problem

Choice problems with binary variables that are interdependent impose severe computational challenges. The option to enumerate all the potential choices—brute force—is not practical even because the number of possible combinations grows exponentially with the number of choices. This rather quickly leads to exhaustion of the available computer memory and exploding computation time.

We use the Gurobi optimizer, a commercial optimizer that is available for free to licensed academic users. Gurobi does not make public the precise algorithms that it uses to solve mixed-integer linear programs (MILPs). However, the optimization log shows that it uses linear relaxations, heuristics, cutting planes, and finally the branch and bound algorithm.[12] There is an open-source alternative to Gurobi called HiGHS. However, Gurobi tends to be five to ten times faster than HiGHS.

The main solution method we employ is the path-based profit maximization approach outlined above. This approach guarantees a global optimum but as the number of stages increases, it becomes increasingly demanding in terms of time and memory. An alternative way to incorporate the endogenous decision on market access and quantity supplied, is to apply a nested procedure where the inner loop solves the cost-minimization MUFLP, and the outer loop updates the endogenous quantity based on equation (1) and (**??**) with a check on whether the market access condition is met as in equation (**??**). We solve the two loops iteratively until demand in the distributed markets converge. The advantage of such an algorithm is that the cost-min MUFLP can be formulated based on arcs (edges) rather than complete paths. The disadvantage is that the algorithm frequently converges to a local optimum. In most cases the solutions from the two methods are quite close but there are no guarantees. We have found instances where the arc-based algorithm converges to solutions that are seriously underperform.

A.2 The Arc-Based-Formulation of the cost-minimization problem

Let $x_{nab}^{km} = 1$ if consumers at n uses the edge $\{a, b\} \in E_k$ for model m where the sets E_k contain all possible edges between $a \in L_{k+1}$, $b \in L_k$ and $k = 1, ..., K$. Since we allow for an endogenous choice of market access, we introduce a third binary decision $y_n^m = 1$ if and only if the car model m is sold at location n. The assignment and distribution

[12] https://www.gurobi.com/resources/mixed-integer-programming-mip-a-primer-on-the-basics/

decisions $\{x_{nab}^{km}, y_n^m\}$ are model specific, and the location decisions $\{y_{\ell_k}^f\}$ for $k = 1, ..., K$ are firm-specific.

Holding fix the set of models with market access y_n^m and endogenous quantity supplied q_n^m, the firm's profit maximization problem can be reformulated to a problem of cost minimization. Each firm solves its production location and assignment decisions knowing its models' variable costs at every level of production for any location configurations c_{ab}^{km} as in equation (3), and fixed cost for every location-stage $\phi_{\ell_k}^f$.

$$\min_{\{\boldsymbol{x}, \boldsymbol{y_\ell}\}} \sum_{n \in N} \sum_{m \in M_f} \sum_{k=1}^{K} \sum_{\{a,b\} \in E_k} q_n^m y_n^m c_{ab}^{km} x_{nab}^{km} + \sum_{k=1}^{K} \sum_{\ell_k \in L_k} \phi_{\ell_k}^f y_{\ell_k}^f \tag{A.1}$$

subject to

$$\sum_{b \in L_K} x_{nnb}^{Km} = y_n^m, \quad n \in N, m \in M_f \tag{A.2}$$

$$\sum_{b \in J_{k-1}} x_{nab}^{(k-1)m} = \sum_{b' \in L_{k+1}} x_{nb'a}^{km}, \quad n \in N, m \in M_f, a \in L_k, k = 2, ..., K \tag{A.3}$$

$$\sum_{a \in L_{k+1}} x_{nab}^{km} \leq y_b^f, \quad n \in N, m \in M_f, b = \ell_k \in L_k, k = 1, ..., K \tag{A.4}$$

$$x_{nab}^{km} \geq 0, \quad n \in N, m \in M_f, \{a, b\} \in E_k, k = 1, ..., K \tag{A.5}$$

$$y_{\ell_k}^f \in \{0, 1\}, \quad \ell_k \in L_k, k = 1, ..., K. \tag{A.6}$$

where

$$c_{ab}^{km} = \beta_k a_k w_b + \beta_k t_{ab} - \beta_k u_{ab}^{km}. \tag{A.7}$$

Equations (A.2) and (A.3) are *flow conservation constraints*. In constraint (A.2), each consumer, if chosen to sell, must be served by one facility. In constraint (A.3), total flow that belongs to consumers in n and comes into a facility, should be equal to the amount from this facility. Equations (A.4) are *activity constraints* which govern a facility must be open to serve a client or a downstream producer. The inequalities guarantee that flow can only exist on the constructed edges. Finally, equations (A.5) and (A.6) state the aforementioned decision variables.

B Additional empirical results

B.1 Stage-level regression results

Table B1: Cell Sourcing decision

Dependent Variable:	Cell sourcing decision					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
Border	-1.65[a]	-1.94[a]	-2.58[a]	-2.46[a]	-2.47[a]	-2.45[a]
	(0.472)	(0.491)	(0.594)	(0.518)	(0.511)	(0.513)
log distance	-0.565[a]	-0.427[a]	-0.182[a]	-0.184[a]	-0.158[a]	-0.128[b]
	(0.011)	(0.026)	(0.067)	(0.067)	(0.049)	(0.059)
RTA	1.04[b]	1.04[b]	1.10[b]	1.30[a]	1.67[a]	1.66[a]
	(0.406)	(0.404)	(0.433)	(0.441)	(0.390)	(0.395)
Intrafirm		1.89[a]	0.849[a]	0.501[a]	0.148	0.312[c]
		(0.117)	(0.182)	(0.153)	(0.149)	(0.167)
Intraplant			3.00[a]	2.72[a]	3.12[a]	3.58[a]
			(0.419)	(0.366)	(0.279)	(0.329)
Observations	17,294	16,655	16,655	11,203	8,295	7,762
Squared Correlation	0.511	0.515	0.529	0.472	0.523	0.537

Column (4) constrains the set of choosers to be car models with positive sales (strictly defined). Column (5) adds the constraint that the choice set constitutes of plants already producing the required cell shape. Column (6) adds the constraint that the choice set constitutes of plants already producing the required cell in terms of chemistry/material (NCM vs LFP). Clustered (dyad) standard-errors in parentheses, Signif. Codes: a: 0.01, b: 0.05, c: 0.1.

Table B2: Pack Sourcing decision

Dependent Variable:	Pack sourcing decision				
Model:	(1)	(2)	(3)	(4)	(5)
Border	-3.74[a]	-3.84[a]	-3.84[a]	-3.98[a]	-3.49[a]
	(0.898)	(0.919)	(0.919)	(0.973)	(0.774)
log distance	-0.231[a]	-0.215[a]	-0.216[a]	-0.231[a]	-0.283[a]
	(0.018)	(0.021)	(0.016)	(0.015)	(0.042)
RTA	0.965[b]	1.02[b]	1.02[b]	1.02[b]	0.986[a]
	(0.475)	(0.481)	(0.481)	(0.504)	(0.381)
EC(module)	-0.067	-0.108[b]	-0.108[b]	-0.049	-0.189[c]
	(0.052)	(0.047)	(0.047)	(0.051)	(0.102)
Intrafirm		0.601[a]	0.601[a]	0.983[a]	0.659[b]
		(0.093)	(0.097)	(0.119)	(0.263)
Intraplant			-0.006	-0.118	-0.146
			(0.173)	(0.213)	(0.346)
Observations	16,005	16,005	16,005	10,591	8,190
Squared Correlation	0.320	0.326	0.326	0.347	0.339

Column (4) constrains the set of choosers to be car models with positive sales (strictly defined). Column (5) constrains the set of choosers to be owned by one of the top 13 manufacturers. Clustered (dyad) standard-errors in parentheses, Signif. Codes: a: 0.01, b: 0.05, c: 0.1.

Table B3: BEV Sourcing decision

Dependent Variable:	BEV sourcing decision				
Model:	(1)	(2)	(3)	(4)	(5)
Border	-1.73[a]	-1.72[a]	-1.88[a]	-1.72[a]	-1.84[a]
	(0.437)	(0.459)	(0.394)	(0.460)	(0.412)
log distance	-0.069	-0.069	-0.013	-0.069	-0.104
	(0.053)	(0.054)	(0.051)	(0.054)	(0.071)
RTA	0.863[a]	0.833[a]	1.01[a]	0.833[a]	1.13[a]
	(0.171)	(0.177)	(0.163)	(0.177)	(0.225)
EC(pack)		-0.485[a]	-0.424[a]	-0.485[a]	-0.049
		(0.092)	(0.053)	(0.093)	(0.084)
Observations	54,696	54,696	54,933	54,933	36,107
Squared Correlation	0.223	0.226	0.241	0.244	0.243
FE origin	ctry	ctry	cluster1	cluster2	cluster2
Sample	all	all	all	all	NRdest

Clustered (dyad) standard-errors in parentheses, Signif. Codes: a: 0.01, b: 0.05, c: 0.1.