Modeling Demand in the Electric Vehicle Lithium-Ion Battery Market

Alexandra Paulus

April 07, 2025

Abstract

The electric vehicle lithium-ion battery market is highly relevant to decarbonization and electrification efforts. It is also complex and dynamic, involving differentiated products and within-manufacturer substitution patterns. In this study, battery prices are systematically estimated as a function of raw material prices; subsequently, a nested logit demand model is built and run in order to better understand the factors that drive consumer-level choices and market-level welfare. It is found that prices and energy density of batteries are major drivers of consumer choices and that within-manufacturer substitution patterns, likely derived from supply agreements, play a key explanatory role. Counterfactual analyses reveal that the introduction of lithium iron phosphate batteries has had a large, positive effect on consumer welfare, while the stoichiometric evolution of lithium nickel manganese cobalt batteries has had a positive, moderate effect.

1 Introduction

The electric vehicle (EV) revolution is upon us. In the U.S. alone, the number of EVs on the road grew from approximately 22,000 in 2011 to over 2 million in 2021, according to the U.S. Bureau of Labor Statistics. This can be attributed to a combination of interplaying factors: decreasing supply of fossil fuels, rising concerns about the rate of greenhouse gas emissions, increased government incentives and subsidies to spur the production and adoption of EVs, and improved technology and greater variety of models. However, this didn't occur in isolation. The EV market has largely been driven by the EV lithium-ion battery market. This is because automotive original equipment managers (OEMs) entering the EV space have, at least up until this point, lacked the capacity to manufacture their own lithium-ion batteries, instead purchasing directly from battery manufacturers. Demand for EV batteries has grown approximately fourteen-fold from 51.5 GWh in 2016 to 772 GWh in 2023, according to Statista.

In and of itself, the EV lithium-ion battery market is highly dynamic. In tandem with a sharp increase in demand, the market has experienced a sharp decline in prices—by 2023, prices had dropped to 18% of their 2013 levels, according to BloombergNEF. It is a truly international market—of the ten largest battery manufacturers, five are Chinese, three are Korean, and two are Japanese, and the key EV OEMs that purchase from them are multinational companies. The state of the market depends heavily on global, volatile supply chains, due to the involvement of mining and processing of critical materials. Furthermore, EV battery manufacturing is receiving growing support from governments, many of which seek to simultaneously increase and nearshore the production of cathode and anode active materials, the batteries themselves, and the EVs they're incorporated into.

On a more micro level, the EV market is particularly complex for two main reasons. Firstly, it consists of differentiated products: each major battery manufacturer produces some combination of lithium nickel manganese cobalt (NMC) batteries, lithium nickel cobalt aluminum (NCA) batteries, and lithium iron phosphate (LFP) batteries. These products have innately different characteristics based on their chemical compositions. NMC and NCA batteries tend to have higher energy and power densities than LFP batteries, meaning that they can deliver more power and be charged at a faster rate; on the other hand, LFP batteries are typically cheaper, safer, and have longer cycle lives. However, even within a chemistry type, there is room for further differentiation, through innovation and optimized design and manufacturing processes. Secondly, the market relies at least partly on supply agreements between battery manufacturers and EV OEMs, which can take the form of more flexible contracts or more rigid joint ventures. An implication of this is that EV OEMs might be bound to, or have more loyalty towards, a certain manufacturer's products, which can affect substitution patterns in the market.

Gaining a better understanding of how the EV battery market operates, specifically the factors that drive consumer utility and purchasing choices, is important for a number of reasons. Firstly, it may guide future decisions by battery manufacturing firms to make themselves more competitive, as well as guide future policy decisions to spur growth on an industry-wide level. Secondly, the uniqueness of the EV battery market has two implications: (1) questions about the factors that drive it can't be simply answered using previous studies of other technologies (e.g., semiconductors, solar PVs), and (2) answering them would contribute significantly to our understanding of the economics of other technologies that are also emerging, especially those that are linked to decarbonization or electrification efforts and/or rely on global supply chains. Thirdly, given that the EV market is deeply connected to the EV lithium-ion battery market, better understanding the latter would allow us to build more nuanced, realistic models of the former, as well as to better predict the trajectory of the EV revolution.

In this study, the EV battery market is modelled using a nested logit demand model. Nested logit models are used to model consumer choice behavior in differentiated product markets with flexible substitution patterns. Firstly, on the micro level, parameters that explain how price and other product characteristics affect consumer choices to buy certain products will be estimated, as well as corresponding price and characteristic elasticities of demand. Secondly, on the macro level, counterfactual scenarios will be modeled and analyzed in order to understand what industry-wide trends have contributed most to consumer welfare.

The rest of the study is organized as follows. Section 2 outlines the current state of the literature on the EV lithium-ion battery market and explains how this study fills in some of its gaps. Section 3 provides an explanation of the nested logit demand model, and Section 4 outlines the empirical framework. Section 5 reviews the data used in the study, and Section 6 outlines the estimation method. Section 7 goes through the Results, and Section 8 provides a Conclusion.

2 Electric Vehicle Lithium-Ion Battery Market

Demand for EV batteries has grown substantially, from 51.5 GWh in 2016 to 772 GWh in 2023, according to Statista. This can be directly linked to the rapid growth of general EV demand. Although EVs were invented in the late 19th century, they remained commercially dominated by gasoline-powered cars from 1908, with the introduction of Henry Ford's Model T, through the end of the 20th century. However, EVs reentered the market in the early 2000s, with the release of the Toyota Prius in 2000 and the announcement of a Tesla Motors luxury electric sports car in 2006, and demand for them has been growing rapidly since.

Simultaneously, EV battery prices have fallen significantly, from \$780/kWh in 2013 to \$139/kWh in 2023, according to BloombergNEF. Indeed, much of the existing literature on the EV battery market has focused on causes of cost declines, and, by extension, price declines. One camp has done so by applying a bottom-up approach, through building granular cost models. Ziegler, Song, and Trancik (2021) investigate the causes of cost declines from the late 1990s through the early 2010s and find that public and private R&D, especially regarding cell chemistries and electrode materials, were key drivers, notably to a greater extent than economies of scale. Chung, Elgqvist, and Santhanagopalan (2016) apply a similar approach to six representative firms—a U.S startup, a U.S. transplant in Korea, a Japanese firm, a Korean firm, a more established Chinese firm, and a less established Chinese firm—and find that differences in material and labor costs are the main contributors to differences in production costs. Orangi, et al. (2024) forecast that battery costs could fall between 43.5% and 52.5% by 2030, with cathode material cost reductions and scrap rate minimizations likely to be the primary drivers. Different studies have yielded vastly different findings, which can likely be attributed to the fact that the accuracy and uniformity of inputs are limited by data availability across the industry, and that such bottom-up studies rely heavily on assumptions about how variables should be weighted.

A second camp has investigated causes of cost declines by correlating cost with cumulative production or research investment levels and quantifying the rate of production improvement; typically, this is reported as the decrease in cost that corresponds to each doubling of production volume, i.e., the "learning rate." Nykvist and Nilsson (2015) determine a learning rate of 6-9% from 2007-2014. Hsieh, Pan, Chiang, and Green (2019) focus on a two-level learning curve (i.e., one that takes raw material costs as the lower bound) and find that as batteries continue to become cheaper, materials costs will become increasingly important in limiting further production cost decreases. A significant issue with these studies is that correlation doesn't imply causation—the causes that drive cost decline (whether learning-by-doing, economies of scale, or R&D investment) are assumed beforehand and can't be rigorously determined. Furthermore, all of the aforementioned cost-focused studies are market-wide—they don't tell us anything about within-market dynamics, such as how cost differences across firms affect each of their demand levels.

A smaller subset of studies have focused on analyzing market dynamics, beyond these two overarching trends of demand growth and price declines. However, such studies typically analyze one factor at a time and focus on the extended impact on the broader EV market, rather than the EV battery market specifically. One key driver of year-to-year dynamism in the battery market is the volatility of raw material prices, which affects different battery manufacturers differently, based on the makeup of their product portfolios; for example, a manufacturer that uses more materials with more upwardly volatile prices will likely incur greater costs and be forced to pass them onto consumers via pricing. The effect of raw material prices on the EV battery market is so pronounced that they secondarily affect the EV market—Cagli (2023) finds that there exists significant spillover of volatility from raw material prices to EV stock returns (as battery costs are passed through to battery prices, which become EV costs and are passed through to EV prices), predominantly from base metals in the short-run and lithium and manganese in the long-run. Indeed, Mauler, Lou, Duffner, and Leker (2022) determine that from now through 2030, the negative effect of technological innovation on cell costs might be fully outweighed by the positive effect of tightening raw material markets.

Secondly, while the industry is trending sharply upwards in terms of innovation, some companies make larger year-to-year marginal improvements than others, based on the extent and quality of their R&D, as well as the characteristics they choose to prioritize. For example, CATL incorporated innovative cell-to-pack technology into its LFP battery in 2019, while BYD was only able to do so in 2020. This enabled CATL to launch a battery with higher energy density one year earlier and, as a direct result, expand its partnership with VDL Bus & Coach. Zhu, Liu, Liu, and Yan (2022) investigate the effect of government subsidies for EV battery R&D by constructing a two-tier supply chain consisting of an incumbent EV OEM and a competitive EV OEM that purchase from the same battery supplier. They find that R&D subsidies increase innovation and decrease production costs for the battery manufacturer, which, in turn, encourages more EV OEMs to enter the market; however, the resultant price competition among EV OEMs may actually offset profit gains for both the battery supplier and the OEMs themselves.

Thirdly, the location of the manufacturer can affect how readily EV OEMs purchase from it in a given year, depending on government incentives and subsidies, supply chain disruptions, and geopolitics, among other factors. For example, from 2009 to 2022, the Chinese government invested 200 billion RMB ($29 billion) into subsidies and tax breaks across companies involved in the EV supply chain, according to the MIT Technology Review. These incentives likely made Chinese battery manufacturers more appealing to EV OEMs globally, not only because they lowered production costs but also because they likely served as an indicator of future innovation and improvements in product quality among these firms. Additionally, these incentives catalyzed the broader Chinese EV market, which an EV OEM would likely gain better access to by purchasing more from a Chinese battery supplier. However, in 2022, the U.S. passed the Inflation Reduction Act, which provides up to $7,500 in tax credits for EV OEMs that sourced batteries not from a "foreign entity of concern," including China. Less explicitly, the EU Batteries Regulation, which passed in 2023, requires all manufacturers that sell batteries in the EU to regularly release information on, and eventually meet targets related to, carbon footprint and ethical sourcing of raw materials, which will be more difficult for Chinese companies to adhere to, based on the nature of their supply chains. A few studies have analyzed the effects of such subsidies on the market at large. Hu, Yin, and Zhao (2015), for example, build a structural demand model to analyze the effect of the battery electric vehicle subsidy program in China on the EV market (not the EV battery market), specifically on current and future price elasticities, a demand-side peer effect, and supply-side changes in production costs. They emphasize the use of a dynamic (rather than static) model, given that both prices and product characteristics change over time, and consumers have the choice to buy electric vehicles now and enjoy higher subsidies or later and enjoy higher-quality products. They find that the static model overestimates demand elasticity, relative to the dynamic model, and that a phase-out policy is preferrable to other policies that would offer larger subsidies over longer periods of time, in terms of cost minimization and sales maximization.

As mentioned in Section 1, the EV battery market is uniquely structured because it is partly built on supply agreements and joint ventures between EV OEMs and battery manufacturers. Supply agreements are contracts that specify the volume of batteries that the battery manufacturer will supply to the EV OEM and the timespan over which it will do so; some contracts are more flexible than others, with regards to the chemistry type supplied and the rigidity of the terms (e.g., if either party wants to change the volume supplied). Joint ventures involve an EV OEM building a new plant and licensing a battery manufacturer's technology, often specifying a certain chemistry type beforehand (which can, however, change over time). EV OEMs also make one-time or recurring purchases from battery manufacturers to supplement their existing agreements and ventures. As of 2022, according to a Frost & Sullivan report, key supply agreements existed between LG Energy and Isuzu (four-year agreement to supply cylindrical batteries) and Samsung SDI and Hyundai (seven-year agreement to supply prismatic batteries); key joint ventures existed between CATL and Ford (Michigan plant with 35 GWh capacity), CATL and Tesla (U.S. plant with 80 GWh capacity), LG Energy and General Motors (3 U.S. plants with 130 GWh total capacity), and Samsung SDI and Stellantis (Indiana plant with 23 GWh capacity). A few studies have analyzed the effect of such cooperation on the market at large. Fan, Huang, and Wang (2021), for example, investigate the effect of joint ventures by constructing a supply chain consisting of a battery supplier, a premium brand EV OEM, and a normal brand EV OEM; modeling three strategies (a joint venture between the supplier and premium OEM, joint venture between the supplier and the normal OEM, and no joint ventures); and analyzing the demand and profit of each OEM. They find that when OEMs engage in joint ventures, their demand and profit increase significantly, but there is no substantial effect on overall profitability of the supply chain.

This study fills gaps in the literature in three ways. Firstly, it analyzes the effects of many price and product characteristics on the EV battery market simultaneously, rather than one-by-one, as many of the aforementioned studies have. This is a more realistic approach, given that EV OEMs observe all of these factors simultaneously and make their purchasing decisions accordingly. Secondly, this study focuses specifically on the EV battery market, in which the producers are battery manufacturers and the consumers are EV OEMs, rather than focusing purely on the supply side (i.e., how manufacturing costs translate to battery prices) or the demand side (i.e., how factors are passed through to the EV market). Thirdly, this study uses a company-specific approach, i.e., incorporates real, company-specific data on production and product characteristics, rather than following a bottom-up or industry-wide approach, either of which would require heavy reliance on assumptions and industry-wide estimates throughout the analysis.

3 Nested Logit Demand Model

This study uses the nested logit model outlined by Berry (1994) and Cardell (1997).

Suppose there are $t = 1, 2, \ldots, T$ markets. In each market, there are $j = 1, 2, \ldots, J_t$ products produced by $f = 1, 2, \ldots, F_t$ firms, which sum to a total of N products across all markets.

Take a consumer i who purchases product j in market t. The indirect utility of that consumer can be expressed as:

$$U_{ijt} = \alpha p_{jt} + \beta x_{jt} + \xi_{jt} + \epsilon_{ijt},$$

where p_{jt} is the price of product j in market t, x_{jt} is a K-dimensional vector of observable product characteristics for product j in market t, ξ_{jt} is the valuation of the unobserved product characteristics of product j in market t, and ϵ_{ijt} is a random error term that accounts for other unobserved factors that affect each consumer's preference for product j in market t. Put more simply, the first and second terms of the indirect utility formulation make up the deterministic component of utility, i.e., that determined by prices and characteristics observed by the econometrician; the third and fourth terms make up the stochastic component of utility, i.e., that determined by characteristics unobserved by the econometrician and differences unique to each consumer. α is the estimated price sensitivity parameter, with a high $|\alpha|$ indiciating high price sensitivity and a low $|\alpha|$ indicating low price sensitivity. β is the estimated vector of coefficients on product characteristics, with a high β_k indicating that consumers have a preference for products with (more of) characteristic x_k.

Suppose, also, that products within a nest h are correlated, i.e., are better substitutes for each other than for products outside of that nest. The random error term can thus be rewritten as:

$$\epsilon_{ijt} = \xi_{ih(j)t} + (1 - \rho) \times \bar{\epsilon}_{ijt},$$

where $\xi_{ih(j)t}$ is the idiosyncratic error specific to (or, the idiosyncratic preference for all products in) each nest h, which is a function of j; $\bar{\epsilon}_{ijt}$ is the remaining idiosyncratic error (unexplained by nests) specific to each product for each individual; and ρ is a nesting parameter that controls the extent of substitution within each nest h. Notably, ρ only appears in front of the product-specific idiosyncratic error term and not the nest-specific idiosyncratic error term because it is already implicit in the latter—its distribution function depends on ρ. In the limit of $\rho \to 1$, consumers only substitute between products that are in the same nest; in the limit of $\rho \to 0$, consumers don't substitute within-nest any more than they do between nests, and the model collapses to the simpler multinomial logit model.

Thus, we can rewrite indirect utility as:

$$U_{ijt} = \alpha p_{jt} + \beta x_{jt} + \xi_{jt} + (1 - \rho) \times \bar{\epsilon}_{ijt},$$

and, for simplicity's sake, we can aggregate the first three terms into a single term that represents the average utility for product j in market t:

$$V_{jt} = \alpha p_{jt} + \beta x_{jt} + \xi_{jt}.$$

We assume there is an outside good, $j = 0$. The mean utility of the outside good is forced to be zero, so that the indirect utility derived from purchasing it is just the associated random error term:

$$U_{ijt} = \epsilon_{i0t}.$$

The aggregate market share of product j in market t is equal to the probability that consumer i chooses that product, given that all consumers are identical. It can be rewritten as the product of the probability of choosing product j, conditional on being in nest h, and the unconditional probability of choosing nest h:

$$s_{jt} = \bar{s}_{(j|h)t} \times \bar{s}_{ht}.$$

By the definition of conditional probability, the probability of choosing product j, conditional on being in nest h (i.e., $j \in N_h$), can be expressed as the ρ-adjusted exponentiated utility of choosing product j, divided by the ρ-adjusted exponentiated utility of choosing nest h. The former can be expressed as:

$$W_{jt} = \exp\left(\frac{V_{jt}}{1-\rho}\right)$$

and the latter as:

$$W_{ht} = \sum_{k \in N_h} \exp\left(\frac{V_{kt}}{1-\rho}\right).$$

Thus, the probability of choosing product j, conditional on being in nest h, can be expressed as:

$$\bar{s}_{(j|h)t} = \frac{W_{jt}}{W_h t} = \frac{\exp\left(\frac{V_{jt}}{1-\rho}\right)}{\sum_{k \in N_h} \exp\left(\frac{V_{kt}}{1-\rho}\right)}.$$

Note that the use of exponentiated utilities and the way that each is weighted by ρ is a function of the assumption that the error term is distributed IID with the Type I Extreme Value (Gumbel) distribution.

Then, the unconditional probability of choosing nest h can be expressed as the ρ-adjusted exponentiated utility of choosing nest h, divided by the sum of the ρ-adjusted exponentiated utilities of all nests in the estimation:

$$\bar{s}_{ht} = \frac{(W_{ht})^{(1-\rho)}}{\sum_h (W_{ht})^{(1-\rho)}}.$$

By multiplying and canceling like terms, the aggregate market share of product j in market t can be expressed as:

$$s_{jt} = \frac{W_j t}{W_h t} \times \frac{(W_{ht})^{(1-\rho)}}{\sum_h (W_{ht})^{(1-\rho)}} = \frac{W_{jt}}{W_{ht}^{\rho} \sum_h (W_{ht})^{(1-\rho)}}.$$

Given that the mean utility of the outside good is normalized to zero (i.e., $V_{j0t} = 0$), and that it may be considered the only product in its nest ($h = 0$), both $W_{(j=0)t}$ and $W_{(h=0)t}$ go to 1. Thus, the aggregate market share of the outside good in market t can be expressed as:

$$s_{0t} = \frac{1}{\sum_h (W_{ht})^{(1-\rho)}}.$$

Directly following the approach in Berry (1994), we can calculate logs of the aggregate market shares of product j and the outside good in market t and take their difference to yield:

$$\log(s_{jt}) - \log(s_{0t}) = \frac{V_{jt}}{1-\rho} - \rho \log(W_{ht}).$$

We can then take the log of the unconditional probability of choosing nest h, $\bar{s}_{ht}$, to yield an expression for the ρ-adjusted exponentiated utility of choosing nest h, W_{ht}:

$$\log(W_{jt}) = \frac{\log(\bar{s}_{ht}) - \log(s_{0t})}{1-\rho}.$$

With substitution into prior equations and consolidation of terms into the probability of choosing product j, conditional on being in nest h, $\bar{s}_{(j|h)t}$, we yield:

$$V_{jt} = \log(s_{jt}) - \rho \times \log(\bar{s}_{(j|h)t}) - \log(s_{0t}).$$

Finally, we can substitute in for V_{jt} and rearrange to get:

$$\log(s_{jt}) - \rho \times \log(\bar{s}_{(j|h)t}) - \log(s_{0t}) = \alpha p_{jt} + \beta x_{jt} + \xi_{jt}.$$

Thus, values for α, β, and ρ can be estimated by regressing differences in log market shares and within-group shares on prices and product characteristics.

4 Empirical Framework

4.1 Marginal Cost and Price Prediction Model

Any demand model requires, at a minimum, two inputs: demand levels and prices. However, prices for each battery product are not observed, for a number of reasons. Actual market prices don't exist, as EV battery manufacturers sell a large proportion of their batteries as part of fixed supply contracts, which may set different prices depending on the terms of the contract. Still, EV battery manufacturers don't report any contract-specific prices, nor any kind of suggested or proxy or average price, likely due to the fact that contract information is typically proprietary, as well as the fact that it is not competitively advantageous for a firm to reveal its pricing strategy to other firms operating in the same market.

Marginal costs are also not observed. The majority of EV battery manufacturers operate in multiple lithium-ion battery segments (most commonly for EV applications, energy storage applications, and electronic devices), and some operate in wholly different industries (for example, BYD Company and Panasonic are considered multinational conglomerates). Although these kinds of companies typically report their annual cost of goods sold (COGS) on the basis of being public, they often don't break out their COGS by segment nor industry, making it impossible to reliably extract information about the marginal costs of producing an EV lithium-ion battery, let alone an LFP vs NMC vs NCA battery.

Therefore, in this study, the empirical framework used consists of two parts: a price prediction model, followed by a full demand model. The price prediction model involves (1) estimating the raw materials component of marginal cost for each chemistry type produced by Chinese versus rest-of-world firms, (2) factoring in reported industry-wide cell and pack manufacturing costs to yield battery pack marginal costs, and (3) estimating markups of price over marginal costs for Chinese versus rest-of-world firms and factoring those in, too.

Step (1) involves calculating the cost per kg of the cathode, anode, and electrolyte (the choice of which is represented by i) for a certain chemistry type c, according to the formula:

$$\text{per kg cost of component}_{i_c} = \sum_{m \in i_c} (\text{percent composition}_m \times \text{per kg raw material price}_m)$$

where m is the raw material and could represent lithium, aluminum, cobalt, iron, manganese, nickel, or phosphorus when i = cathode; graphite when i = anode; and lithium hexafluorophosphate when i = electrolyte.

For each cathode type under consideration, percent compositions by mass can be easily calculated by first determining the stoichiometric composition by element (e.g., the NMC(111) cathode has a stoichiometric composition of 25% lithium, 8.3% nickel, 8.3% manganese, 8.3% cobalt, and 50% oxygen) and subsequently using molar masses to calculate material composition by % of total weight. Notably, the anode is considered to be made of 100% graphite, and the electrolyte is considered to be made of 100% lithium hexafluorophosphate, as is typical across the industry.

Subsequently, the cost per kg of battery is calculated according to the formula:

$$\text{per kg cost of battery}_c = \sum_i (\text{percent of battery mass}_{i_c} \times \text{per kg cost}_{i_c})$$

where the first term in the summation represents the component's contribution to the total mass of the battery, and the second term is what was calculated directly above. Finally, the cost per kWh of battery is calculated as per kg cost of battery$_c$ × energy density$_c$, where the latter is a market-wide value for all batteries of the same chemistry type.

Step (2) is dependent on data from the literature; industry-wide cell and pack manufacturing costs are added uniformly to all products, regardless of chemistry type or manufacturer.

Step (3) involves the calculation of the ratio of prices over marginal costs for companies headquartered in China versus other countries. To do so, the difference in sensitivity of demand to marginal cost (i.e., the difference in the effect of a 1 unit increase in marginal cost on demand) between an average Chinese and average rest-of-world manufacturer is measured. The rationale behind this is that demand tends to be more elastic at higher prices, meaning that the same $1 increase in marginal cost would induce greater substitution from an already highly priced product, compared to a lower priced one. An average Chinese firm and an average rest-of-world firm are used in the analysis, because, based on the way marginal costs are constructed (i.e., as a function of material prices), each company in each region is assumed to face the same marginal costs per chemistry type. The demand for a certain battery product in a certain year manufactured by the average Chinese firm is calculated by taking the average of the sales volumes for that product in that year across all Chinese firms in the dataset. The demand for an average rest-of-world firm is calculated analogously.

To do this, the following regression is carried out:

$$\text{Demand}_{ict} = \beta_0 + \beta_1 \times \text{Marginal Cost}_{ict} + \beta_2 \times (\text{Marginal Cost}_{ict} \times \text{Group}) + \alpha_c + \alpha_t + \epsilon_{it}$$

Here, Demand_{ict} is the sales volume of a given chemistry type c by a given firm i in a given year t, in GWh. It is regressed on both the marginal cost of producing this good and an interaction term. The interaction term is defined as the product of the marginal cost of producing this good and a group dummy, which equals 1 for firms headquartered in China and 0 for firms headquartered elsewhere; thus, the interaction term is only nonzero for Chinese firms.

For the average rest-of-world company, the regression specification becomes:

$$\text{Demand}_{ict} = \beta_0 + \beta_1 \times \text{Marginal Cost}_{ict} + \alpha_c + \alpha_t + \epsilon_{it}$$

For the average Chinese company, the regression specification becomes::

$$\text{Demand}_{ict} = \beta_0 + (\beta_1 + \beta_2) \times \text{Marginal Cost}_{ict} + \alpha_c + \alpha_t + \epsilon_{it}$$

β_1 can be interpreted as the marginal cost sensitivity of an average rest-of-world company, while $\beta_1 + \beta_2$ can be interpreted as that of an average Chinese company. Thus, the value of β_2 that is output by the regression can be interpreted as the difference in the marginal cost sensitivities of the two companies. Given that β_1 is expected to be negative, if $\beta_2 > 0$, then the average Chinese company is less price sensitive than the average rest-of-world company, and vice versa if $\beta_2 < 0$.

If the difference in the marginal cost elasticities of demand of the two firms is significant, it can be confidently attributed to differences in the extent to which they mark up their prices over their raw material costs. This is due to the inclusion of various fixed effects, which account for other potentially confounding factors. Chemistry fixed effects, α_c, are included to account for inherent differences in demand levels and patterns for LFP versus NMC versus NCA batteries. Time fixed effects, α_t, are included to account for the time-varying nature of EV technological development, demand for EV batteries, and subsidies and incentives for EV battery production. Notably, a group effect (dummy variable for Chinese firms, for example) is not included because it is assumed that all relevant region-specific effects are included in the β_2 term. Key differences between operating in China versus other countries include lower labor costs, lower real estate costs, lower supply chain costs, and, in general, greater subsidies and incentives for EV battery production, all of which factor into a lower markup of price over raw material marginal costs. It is possible that, over the last decade, some firms have become more wary of purchasing from China, due to geopolitical and supply chain risks, meaning they have become more willing to substitute from a Chinese firm to a rest-of-world firm (which would drive up own-price elasticity of Chinese firms) or less willing to substitute from a rest-of-world firm to a Chinese firm (which would drive down own-price elasticity of rest-of-world firms). However, if this time variation in general consumer sentiment is indeed present, it will be controlled for by the time fixed effects.

Thus, as a crude estimation, the ratio of $(\beta_1 + \beta_2)$ to β_1 can be interpreted as the ratio between the markup of price over marginal cost of Chinese companies versus of rest-of-world companies. Given that, according to Hall (2018), the estimated average markup of price over marginal cost for all large firms in a given industry (over 10,000 employees) is 30%, the price markup of the average rest-of-world firm over marginal costs, x_{RoW}, can be calculated using the formulation:

$$x_{RoW} = \frac{0.3}{s_{\text{Chinese}} \times (\frac{\beta_1+\beta_2}{\beta_1}) + (1 - s_{\text{Chinese}})},$$

where s_{Chinese} is the share of total production by Chinese firms (which is used as a weight). The price markup of the average Chinese firm can be calculated as:

$$x_{Chinese} = \frac{\beta_1 + \beta_2}{\beta_1} \times x_{RoW}.$$

4.2 Nested Logit Demand Model

A logit demand model is constructed in which an EV OEM i purchases a battery product j from manufacturer f (e.g., purchases CATL's BYD battery, or Panasonic's NCA battery) in market t, which is defined as every year from 2015-2023. Nests are introduced and defined per manufacturer, in order to capture the effect of supply agreement and joint ventures on substitution patterns. Specifically, EV OEMs can be thought of as tending to be more willing to substitute across products produced by a single battery manufacturer with which it already has a supply contract, before considering products produced by other manufacturers.

Three observable product characteristics are considered: energy density, the number of new patents filed by the battery manufacturer, and whether the battery manufacturer is based in China. Energy density is typically viewed as the key metric for EV battery performance. It directly dictates the amount of energy that can be stored by a battery of a given weight, which, in turn, can define characteristics such as power and size/weight efficiency. Importantly, it is also directly correlated with other characteristics that consumers care about but that battery manufacturers rarely report data on, such as charging rates, safety, and cyclability. For example, it is the nickel-based layered structure of NMC cathodes that gives it its high energy density—it allows for high degrees of lithium ion mobility and intercalation/deintercalation, as the ions can move in two dimensions between layers of transition metal oxides, which improves power storage and charging/discharging rates. However, this layered structure is also inherently more unstable, as substantial intercalation/deintercalation can introduce lithium impurities and induce irreversible structural changes to the crystal structure, which increases thermal instability (i.e., fire risk) and reduces cyclability. The number of new patents is a proxy for firm-level innovation and investment into R&D, which gives an indication of future improvements of characteristics that EV OEMs care about, as well as general longevity and competitiveness of the firm. Whether the firm is based in China is a variable that captures the benefits or drawbacks (unrelated to differences in prices set by Chinese versus rest-of-world firms, assuming that prices are scaled in the way described in Section 4.1) of purchasing from a Chinese firm, based on geopolitics, supply chain disruptions, and government policies.

Notably, the model doesn't incorporate any heterogeneity on the individual consumer level in terms of tastes for price and product characteristics, i.e., there are no random coefficients in utility function. The model is constructed in this way because data on individual EV OEM purchasing decisions cannot be accessed—only aggregate market shares are observed—and it is reasonable to assume relatively homogeneous preferences across consumers, especially about characteristics such as energy density and safety (EV OEMs will want to maximize both). As a result, the model estimates a single value for each of the parameter α and K elements of vector β.

The indirect utility of EV OEM i purchasing in the battery market can be written as follows:

$$U_{ijt} = \alpha p_{jt} + \beta x_{jt} + \xi_{jt} + \xi_{if(j)t} + (1-\rho) \times \bar{\epsilon}_{ijt}$$

where $\xi_{if(j)t}$ represents the idiosyncratic error specific to each manufacturer nest, f. βx_{jt} is broken out in the following way:

$$\beta x_{jt} = \beta_1(\text{energy density})_{jt} + \beta_2(\text{firm's number of patents})_{f(j)t} + \beta_3 \mathrm{D}_{f(j)}$$

where $D_t = 1$ if the firm is headquartered in China and $D_t = 0$ if the firm is not; this value is considered fixed for each firm over time.

The average utility for product j in market t can be written as follows:

$$V_{jt} = \alpha p_{jt} + x_{jt}\beta + \xi_{jt}.$$

Now, the unobserved product characteristics are discussed. Such characteristics can be thought of as falling under three distinct categories: firm-specific, year-specific, and chemistry-specific.

Unobserved firm-specific product characteristics are directly linked to supply contracts. Supply contracts lead EV OEMs to have a greater preference for certain battery manufacturers for reasons that are additional to the observable firm-level characteristics included in β (i.e., whether the firm is in China, and how many patents the firm has filed). For example, an EV OEM might have chosen to sign a 15-year supply contract with a battery manufacturer 5 years ago based on, for example, the strength of its patent portfolio 5 years ago. Even if its portfolio has weakened since then, relative to competitors' portfolios, the EV OEM would still be constrained to purchasing from that battery manufacturer or would prefer to, based on loyalty and compatability of that manufacturer's batteries with its own products. Unobserved year-specific effects are partly linked to supply contracts, which are established in a certain year and affect the market in the years thereafter. However, such effects are also linked to other external shocks that may drive an EV OEM to purchase more batteries, such as a raise in government subsidies for automobile OEMs that sell X number of EVs. The way that the nests are structured, i.e., by firm and year, effectively controls for both firm- and year-specific effects—the model estimates a value for ρ, which represents the degree of correlation between products in a nest that, by definition, all share the same firm- and year-specific unobserved characteristics. Therefore, there is no need to explicitly include firm fixed effects and year fixed effects.

Unobserved chemistry-specific effects are ones that would cause an EV OEM to purchase a certain battery chemistry for reasons that are seemingly unrelated to energy density (and inherently related characteristics, such as safety and lifespan), and rather are, for example, specific to the EV models they produce. For example, an EV OEM may produce an EV model that requires batteries of a certain shape or size that are more common for one chemistry type than for another. Such preferences have no systematic relationship with any of the observable product characteristics and are hence left as part of the error term.

Now, the outside good is discussed. In the EV battery market, EV OEMs don't truly have the option to purchase an outside good, as it is typically defined. Other batteries suitable for EVs, such as sodium-ion and solid-state batteries, are still emerging and not fully commercialized. Given that EV OEMs don't manufacture their own batteries, they must purchase batteries every year in order to continue manufacturing and selling their own EV products. Hence, here, the outside good is considered to be a battery produced by any other major players that are not considered in the analysis but are in the EV battery space. Practically, the outside good isn't explicitly modeled out, but it is present based on the fact that in each market, the market shares sum to less than one (as they sum to one when the shares of all other firms are included).

Following the model presented in Section 3, a linear equation is derived. From this, two regression specifications are developed. The first only has price and a single observed product characteristic—energy density—as explanatory variables:

$$\log(s_{jt}) - \rho \times \log(\bar{s}_{(j|h)t}) - \log(s_{0t}) = \alpha p_{jt} + \beta_1(\text{energy density})_{jt} + \xi_{jt}.$$

The second also includes the firm's number of patents and a China dummy as observed product characteristics:

$$\log(s_{jt}) - \rho \times \log(\bar{s}_{(j|h)t}) - \log(s_{0t}) = \alpha p_{jt} + \beta_1(\text{energy density})_{jt} + \beta_2(\text{firm number of patents})_{f(j)t} + \beta_3 \mathrm{D}_{f(j)} + \xi_{jt}.$$

The second regression is run after the first in order to break out observed firm-level characteristics and unobserved ones related to supply agreements and analyze them separately. Additionally, doing so allows us to confirm whether breaking these out has an effect on incremental robustness and model fit.

For a product j in market t, its own-price elasticity, i.e., the change in the market share of product j, s_{jt}, when its own price, x_{jt}, increases by 1%, can be calculated as:

$$\epsilon_{jt} = \frac{x_{jt}}{s_{jt}} \frac{\partial s_{jt}}{\partial x_{jt}}.$$

For products j and k in market t, their cross-price elasticity, i.e., the change in the market share of product j when the price of product k increases by 1%, can be calculated as:

$$\epsilon_{jt} = \frac{x_{kt}}{s_{jt}} \frac{\partial s_{jt}}{\partial x_{kt}}.$$

Consumer surpluses for each market t can also be calculated as a market share-weighted sum of the difference in mean utility and price of each product j in the market:

$$\text{CS}_t = \sum_j [(U_{jt} - p_{jt}) \times s_{jt}].$$

Aggregate consumer surplus is defined as the sum of consumer surpluses for each market. Average consumer surplus is defined as an unweighted average of the consumer surpluses for each market.

4.3 Instruments

Given the simultaneity of price and demand, there exists an inherent endogeneity problem; specifically, prices may be correlated with the error term, which represent unobserved factors (product characteristics) that provide shocks to demand. As a result, instruments for price must be included in order to isolate exogenous variation in prices. For each battery chemistry, four instruments are included. Two are the same across all chemistries: current global lithium carbonate prices, and one-year lagged global lithium carbonate prices. Current lithium iron phosphate and one-year lagged iron ore prices are included as instruments for LFP batteries, current and one-year lagged nickel ore prices are included for NMC batteries, and current and one-year lagged aluminum ore prices are included for NCA batteries. Using a mixture of current and lagged prices as instruments serves to balance relevance (current raw material prices are much more closely correlated with current battery prices than lagged raw material prices are) with minimizing endogeneity (lagged raw material prices aren't at risk of being affected by current EV battery demand in the same way that current raw material prices are).

5 Data

The data consists of market shares, prices, and product- and firm-level characteristics for 11 different EV lithium-ion battery products (either NMC, NCA, or LFP) sold in the market from 2015-2023. This translates to 90 product-year data points in total.

5.1 Market Shares

Firstly, six firms with sufficient data on both total annual sales volumes and breakdowns of annual production by chemistry type are identified: CATL, BYD, LG Energy, Panasonic, SK On, and Samsung SDI. These also happen to be the six largest firms in the market, collectively accounting for 87.2% of market share over the years considered, on average.

Data on annual sales volumes from 2015-2023 is sourced from Frost & Sullivan reports, which sourced their data directly from data vendor EV Volumes. Data on breakdowns of sales by chemistry type in 2018 is also sourced from Frost & Sullivan, and that in 2023 was sourced from S&P Global Ratings and S&P Mobility. In the years prior to 2018, the same per-company breakdown as in 2018 was assumed; this is a reasonable assumption because most companies were primarily operating in a single battery space (>90% of total production) up until 2018: CATL, LG Energy, SK On, and Samsung SDI in the NMC space; Panasonic in the NCA space; and BYD in the LFP space (81% of total production). Between 2018-2023, three out of six companies (CATL, BYD, and LG Energy Solution) experienced significant changes in the makeups of their product portfolios. CATL and BYD entered the LFP space in 2020, meaning that, for each company starting in the following year, LFP's share of total production increased, and NMC's share decreased accordingly (neither firm ever operated in the NCA space); therefore, the shares of these two chemistry types were assumed to linearly increase and decrease, respectively, from 2021 onwards. Analogously, LG Energy entered the NCA space in 2021, meaning that NCA's share of total production increased, and NMC's share decreased accordingly (it never operated in the LFP space); therefore, the shares of these two chemistries were assumed to linearly increase and decrease, respectively, starting in 2022. The other three companies experienced no such change. Using the compiled data on both annual sales volumes and sales breakdowns by chemistry type, the market share of a given chemistry type, c, manufactured by a given firm, f, in a given year, t, is calculated in the following way:

$$s_{cft} = \frac{(\text{sales volume})_{ft} \times (\text{chemistry share of sales volume})_{ft}}{\sum_f (\text{sales volume})_{ft}}$$

5.2 Prices

Data on the makeups of typical LFP, NMC(111), NMC(622), NMC(811), and NCA cells is sourced from the Battery Manufacturing Cost Estimation (BatPaC) software modeling tool developed by Argonne National Laboratory. Specifically, the tool reports the typical weight (g) of the cathode, anode, electrolyte, foil, separator, terminals, and total cell for each chemistry type. From this, each component (cathode, anode, electrolyte) as a percentage of battery mass can be calculated.

Data on raw material prices for the cathode, anode, and electrolyte is collected from a number of sources. A single global price is assumed for all cathode materials except for lithium, for which lithium carbonate (used in LFP batteries) and lithium hydroxide (used in

NMC/NCA batteries) are considered separately and each broken out into China-specific and global/Asia-wide prices. Daily prices for lithium carbonate, lithium hydroxide, aluminum ore, cobalt ore, iron ore, manganese ore, and nickel ore are sourced from CapitalIQ Pro and averaged per calendar year. Phosphoric acid prices aren't made publicly available and are instead calculated by taking the current Guangxi price, as reported by Shanghai Metals Market, and backcalculating historical prices using the Producer Price Index reported by Federal Reserve Economic Data (FRED). Manganese sulfate prices are directly calculated by using the historical ratio of manganese ore to HPMSM prices of 6:1,000, as estimated by Evolution Capital. Processing costs in $/kg for aluminum ore, cobalt ore, and nickel ore into battery-grade materials are sourced from market reports and explicitly factored into the analysis. The processing cost of iron into iron phosphate is assumed to be largely covered by the cost of phosphoric acid. Graphite prices are sourced from the United States Geological Survey (for 2015-2022 prices) and an IndexBox report on U.S. Graphite (for 2023 prices). Lithium hexafluorophosphate prices are also sourced from CapitalIQ Pro. Full details on data sourcing are included in Table 1.

From 2015-2017, the only type of NMC battery that was mass-produced was NMC(111) (which has a 1 to 1 to 1 stoichiometric ratio of nickel to manganese to cobalt). In 2018, NMC(622) batteries (which has a 6 to 2 to 2 stoichiometric ratio) were introduced, and in 2021, NMC(811) batteries (which has an 8 to 1 to 1 stoichiometric ratio) were introduced. Therefore, from 2015-2017, the cost per kWh of a general NMC battery is calculated as that of an NMC(111) battery; from 2018-2020, it is calculated as the average of the costs of an NMC(111) battery and NMC(622) battery; and from 2021-2023, it is calculated as the average of the costs of an NMC(111) battery, NMC(622) battery, and NMC(811) battery. This is done at a industry-wide level because reliable data on the years in which each manufacturer started producing NMC(622) and NMC(811) batteries and in what proportions of total NMC battery production is not available.

According to a market intelligence report published by Thunder Said Energy, the average cost of manufacturing a battery cell from processed raw materials is reported to be $32/kg. According to BloombergNEF, the average cost of manufacturing a battery pack from battery cells is 26.7% of the cost cell. These two data points are used to carry out step (2) of the marginal cost and price prediction model.

Table 1: Sourcing Data on Raw Material Prices

Raw Material	Component	Chemistry	Price Series	Processing Cost
Lithium Carbonate (China)	Cathode	LFP	Lithium Carbonate EXW China Battery	
Lithium Carbonate (RoW)	Cathode	LFP	Lithium Carbonate Global Average	
Lithium Hydroxide (China)	Cathode	NMC, NCA	Lithium Hydroxide EXW China	
Lithium Hydroxide (RoW)	Cathode	NMC, NCA	Lithium Hydroxide CIF Asia	
Aluminum Ore	Cathode	NCA	LME-Aluminum 99.7% Cash	$1.77/kg
Cobalt Ore	Cathode	NMC, NCA	LME-Cobalt Cash	$6.00/kg
Iron Ore	Cathode	LFP	Iron Ore 62% FE	
Phosphoric Acid	Cathode	LFP	Guangxi Price (Current) + FRED Price Series (Historical)	
Manganese Ore	Cathode	NMC	Manganese Ore Mn32% Fe20%,Tianjin-SA	times factor of 1,000/6
Nickel Ore	Cathode	NMC, NCA	LME-Nickel Cash	$2.00/kg
Graphite	Anode	LFP, NMC, NCA	USGS (2015-2022), IndexBox estimate (2023)	
Lithium Hexafluorophosphate	Electrolyte	LFP, NMC, NCA	Mainstream Price (RMB/ton)	

5.3 Product Characteristics

Data on the energy density of a given chemistry type by a given manufacturer in a given year is scraped from company filings, company press releases, market intelligence reports, and battery seller websites. Data on the number of new patents filed in a given year is sourced from the World Intellectual Property Organization (WIPO) PATENTSCOPE database. Patents including the phrase "electric vehicle" and either the word "battery," "cathode," "anode," or "electrolyte" in their main text were screened for. Every simple patent family is counted as a single patent.

Table 2 reports the descriptive statistics of key non-marginal cost and non-price variables that are input into the price prediction model, and, subsequently, the nested logit demand model.

Table 2: Descriptive Statistics of Model Inputs

	Mean	**Std**	**Min**	**Median**	**Max**
Sales Volume (MWh)	20.09	33.58	0.05	4.53	144.56
Market Share (%)	10.46	9.99	0.54	7.59	40.03
Energy Density (Wh/kg)	167.23	37.04	103.80	162.40	255.00
Number of Patents Filed	219.83	203.71	2.00	188.00	1028.00

6 Estimation

To carry out the regressions specified at the end of Section 4.2, the linear instrumental variables (IV) generalized method of moments (GMM) algorithm built into the PyBLP Python package is used. The goal of the estimation is to minimize the difference between the observed market shares, i.e., those that are input into the model, and the predicted market shares, which are functions of the parameters estimated by the model.

A population moment condition, which corresponds to product j in market t, is defined as:

$$E[\omega_{jt}(\theta^*) \cdot Z_{jt}] = 0.$$

Here, $\omega_{jt}(\theta^*)$ is a scalar and represents the actual "market-specific valuation of the unobserved product characteristics," being a function of the population parameters θ^*. Z_{jt} is a vector that contains the values of the instrumental variable in question for product j in market t.

To construct an estimator, $\omega_{jt}(\theta)$, which is a function of the estimated parameters θ, must be defined. This was defined earlier to be ξ_{jt}. With this, the GMM estimator can be defined as:

$$\hat{\theta} = \arg\min_{\theta}[\hat{g}(\theta)\text{'}W\hat{g}(\theta)].$$

Here, $\hat{g}(\theta)$ is an M x 1 vector of sample moment conditions, and W is an internally computed M x M weighting matrix. Thus, the output of the function that is minimized, $\hat{g}(\theta)\text{'}W\hat{g}(\theta)$, is another M x 1 vector that represents the weighted sum of squared moment conditions. Given that each population moment condition requires convergence to zero, the vector of parameters, θ, which minimize the aforementioned function, also minimize the difference between the population weighted sum and the sample weighted sum, i.e., are the optimal parameters.

The basic steps in this iterative optimization process are as follows. It starts with some initial guess for the estimated set of parameters, θ. For this initial guess, the mean utility of product j in market t, V_{jt}, is calculated for all values of j and t by finding that which forces the estimated market share to equal the observed share. This optimization step ceases when the change in the following formulation is smaller than an already very low certain threshold:

$$V_{jt} + \log(s_{jt}) - \log(s_{jt}(V_{jt}, \theta))$$

Next, using these shares and following the linear expression specified at the end of Section 4.2, each value of $\omega_{jt}(\theta)$ is calculated as:

$$\omega_{jt}(\theta) = [\log(s_{jt}) - \rho \times \log(\bar{s}_{(j|h)t}) - \log(s_{0t})] - [\alpha p_{jt} + \beta x_{jt}].$$

Then, these values of $\omega_{jt}(\theta)$ are used to construct sample moment conditions, which are subsequently used to calculate the GMM estimator. Based on this output value, the program uses the built-in Trust-Region Constrained Algorithm to identify a better guess for θ, and the estimation is rerun. This is repeated until the GMM estimator no longer changes to a significant extent between iterations.

The output is a vector $\hat{\theta}$ that contains the best estimates for price sensitivity α, the vector of coefficients on observed product characteristics β, a nesting parameter ρ, and all corresponding standard errors.

7 Results

7.1 Estimation of Marginal Costs

Step (1) outlined in Section 4.1 is carried out here to determine the raw materials component of marginal cost of each battery chemistry in each region (China versus rest-of-world) in each year, as a function of raw material prices. As shown in Figures 1 and 2, there is significant variation in the marginal costs of batteries of different chemistries. There are also minor differences in the marginal costs of batteries of the same chemistry type but produced in China versus other countries, which are purely due to differences in lithium carbonate (for LFP batteries) and lithium hydroxide (for NMC/NCA batteries) prices across both regions. Notably, an equivalent figure to Figure 2 could have been generated for rest-of-world marginal costs, but it would have appeared essentially identical, based on the similarities between per-chemistry Chinese and rest-of-world marginal costs shown in Figure 1.

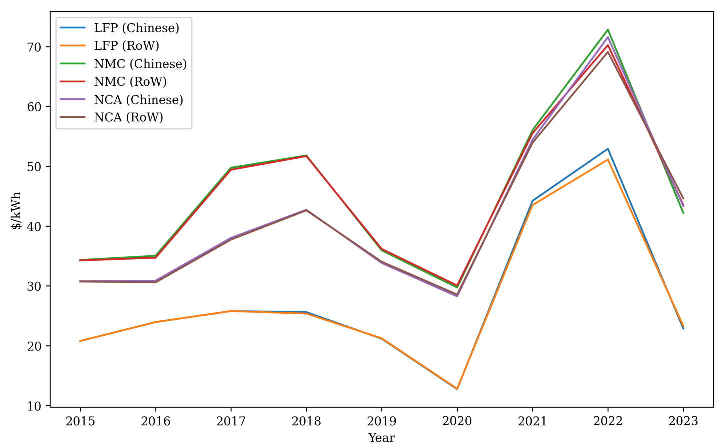

Figure 1: Marginal Cost Time Series by Battery Chemistry and Region.

Calculated values for the raw materials component of marginal cost are shown to be highly correlated with raw material input prices, as expected, based on how they are determined. As shown in Figures 3 and 4, marginal costs of LFP batteries are strongly correlated with lithium carbonate prices and moderately correlated with iron phosphate prices. Notably, although lithium carbonate only makes up 7.4% of the LFP cathode by weight, lithium carbonate prices heavily control LFP marginal costs, largely on the basis of being so expensive. As shown in Figures 5 and 6, NMC marginal costs are strongly correlated with cobalt, lithium hydroxide, and nickel prices, and weakly correlated with manganese sulfate prices. Notably, as companies transitioned from NMC(111) to NMC(622) to NMC(811) over time, cobalt (and manganese) made up a decreasing % of cathode weight, and nickel made up an increasing % of cathode weight. Finally, as shown in Figures 7 and 8, NCA marginal costs are strongly correlated with nickel, lithium hydroxide and aluminum prices, and moderately correlated

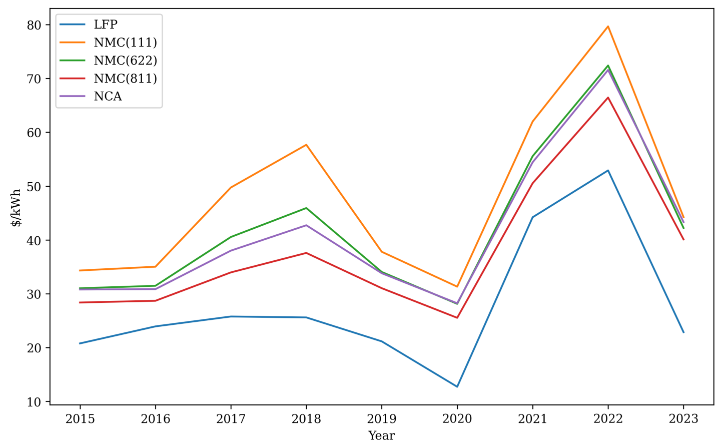

Figure 2: Marginal Cost Time Series by Battery Chemistry (China Only).

with cobalt prices. Using these raw material marginal cost values, Step (2) outlined in Section 4.1 is carried out at this point in order to determine battery pack marginal costs, based on the cell and pack manufacturing markup values reported in Section 5.2.

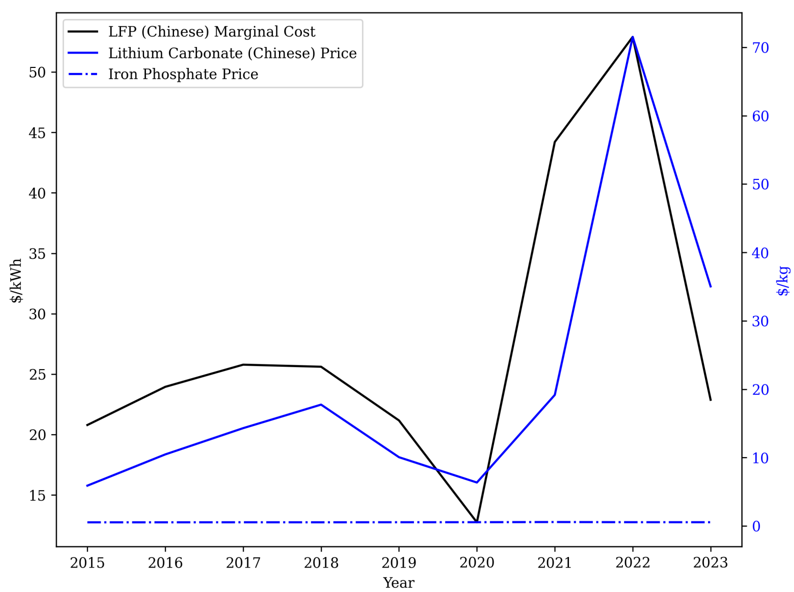

Figure 3: Marginal Cost of Chinese LFP Battery and Relevant Raw Material Input Prices.

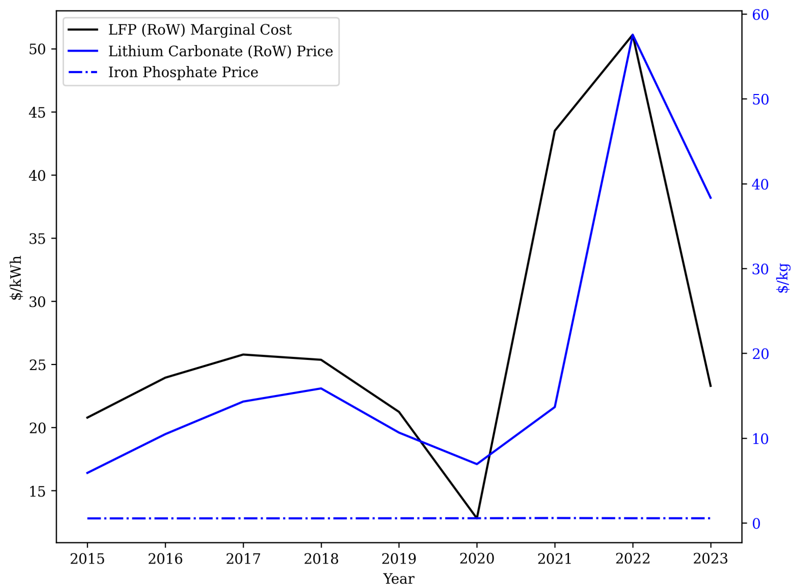

Figure 4: Marginal Cost of Rest-of-World LFP Battery and Relevant Raw Material Input Prices.

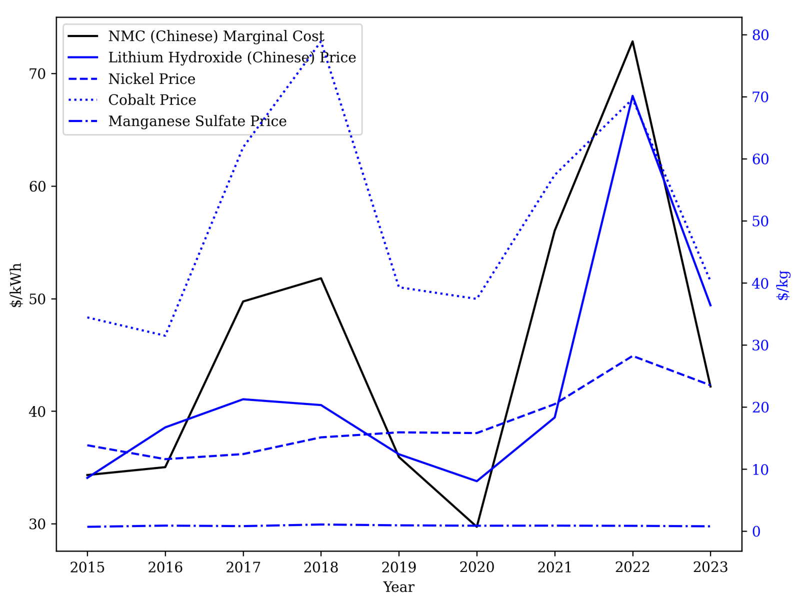

Figure 5: Marginal Cost of Chinese NMC Battery and Relevant Raw Material Input Prices.

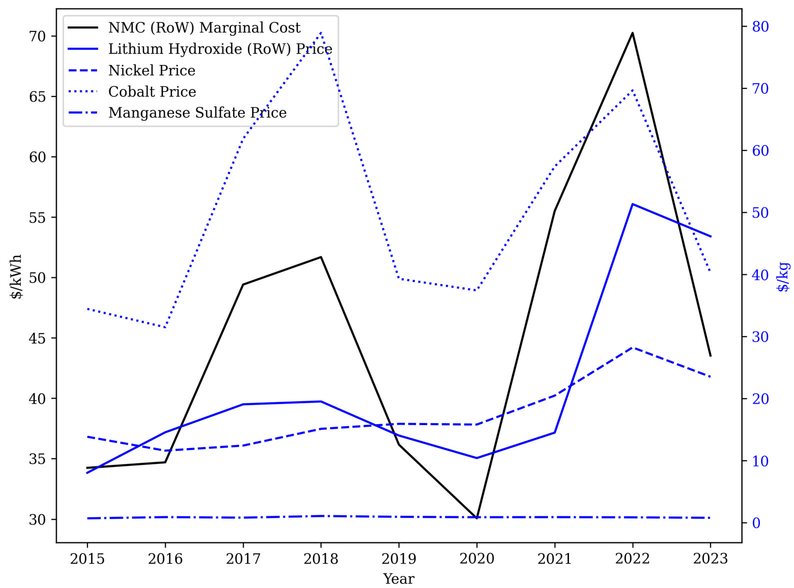

Figure 6: Marginal Cost of Rest-of-World NMC Battery and Relevant Raw Material Input Prices.

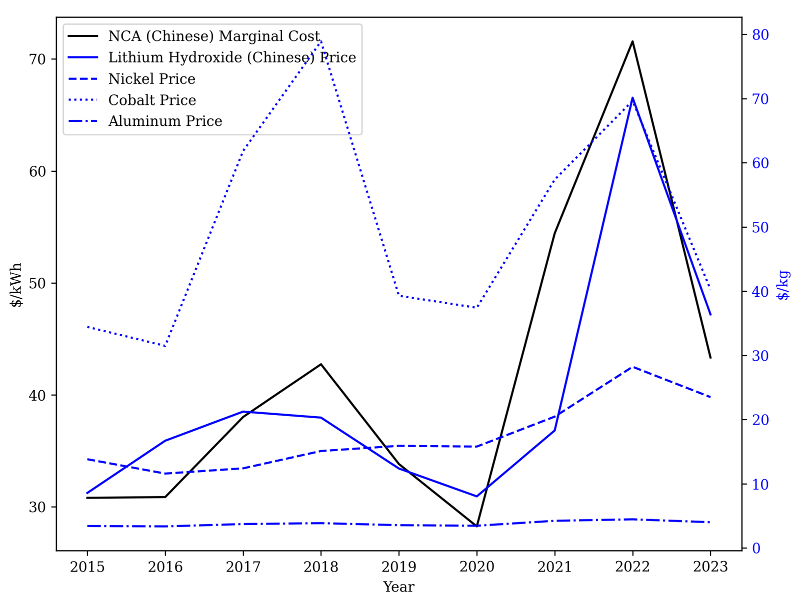

Figure 7: Marginal Cost of Chinese NCA Battery and Relevant Raw Material Input Prices.

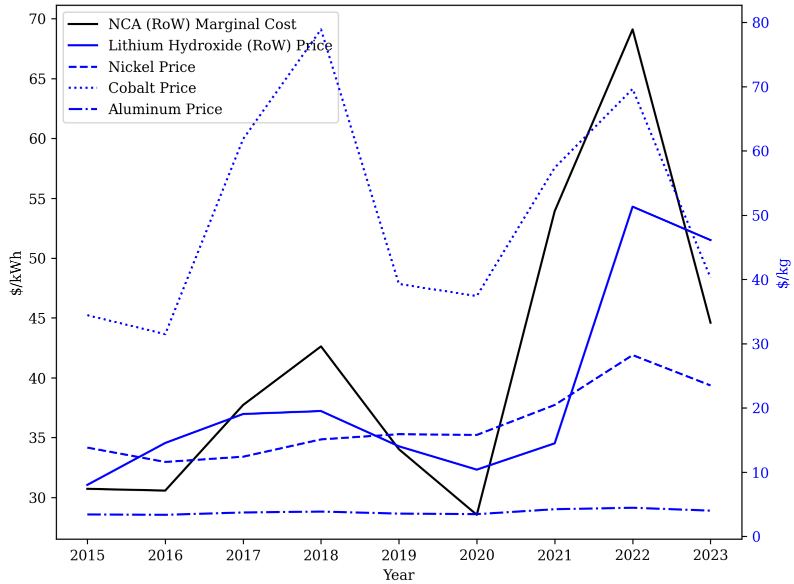

Figure 8: Marginal Cost of Rest-of-World NCA Battery and Relevant Raw Material Input Prices.

7.2 Estimations of Price Markups

In accordance with Step (3) outlined in Section 4.1, demand is regressed on marginal cost and an interaction term (constructed as marginal cost, multiplied by the China dummy variable) for an average Chinese firm and an average rest-of-world firm. The results are presented in Table 3.

It is found that, for the average rest-of-world firm, the sensitivity of demand to marginal cost is –0.671, a result that is statistically significant at the 5% confidence level. It is also found that the difference in sensitivity of demand to marginal cost for the average rest-of-world firm versus the Chinese firm is +0.410, which is statistically significant at the 1% confidence level. This means that the sensitivity of demand to marginal cost for the average Chinese firm is equal to (–0.671 + 0.410) = –0.261. Time and chemistry fixed effects are relatively large, and three out of the four are statistically significant at the 1% or 5% significance level. Put in simpler terms, a $1 increase in marginal cost leads to a 0.261 GWh decrease in sales volume for the average Chinese firm, versus a 0.671 GWh decrease for the average rest-of-world firm.

This result can indeed be primarily attributed to differences in markups of price over raw material marginal costs. It cannot be attributed to the fact that Chinese firms have entered the LFP market more, and LFP batteries are cheaper, because chemistry fixed effects are included. It cannot be attributed to differences in product characteristics, as the comparison is being made between a single average Chinese firm and a single average rest-of-world firm, the construction of which erases firm-level heterogeneity of characteristics, and product characteristics are assumed to be uncorrelated with location. One may think

Table 3: Regressing Demand on Marginal Cost for Average Chinese Firm vs Average Rest-of-World Firm

Variable	**Estimate**
Constant	7.281 (7.857)
Marginal Cost	-0.671** (0.283)
Marginal Cost x China	0.410*** (0.131)
2018-2020	7.047 (5.733)
2021-2023	44.690*** (7.560)
NMC Chemistry	14.623** (7.117)
NCA Chemistry	25.619*** (7.889)
R-squared	0.6065
Observations	45

Note: *** p<0.01, ** p<0.05, * p<0.1

it could include unobserved, non-price-related preferences regarding purchasing from China (e.g., the marginal cost elasticity of the average Chinese firm should actually be lower, but EV OEMs are wary of purchasing from China); however, given that policies and attitudes regarding purchasing from China have changed over time, these should be captured by the included time fixed effects.

Hence, the markup of price over marginal cost for the average Chinese firm can be thought of as being $\frac{0.261}{0.671} = 38.93\%$ of that for the average rest-of-world firm, i.e., $\frac{\beta_1+\beta_2}{\beta_1} =$

0.3893. Thus, given that the Chinese firms in the sample produced 51.84% of all batteries, the average markup of price over marginal cost, as a percentage of marginal cost, can be calculated using the following formulation (revisit Section 4.1 for further details):

$$x_{RoW} = \frac{0.3}{0.5184 \times 0.3893 + (1 - 0.5184)} = 0.4390$$

and that of the average Chinese firm can be calculated as:

$$x_{\text{Chinese}} = 0.3893 \times 0.4390 = 0.1709$$

Thus, Chinese firms are assumed to markup their prices over their marginal costs by 17.09%, while rest-of-world firms are assumed to do so by 43.90%. These markup values are used to scale prices for the subsequent analysis, as detailed in Section 4.1.

Descriptive statistics of estimates of the raw material component of marginal cost and estimated scaled prices are reported in Table 4. Notably, for the nested logit model, input data on market shares ranges from 0 to 1 (0% to 100%), and the China dummy variable can take on the value of 0 or 1. Therefore, all other variables, reported in both Table 2 and Table 4, are scaled down to normalize their magnitudes. Specifically, all price and energy density values are scaled down by a factor of 100, and the number of patents filed are scaled down by a factor of 1,000. The estimated coefficients are interpreted accordingly.

Table 4: Descriptive Statistics of Marginal Cost and Price Inputs

	Mean	**Std**	**Min**	**Median**	**Max**
Raw Materials Component of Marginal Cost (\$/kWh)	38.69	14.95	12.73	35.03	72.85
Prices, Scaled (\$/kWh)	112.90	27.78	66.35	113.15	184.33

7.3 Nested Logit Demand With Scaled Prices

As described in Section 4.2, $\log(s_{jt}) - \rho \times \log(\bar{s}_{(j|h)t}) - \log(s_{0t})$ is regressed on estimated prices and various observed product characteristics. The results are presented in Table 5. To produce the results in column (1), only energy density is included in the regression. To produce the results in column (2), two additional manufacturer-level characteristics are included—the number of patents filed by the firm, and whether the firm is Chinese.

In both regressions, the price and energy density coefficients are statistically significant and are of the expected signs—a higher price decreases the utility, and hence likelihood, of

Table 5: Nested Logit Demand Results With Scaled Prices

Variable	(1) **Estimate**	(1) **Odds Ratio**	(2) **Estimate**	(2) **Odds Ratio**
Constant	-2.002*	0.135	-1.833	0.160
	(1.068)		(1.128)	
Price	-1.401**	0.246	-1.228**	0.293
	(0.568)		(0.600)	
Energy Density	1.852***	6.371	1.611**	5.006
	(0.650)		(0.775)	
Is Chinese			0.139	1.150
			(0.277)	
Number of Patents Filed			0.109	1.115
			(0.164)	
Rho Estimate	0.371		0.461	
	(0.303)		(0.341)	
Hansen J-Stat	4.273		4.586	
Hansen J-Stat p-value	0.511		0.469	

Note: *** p<0.01, ** p<0.05, * p<0.1

choosing that product, and a higher energy density increases the utility, and hence likelihood, of choosing it. Both estimates of ρ, i.e. the within-nest correlation parameter, are positive and relatively large, indicating that consumers indeed substitute more readily across products produced by the same manufacturer. Finally, the p-value of each Hansen J-statistic is highly non-significant, indicating that the instruments are valid and fulfill the exogeneity condition and, hence, the model is well-specified.

According to the odds ratios reported for (1), a \$100 increase in price reduces the likelihood of selection by (1 – 0.246) x 100% = 75.4%, while a 100 Wh/kg increase in energy density increases this likelihood by (6.371 – 1) x 100% = 537.1%. The effects are slightly smaller, but still significant and substantial, for (2): a \$100 increase in price reduces the likelihood of selection by (1 – 0.293) x 100% = 70.7%, while a 100 Wh/kg increase in energy density increases it by (5.006 – 1) x 100% = 400.6%. Based purely on the magnitude of their coefficients, a unit increase in energy density has a greater effect on utility than a unit decrease in prices do. This can be interpreted to mean that EV OEMs highly value batteries with higher energy density, but there is a competing price effect that causes them to readily substitute for cheaper batteries. This likely occurs both within nests (e.g. an EV OEM's choice between NMC versus LFP batteries, both manufactured by CATL) and across nests. The latter likely occurs on two levels: within a chemistry (e.g., an EV OEM's choice between two LFP battery products, one of which may be marginally cheaper but marginally less energy dense than the other) and across chemistries (e.g. an EV OEM's choice between an LFP battery and an NMC battery, the former of which is significantly cheaper but significantly less energy dense than the latter).

Given that the price effect of being based in China has already been calculated and incorporated (see Section 7.2), the coefficient on the China dummy of 0.139 can be interpreted as the remaining non-price effect. The coefficient is positive, small, and statistically insignificant, which can be interpreted in two different ways. It may mean that EV OEMs generally don't have strong preferences about whether the manufacturing firm is based in China or not, beyond how it can lead to lower costs, which are passed on as lower prices. Alternatively, it may mean that there is substantial heterogeneity across EV OEMs regarding whether manufacturing in China is favorable (due to lower costs, government incentives, easier access to the Chinese EV market, etc.) or unfavorable (due to geopolitical and supply chain risks), and the coefficient thus converges to zero. The coefficient on the number of patents is small and positive, but it is also not statistically significant. We may interpret it to mean that a higher number of patents leads to only a minimal increase in the likelihood of selecting a product manufactured by that firm. At first, this coefficient may be expected to be larger, but there are plausible reasons for why it is nearly zero—for example, the firm may be overinvesting in R&D instead of focusing on optimizing its production and supply chain strategies.

Between (1) and (2), the incorporation of the two additional observable manufacturer-level characteristics dampens the estimated effects of prices and energy density on demand, i.e., reduces the absolute values of their coefficients. Given that both coefficients remain statistically significant, this likely indicates that explanatory power for the variation in demand is now being allocated more accurately to different characteristics, and the model is generally better able to explain consumer choices. Secondly, the rho estimate increases from 0.371 to 0.461; this indicates that when observable manufacturer-level characteristics are accounted for, the correlation between products in a nest, on the basis of unobserved characteristics, is better captured when more observed characteristics are accounted for. Therefore, the model that produced the results in (2) will be used as the main specification for the remainder of the analysis.

Using this model, own- and cross-price elasticities for all 11 products included in the sample are calculated and reported in Table 6. Each element is the mean of the elasticities calculated for each of the nine markets (years). Each matrix element, $(j,\, k)$, represents the % change in market share of product j in response to a 1% increase in the price of product k; in other words, it represents the price elasticity of product j with respect to product k. As expected, all own-price elasticities (which are represented by the diagonal elements) are found to be negative and relatively large. This implies that an increase in the price of a product results in a relatively large decrease in demand for that product. All cross-price elasticities (which are represented by the off-diagonal elements) are positive and relatively small. This implies that an increase in the price of one product results in a slight increase in demand for every other product, meaning that they are substitutes.

Given that price elasticity is defined as the percentage change in demand of one product due to a 1% increase in price of that product (in the case of own-price elasticity) or a different product (in the case of cross-price elasticity), and different products sold by different manufacturers in different years have different prices, it is futile to compare the absolute values of elasticities across rows. Rather, we can focus on patterns down a single column. Generally, cross-price elasticities for products produced by the same manufacturer are larger than those for products produced by different manufacturers—for example, the price elasticity of BYD's NMC battery (row) with respect to BYD's LFP battery (column) is 0.9106, while the price elasticities of all other products with respect to BYD's LFP battery ranges from 0.1188 to 0.2006. This confirms that EV OEMs substitute more readily between the same manufacturer's products (i.e., within the nest from which they already purchase) and, hence, that the nesting structure is correctly capturing such substitution patterns.

It is worth noting that the cross-price elasticities of LGES's NCA battery appear extreme, relative to other values that are in the same columns. This is likely due to the fact that it was introduced to the market in 2022, meaning that each mean price elasticity value is calculated based on only two markets' worth of data.

Beyond within-nest substitution (and the abnormalities of LGES's NCA battery), there is some additional, albeit much less substantial, heterogeneity. For example, the price elasticity of BYD's NMC battery with respect to LGES's NMC battery is relatively large (0.2628, versus other values in the column), and the price elasticities of BYD's NMC battery and LGES's LFP battery with respect to Panasonic's NMC battery are relatively large (0.0458, versus other values in the column). This is an indication that a range of factors (specifically, a combination of observed and unobserved characteristics) are indeed at play and causing differentiation between battery products, making some better substitutes than others.

Table 6: Mean Own- and Cross-Price Elasticities

	BYD LFP	BYD NMC	CATL LFP	CATL NMC	LGES LFP	LGES NCA	LGES NMC	Panasonic NCA	Panasonic NMC	SK NMC	Samsung NMC
BYD LFP	-1.0954	0.2204	0.0519	0.2610	0.0040	0.0934	0.2507	0.3571	0.0423	0.0610	0.1036
BYD NMC	0.9106	-2.3776	0.0405	0.2580	0.0040	0.0810	0.2628	0.3855	0.0458	0.0560	0.1054
CATL LFP	0.1258	0.0284	-1.8110	1.2536	0.0040	0.0934	0.2507	0.3571	0.0423	0.0610	0.1036
CATL NMC	0.1258	0.0284	0.2090	-1.3304	0.0040	0.0934	0.2507	0.3571	0.0423	0.0610	0.1036
LGES LFP	0.1188	0.0284	0.0405	0.2580	-2.4694	0.4679	1.6516	0.3855	0.0458	0.0560	0.1054
LGES NCA	0.2006	0.0177	0.1493	0.3636	0.0192	-3.1101	1.4088	0.1613	0.0175	0.1227	0.1017
LGES NMC	0.1258	0.0284	0.0519	0.2610	0.0250	0.5800	-1.5862	0.3571	0.0423	0.0610	0.1036
Panasonic NCA	0.1258	0.0284	0.0519	0.2610	0.0040	0.0934	0.2507	-1.4252	0.1881	0.0610	0.1036
Panasonic NMC	0.1258	0.0284	0.0519	0.2610	0.0040	0.0934	0.2507	1.6196	-2.9782	0.0610	0.1036
SK NMC	0.1258	0.0284	0.0519	0.2610	0.0040	0.0934	0.2507	0.3571	0.0423	-1.6466	0.1036
Samsung NMC	0.1258	0.0284	0.0519	0.2610	0.0040	0.0934	0.2507	0.3571	0.0423	0.0610	-1.6041

Table 7: Mean Own- and Cross-Energy Density Elasticities

	BYD LFP	BYD NMC	CATL LFP	CATL NMC	LGES LFP	LGES NCA	LGES NMC	Panasonic NCA	Panasonic NMC	SK NMC	Samsung NMC
BYD LFP	2.1774	-0.4114	-0.1081	-0.6255	-0.0059	-0.1593	-0.4579	-0.7260	-0.0592	-0.1047	-0.1539
BYD NMC	-1.7441	4.4381	-0.0757	-0.5938	-0.0059	-0.1169	-0.4745	-0.7852	-0.0629	-0.0925	-0.1523
CATL LFP	-0.2475	-0.0516	3.7710	-2.9696	-0.0059	-0.1593	-0.4579	-0.7260	-0.0592	-0.1047	-0.1539
CATL NMC	-0.2475	-0.0516	-0.4341	3.1917	-0.0059	-0.1593	-0.4579	-0.7260	-0.0592	-0.1047	-0.1539
LGES LFP	-0.2234	-0.0516	-0.0757	-0.5938	3.5533	-0.6753	-2.9658	-0.7852	-0.0629	-0.0925	-0.1523
LGES NCA	-0.3911	-0.0312	-0.3137	-0.9303	-0.0228	5.0630	-2.3865	-0.2668	-0.0301	-0.2081	-0.1645
LGES NMC	-0.2475	-0.0516	-0.1081	-0.6255	-0.0366	-0.9968	2.8860	-0.7260	-0.0592	-0.1047	-0.1539
Panasonic NCA	-0.2475	-0.0516	-0.1081	-0.6255	-0.0059	-0.1593	-0.4579	2.7505	-0.2797	-0.1047	-0.1539
Panasonic NMC	-0.2475	-0.0516	-0.1081	-0.6255	-0.0059	-0.1593	-0.4579	-3.1883	4.5068	-0.1047	-0.1539
SK NMC	-0.2475	-0.0516	-0.1081	-0.6255	-0.0059	-0.1593	-0.4579	-0.7260	-0.0592	2.5874	-0.1539
Samsung NMC	-0.2475	-0.0516	-0.1081	-0.6255	-0.0059	-0.1593	-0.4579	-0.7260	-0.0592	-0.1047	2.4510

Own- and cross-energy density elasticities of demand are similarly calculated and reported in Table 7. Own-energy density elasticities are positive, meaning that an increase in energy density of a product leads to an increase in demand for it. Cross-energy density elasticities are negative and are generally substantially larger for products manufactured by the same firm, due to stronger substitution both to and from products in the same year-manufacturer nests.

7.4 Robustness Checks

As discussed in Section 7.3, the regression results of market shares on price and energy density did not differ significantly from those of market shares on price, energy density, number of patents filed, and whether the firm is based in China—the coefficients remained of similar magnitudes, of the same signs, and statistically significant (for the ones that were originally). This served as a first robustness check, as it confirmed the strength of the explanatory power of both price and energy density, even with the addition of extra explanatory variables.

Two more robustness checks are run. Firstly, $\log(s_{jt}) - \rho \times \log(\bar{s}_{(j|h)t}) - \log(s_{0t})$ is regressed on unscaled prices, which are calculated as a linear function of marginal costs, as described in Section 7.1, without incorporation of the price markups estimated in Section 7.2 (instead, a uniform 30% markup is applied to all batteries), and various observed product characteristics. This is done in order to confirm that it is not the different scaling of prices for Chinese versus rest-of-world firms that gives price and energy density their explanatory power. The results are presented in Table 8.

The coefficients on price and energy density reported in Table 7 have the same sign as and do not differ significantly in magnitude nor statistical significance from those reported for the main model in column (2) of Table 5. The price coefficient in the main model is -1.228, versus -1.222 in the robustness check; the energy density coefficient in the main model is 1.611, versus 1.544 in in the robustness check. From this, we can conclude that the scaling of prices is not responsible for the explanatory power of price and energy density. Also, the p-value of each Hansen J-statistic is highly non-significant, indicating that the model is still well-specified and the results can be interpreted reliably.

Notably, when unscaled prices are used, the coefficient on the China dummy is relatively large and positive (0.422) and statistically significant. In contrast, recall that in the main model, the coefficient is smaller and statistically insignificant (0.139). By incorporating differences in price markups for Chinese versus rest-of-world firms into price calculations, as done in Section 7.2, we essentially break out the effect of being based in China into a price effect and a non-price effect. This difference confirms that the price effect is significantly larger than the non-price effect.

As a second robustness check, the main model is rerun without nesting, changing it from a nested logit model to a standard logit model. This is done to confirm that enforced substitution patterns between products in the same nests is not falsely attributing explanatory power to price and the observed product characteristics. The model is carried out according

Table 8: Robustness Check #1: Nested Logit Demand Results With Non-Scaled Prices

Variable	**Estimate**	**Odds Ratio**
Constant	-1.856* (1.118)	0.156
Price	-1.222** (0.578)	0.295
Energy Density	1.544** (0.747)	4.684
Is Chinese	0.422** (0.193)	1.526
Number of Patents Filed	0.114 (0.159)	1.120
Rho Estimate	0.493 (0.326)	
Hansen J-Stat	4.809	
Hansen J-Stat p-value	0.186	

Note: *** p<0.01, ** p<0.05, * p<0.1

to the following regression specification:

$$\log(s_{jt}) - \log(s_{0t}) = \alpha p_{jt} + \beta_1(\text{energy density})_{jt} + \beta_2(\text{firm's number of patents})_{f(j)t} + \beta_3 \mathrm{D}_{f(j)} + \xi_{jt}.$$

and the results are included in Table 9.

Again, the coefficients on price and energy density have the same sign as and do not differ significantly in magnitude nor statistical significance from those reported for the main model in column (2) of Table 5. The price coefficient in the main model is -1.228, versus -1.792 in the robustness check; the energy density coefficient in the main model is 1.611, versus 2.626 in in the robustness check. Given that both coefficients are smaller in absolute magnitude in the main model, it can be concluded that nesting does take on some explanatory power, but the primary drivers are still price and energy density, and that the main model is robust.

Table 9: Robustness Check #2: Standard Logit Demand Results

Variable	**Estimate**	**Odds Ratio**
Constant	-3.127*** (0.934)	0.044
Price	-1.792** (0.782)	0.167
Energy Density	2.626*** (0.437)	13.820
Is Chinese	-0.042 (0.374)	0.959
Number of Patents Filed	-0.083 (0.187)	0.920
Hansen J-Stat	3.756	
Hansen J-Stat p-value	0.585	

Note: *** p<0.01, ** p<0.05, * p<0.1

The coefficients on the China dummy and the number of new patents filed do change sign but are still similarly small in magnitude. However, it's important to note that they are statistically insignificant in both models, meaning that the difference in their values should not be interpreted as having any implication for robustness.

Collectively, both robustness checks confirm that the model is robust.

7.5 Counterfactuals

Estimated values for utilities and market shares, as well as the prices originally input into the model, can be used to calculate consumer surpluses, in accordance with the equation for consumer surplus included at the end of Section 4.2. Such calculations are valuable because they enable us to carry out counterfactual analysis—by quantifying and comparing what economic welfare would have been under alternative circumstances, we can understand what events and trends have had the largest effects on welfare and generally been most important to consumers in the market.

Two counterfactuals are considered. The first is concerned with the outcome if LFP technology had never been introduced to the market. As mentioned in Section 1, LFP batteries offer a cheaper, albeit less energy dense, alternative to NMC and NCA batteries. In practice, the counterfactual is carried out by excluding LFP products (from CATL, BYD, and LG Energy) from the analysis, recalculating market shares based on estimated aggregate parameters from the base model and updated product-level utilities, and calculating consumer surplus in each market. The second counterfactual considers the outcome if NMC technology had never progressed from NMC(111) to NMC(622) to NMC(811). NMC(622) batteries and, subsequently, NMC(811) batteries offered less cobalt-intense, and hence cheaper, alternatives to NMC(111) batteries. Importantly, it is assumed that if NMC(622) and NMC(811) were never brought to market, the upward trajectory of the energy density of NMC batteries would not have been any different—it is unclear what the exact effect of NMC battery chemistry evolution was on NMC battery energy density generally, given that energy density still improved readily for many companies in the years before NMC(622) was introduced, and improvements in pack-level energy density can be attributed to a wide range of processes and techniques, including cell-to-pack technology and electrolyte and anode composition. As a result, this counterfactual is practically implemented by making general NMC prices equivalent to those for NMC(111) for all years (rather than a 50-50 split of NMC(111) and NMC(622) starting in 2018, and a 33-33-33 split of NMC(111), NMC(622), and NMC(811) starting in 2021).

For the base model and each counterfactual, the aggregate consumer surplus (i.e., summed across all markets) and average consumer surplus (i.e., market average) are calculated. The changes in both surplus measures, relative to those of the base model, are also calculated. Results are reported in full in Table 10. Both counterfactuals lead to lower aggregate and average consumer surplus than the base model. However, counterfactual (1) leads to a significantly larger decrease in aggregate consumer surplus than counterfactual (2) (-0.936 versus -0.258) and in average consumer surplus (-0.104 versus -0.029). This is an intuitive result. Completely removing LFP batteries from the market has three effects on the products that remain in the market, i.e. NMC and NCA batteries: (a) it drives down aggregate utility, U_j, due to the fact that all products now have high prices, (b) it drives up aggregate utility, U_j, due to the fact that all products now have high energy density, and (c) it drives up prices, p_j. Effects (a) and (c) likely substantially outweigh effect (b), as the difference in prices between LFP batteries and NMC/NCA batteries is relatively larger than the difference in energy density. Considering the other counterfactual, halting technological progress of NMC batteries has effects similar to the aforementioned (a) and (c); however, the difference in price of an LFP battery versus an NMC(111) battery is significantly greater than the difference in price of an NMC(111) battery versus an NMC(811) battery, which explains why both effects (a) and (c) are dampened.

Table 11 reports the difference in consumer surplus between each counterfactual and the base model in each year. For LFPs, the difference increased for the first time during the

Table 10: Consumer Surplus Under Counterfactuals

	Base Model	(1) No LFP	(2) No Development of NMC
Aggregate Consumer Surplus	15.143	14.207	14.885
Change in Aggregate Consumer Surplus		-0.936	-0.258
Average Consumer Surplus	1.683	1.579	1.654
Change in Average Consumer Surplus		-0.104	-0.029

years 2016-2018. This may be due to the fact that, as can be gleaned from Figures 3-8, there were large, postivie price shocks to raw material input prices; as a result, it likely would have been particularly favorable for EV OEMs to substitute to a less lithium-intense, cobalt-free, and generally cheaper battery, such as an LFP battery, during this time. The difference also started rapidly increasing starting in 2021, growing from -0.079 in that year to -0.254 in 2023. This can likely be explained by the fact that LFPs were increasingly adopted beginning in 2020, growing from 12.7% of total batteries sold in that year to 36.5% in 2023; given that calculation of consumer surplus is weighted by market share, it makes sense that the negative effect of the removal of LFPs only became more pronounced as its market share increased in the base model. For NMCs, there is no difference in the years 2015-2017, as NMC(622) batteries were only introduced starting in 2018 (and NMC(811) batteries in 2021). The difference spiked in 2021 and 2022. This can likely be attributed to the massive price shocks to lithium hydroxide, lithium carbonate, nickel, and cobalt that occurred in the wake of the energy crisis and are also visible in Figures 3-8; such price shocks would have induced consumers to switch not only to cheaper LFPs, as previously discussed, but also to cheaper NMCs more rapidly.

The results of this counterfactual analysis can be generalized. From counterfactual (1), we may learn that being able to introduce products to the market that offer a significantly cheaper option, without compromising too much in terms of other characteristics that are important to consumer utility (such as energy density), plays an important role in boosting consumer welfare. From counterfactual (2), we may learn that being able to respond to material price shocks with actual R&D to bring a cheaper version of a current product to the market, rather than just, for example, rerouting supply chains for incumbent products, can also boost consumer welfare.

Table 11: Year-by-Year Changes in Consumer Surplus Under Counterfactuals

	(1) **No LFP**	**(2)** **No Development of NMC**
2015	-0.067	0.000
2016	-0.115	0.000
2017	-0.073	0.000
2018	-0.078	-0.041
2019	-0.065	-0.015
2020	-0.045	-0.016
2021	-0.079	-0.084
2022	-0.158	-0.080
2023	-0.254	-0.022

8 Conclusion

This study aims to model demand in the EV lithium-ion battery market. The first part of this study determines prices of EV lithium-ion battery products as a function of raw material prices and estimated markups of price over raw material costs for Chinese versus rest-of-world manufacturers. The second part of this study applies the nested logit demand model to the EV battery market to measure the effects of these determined prices, as well as three product characteristics—energy density, the number of new patents filed by the manufacturer, and whether the manufacturer is based in China—on consumer utility and preferences. Price is found to have a significant negative effect, and energy density to have a significant positive effect. The number of new patents filed by the manufacturer and whether the manufacturer is based in China have small positive effects, but neither are statistically significant. The positive rho estimate, combined with negative own-price elasticities and positive cross-price elasticities, confirm that nesting by manufacturer accurately captures substitution patterns.

Furthermore, two counterfactual scenarios are analyzed: (1) LFP batteries were never introduced to the market, and (2) NMC battery technology never progressed from NMC(111) to NMC(622) to NMC(811). It is found that aggregate and average consumer surplus are lower under both scenarios, compared to the base model, and that counterfactual (1) leads to lower consumer surplus than counterfactual (2), meaning that the introduction of LFP batteries boosted consumer surplus more than the introduction of higher nickel content NMC batteries. It is also determined that, across both counterfactuals, consumer surplus decreased more in years when the battery product in question was in higher demand and when relevant raw material prices spiked the most.

There are a number of potential future directions for this research. Firstly, the price markup estimation model could be enhanced. Marginal battery costs could be calculated as more than a function of raw material prices by including factors such as manufacturer-specific procurement and processing costs and whether the manufacturer applies cell-to-pack or cell-to-chassis technologies; this assumes that data pertaning to both of these things may become more accessible in the future. Relative price markups could also be calculated for more specific groups of firms than "Chinese" versus "rest-of-world," such as Chinese versus Korean versus Japanese firms, or even for groups of firms within certain production capacity ranges (as a proxy for economies of scale).

Secondly, if data on EV OEM-level purchasing decisions can be located, then the assumption of preference homogeneity across consumers can be further relaxed, and a mixed nested logit model can be run, instead. Going from a standard to a nested logit model (as this study has done) allows for more flexible substitution patterns, but explicitly incorporating random coefficients allows for heterogeneity in how consumers value different product characteristics. This may be important, given that some EV OEMs specialize in certain vehicle segments and may therefore care disproportionately about a certain characteristic.

Finally, more counterfactual analyses can be run to analyze other factors that have driven, and may continue to drive, the EV battery market. For example, a counterfactual could be modeled out in order to analyze what kinds of subsidies or incentives would maximize market-wide welfare (e.g., mining subsidies versus production subsidies versus R&D grants). A different counterfactual might investigate which of the battery technologies that have the potential to replace lithium-ion batteries in the next decade, namely sodium-ion and solid-state batteries, will likely have the largest (positive) effect on consumer surplus. Yet another counterfactual could be used to study how firms' increasing adoption of battery recycling processes may affect welfare in the long-run.

References

[1] *Global electric vehicle battery demand from 2016 to 2023, by leading region or country* (2023). Statista. Retrieved from: `https://www.statista.com/statistics/1103229/global-battery-demand-by-region/`.

[2] *The History of the Electric Car* (2023). U.S. Department of Energy. Retrieved from: `https://www.energy.gov/articles/history-electric-car`.

[3] Colato, J., Ice, L. *Charging into the future: the transition to electric vehicles* (2023). U.S. Bureau of Labor Statistics. Retrieved from: `https://www.bls.gov/opub/btn/volume-12/charging-into-the-future-the-transition-to-electric-vehicles.htm`.

[4] Ziegler, M. S., Song, J., Trancik, J. E. (2021). "Determinants of lithium-ion battery technology cost decline." *Energy Environ. Sci.* 14(12), 6074-6098.

[5] Chung, D., Elgqvist, E., Santhanagopalan, S. (2016). "Automotive Lithium-ion Cell Manufacturing: Regional Cost Structures and Supply Chain Considerations." *Clean Energy Manufacturing Analysis Center*.

[6] Orangi, S., Manjong, N., Clos, D. P., Usai, L., Burheim, O. S., Strømman, A. H. (2024). "Historical and prospective lithium-ion battery cost trajectories from a bottom-up production modeling perspective." *Journal of Energy Storage*, 76, 109800.

[7] Nykvist, B., Nilsson, M. (2015). "Rapidly falling costs of battery packs for electric vehicles." *Nature Climate Change*, 5, 329-332.

[8] Kittner, N., Lill, F., Kammen, D. M. (2017). "Energy storage deployment and innovation for the clean energy transition." *Nature Energy*, 2, 17125.

[9] Hsieh, I. L., Pan, M. S., Chiang, Y, Green, W. H. (2019). "Learning only buys you so much: Practical limits on battery price reduction." *Applied Energy*, 239, 218-224.

[10] Cagli, E. C. (2023). "The volatility spillover between battery metals and future mobility stocks: Evidence from the time-varying frequency connectedness approach." *Resources Policy*, 86(A), 104144.

[11] Mauler, L., Lou, X., Duffer, F., Leker, J. (2022). "Technological innovation vs. tightening raw material markets: falling battery costs put at risk." *Energy Advances*, 3.

[12] Zhu, X., Liu, K., Liu, J., Yan, A. (2022). "Is government RD subsidy good for BEV supply chain? The challenge from downstream competition." *Computers Industrial Engineering*, 165, 107951.

[13] Yang, Z. *How did China come to dominate the world of electric cars?* (2023). MIT Technology Review. Retrieved from: `https://www.technologyreview.com/2023/02/21/1068880/how-did-china-dominate-electric-cars-policy/`.

[14] Hu, Y., Yin, H., Zhao, L. (2025). "Subsidy Phase-Out and Consumer Demand Dynamics: Evidence from the Battery Electric Vehicle Market in China." *The Review of Economics and Statistics*, 107(2), 458-475.

[15] Fan, Z., Huang, S., Wang, X. (2021). "The vertical cooperation and pricing strategies of electric vehicle supply chain under brand competition." *Computers Industrial Engineering*, 152, 106968.

[16] Berry, S. (1994). "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics*, 25(2), 242-262.

[17] Berry, S. (1997). "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity." *Econometric Theory*, 13(2), 185-213.

[18] Hall, R. (2018). "New Evidence on the Markup of Prices Over Marginal Costs and the Role of Mega-Firms in the US Economy." *NBER*, Retrieved from: `https://www.nber.org/papers/w24574`.

[19] Stallard, J., Wheatcroft, L., Booth, S., Boston, R., Corr, S., De Volder, M., Inkson, B., Fleck, N. (2022). "Mechanical properties of cathode materials for lithium-ion batteries." *Joule*, 6(5), 984-1007.

[20] Ciez, R., Whitacre, J. F. (2016). "The cost of lithium is unlikely to upend the price of Li-ion storage systems." *Journal of Power Sources*, 320, 310-313.

[21] Pillai, U. (2015). "Drivers of cost reduction in solar photovoltaics." *Energy Economics*, 50, 286-293.

[22] Chen, W., Hsieh, I. L. (2023). "Techno-economic analysis of lithium-ion battery price reduction considering carbon footprint based on life cycle assessment." *Journal of Cleaner Production*, 425, 139045.

[23] Barwick, P., Kwon, H., Li, S., Zahur, N. (2025). "Drive Down the Cost: Learning by Doing and Government Policies in the Global EV Battery Industry." *NBER*, Retrieved from: `https://www.nber.org/papers/w33378`.