Critical Minerals, Geopolitics, and the Green Transition

Tomás Domínguez-Iino
University of Chicago

Jonathan T. Elliott
Johns Hopkins University

Allan Hsiao
Stanford University

May 7, 2026

(Latest version here)

Abstract

The green energy transition will be powered by the mining and processing of lithium, nickel, and cobalt, which are critical for the production of advanced batteries. These minerals are concentrated geographically but traded globally. We study the geopolitical implications of active policy intervention in key mining countries, and we discuss consequences for green technology adoption worldwide. We show that the joint use of critical minerals in advanced batteries generates complementarity, which shapes both producer welfare and battery adoption. We quantify supply chain vulnerability, policy spillovers across mineral markets, and the potential for mineral cartels.

We thank seminar participants at Chicago, Columbia, Cornell, Di Tella, Geneva, Penn State, PSE, Stanford, Toronto, UC Davis, UT Austin, Washington State, and Yale. Xilin Fan, James Han, Xinwen Huang, Sian McAllister, and Shihan Wang provided excellent research assistance. We are grateful for financial support from the Becker Friedman Institute.

1 Introduction

The energy transition calls for the large-scale adoption of advanced battery technologies. This green future depends crucially on the production of lithium, cobalt, and nickel, which are essential inputs for battery technologies. These critical minerals are traded in global markets, but their geological endowments are highly concentrated: the top three producing countries account for 70 to 85% of world output, far exceeding the concentration observed in fossil fuel markets. This paper studies how concentrated mineral resources will shape the green transition.

Geological endowments make concentration an intrinsic feature of this setting. Upstream mineral production is dominated by a small number of countries: Australia for lithium, the Democratic Republic of the Congo (DRC) for cobalt, and Indonesia for nickel. Extraction costs vary widely across mines due to persistent differences in ore grade, deposit types, and access to infrastructure. Downstream battery production also varies in its mineral usage. We focus on three main battery technologies: nickel manganese cobalt (NMC), nickel cobalt aluminum (NCA), and lithium iron phosphate (LFP). Each differs in its use of lithium, cobalt, and nickel as inputs.

This collection of battery technologies allows for both substitution and complementarity across minerals. Suppose Indonesia restricts nickel supply in order to raise nickel prices. A rise in nickel prices triggers two opposing forces for lithium demand. Through substitution, battery demand shifts toward LFP, which uses lithium but not nickel. Lithium demand rises. Through complementarity, battery demand shifts away from NMC and NCA, which use both lithium and nickel. Lithium demand falls. The relative strengths of these forces determine the net effect on lithium prices. This net effect matters for how strategic mineral policies distribute gains and losses across countries, which has implications for geopolitics and green adoption. Whether substitution or complementarity dominates is an empirical question.

To this end, we build a global dynamic equilibrium model that captures these upstream and downstream elements. On the demand side, we estimate an almost ideal demand system using electric vehicle (EV) manufacturer-level data on battery installations by technology. The data span the universe of EV models produced worldwide from 2018 to 2024. The demand system captures how EV manufacturers substitute

across battery technologies in response to changes in battery prices. Battery technologies combine multiple mineral inputs according to recipes dictated by chemical properties, and we observe these recipes with engineering data. The recipes govern the strength of complementarity across minerals, which interacts with substitution across batteries to determine green adoption.

On the supply side, we estimate a dynamic model of mining with a mine-level panel dataset that covers the near universe of global lithium, cobalt, and nickel production. For each mine, we observe production, capacity, ore grade, and detailed cost data. We develop a general framework for dynamic supply estimation that leverages our cost data to address endogeneity bias. Intuitively, our average cost data provide an additional moment that we use to purge unobserved cost shocks, which otherwise lead to biased estimation. The model captures the distribution of producer gains and losses in response to changes in mineral prices. These uneven impacts have geopolitical implications, which we illustrate in our counterfactual analysis.

We apply our empirical framework to three sets of counterfactuals. First, we quantify supply chain vulnerability by removing the top producer from each mineral market. We observe the empirical signature of strong complementarity: removing a producer raises its own mineral price, but it also reduces other mineral prices. Consider Indonesia, which accounts for one-third of global nickel production. The removal of Indonesian nickel leads nickel prices to rise, raising NMC and NCA battery prices. But it also leads lithium and cobalt prices to fall, reducing LFP prices. These opposite effects add richness to the typical energy resource setting, in which competing countries produce substitute goods with prices that rise and fall together.

Second, we simulate unilateral mineral policy. Again consider Indonesia, which exercises its market power by reducing nickel supply and raising nickel prices. This policy action lowers lithium and cobalt prices through complementarity. In terms of geopolitics, Indonesian interests align with those of other nickel producers, which are substitutes that benefit and free-ride as nickel prices rise. But Indonesian interests conflict with those of lithium and cobalt producers, which are complements that suffer as their mineral prices fall. In terms of green adoption, nickel-intensive NMC and NCA batteries become more expensive, while lithium-intensive LFP batteries become less expensive. Indonesian policy effectively taxes the former and subsidizes the latter. Total green adoption declines, but complementarity generates an LFP subsidy that

softens the fall.

Third, we assess the impacts of mineral cartels. Single-mineral cartels coordinate among substitute producers within a given mineral, amplifying the effects of unilateral policy and unambiguously slowing the green transition. Multi-mineral cartels coordinate across complementary minerals. We show that a unified multi-mineral cartel internalizes cross-mineral spillovers and thus sets less aggressive policy than separate competing cartels. Battery prices fall and accelerate green adoption. This result reflects the classic logic of Cournot complementarity (Cournot 1838), by which the merger of complementary producers leads to the removal of double markups. Complementarity can accelerate the green transition.

We build on several bodies of work. First, we contribute to a broad literature on trade and the environment, particularly as it relates to the environmental impacts of unilateral policy (Rauscher 1997, Copeland and Taylor 2003, Kortum and Weisbach 2017, Copeland et al. 2022, Kortum and Weisbach 2024, Hsiao 2026). We show that upstream policies of mineral-endowed countries have global implications through trade. In our setting, high concentration and strong spillovers also have geopolitical implications that relate to a growing literature on geoeconomics (Hirschman 1945, Blackwill and Harris 2016, Mohr and Trebesch 2025, Clayton et al. 2026a,b). We highlight supply chain vulnerability and the potential for cartels in an industry of key importance for global climate targets.

Second, we connect to work on natural resource markets. We relate most directly to studies that emphasize the impacts of market power (Aguirregabiria and Luengo 2016, Asker et al. 2024, Kellogg 2024, De Cannière 2025), including in oil markets with the Organization of the Petroleum Exporting Countries (OPEC).[1] We ask how the same core issue of market concentration will shape the next generation of energy resources. We add an emphasis on complementarity, which introduces a novel margin of global interaction with no direct analog in fossil fuel markets.

Third, we relate to a recent literature on the EV industry (Barwick et al. 2025, Kwon 2025, Allcott et al. 2026, Barwick et al. 2026, Head et al. 2026, Remmy 2026). These studies largely focus on the downstream stages of the supply chain, including

[1] More broadly, this literature also studies the impacts of trade (Farrokhi 2020, Bornstein et al. 2023, Abuin 2025), sanctions (Johnson et al. 2023, Moll et al. 2023, Bachmann et al. 2024, Johnson et al. 2025), and innovation (Popp 2002, Hassler et al. 2021, Alfaro et al. 2025).

the impact of EV subsidies on final consumers. We focus instead on mining and processing upstream, showing how upstream concentration affects downstream adoption. A small set of mineral-endowed countries will shape our progress toward a greener future. We study the geopolitics of that green transition.

2 Background and Data

We introduce our data and document three stylized facts: resource concentration at the country level, heterogeneity in mining costs, and recent growth in battery adoption.

2.1 Institutional details

Critical minerals

Critical minerals have become a strategic priority for governments worldwide. The US Energy Act of 2020 defines criticality broadly, citing "consequences for economic or national security" and vulnerability to disruption (USGS 2026), while Canada's Critical Minerals Strategy emphasizes the "energy transition" and "green economy" (NRC 2022). We thus distinguish between minerals that are critical for the energy transition and those that are critical for national security. We abstract from minerals for national security, including rare earth elements that are more closely tied to military applications, and we focus instead on minerals for the energy transition. This transition calls for the mass production of advanced batteries, which are needed to deploy renewable energy at scale.

Table 1 shows that battery production is a leading driver of global demand for lithium, cobalt, and nickel – our primary minerals of interest. Within batteries, EV batteries dominate at 80% of global battery demand.[2] EV batteries alone account for 65% and 45% of global demand for lithium and cobalt, respectively, with non-EV batteries adding another 20% and 25%. EV batteries are a smaller but fast-growing 11% share of global demand for nickel, which is also a key input for stainless steel production. We exclude graphite because synthetic graphite is already a common

[2] EV batteries account for 80% of annual installed battery capacity. Non-EV battery energy storage systems (BESS) account for another 15%, and portable electronics for 5%.

Table 1: World mineral demand (2022-2024)

	Lithium	Cobalt	Nickel	Graphite	Copper	Manganese
EV batteries	65%	45%	11%	10%	4%	1%
Non-EV batteries	20%	25%	-	15%	-	-
Ceramics and glass	4%	-	-	-	-	-
Metals/alloys	-	19%	67%	-	-	90%
Wiring/infrastructure	-	-	-	-	92%	-
Refractories	-	-	-	40%	-	-
Other	11%	11%	22%	35%	4%	9%

Values for lithium, cobalt, nickel are from IEA (2024). Values for copper, manganese, graphite are from USGS, IRENA, BHP Insights, and Fastmarkets.

alternative to natural graphite, and so supply is not constrained by concentrated geographic endowments. We similarly exclude copper and manganese, which have much smaller shares of demand from battery technologies.

Supply chain

The path from mineral deposits to electric vehicles spans three broad stages: upstream mining and processing, midstream refining and battery manufacturing, and downstream EV manufacturing. Each stage features geographic concentration.

Upstream. Mineral resources take the form of hard rock and liquid brine deposits.[3] Mines extract these mineral resources by excavating the former and surface pumping the latter. Mines then engage in "beneficiation," which refers to initial processing. For hard rock, mineral ore is separated from waste rock. For liquid brine, the extracted liquid is concentrated through evaporation in open-air pools. The product of mining and processing is a mineral concentrate, which mines sell in global markets. For lithium mines, it is spodumene or lithium carbonate depending on their deposit type. For cobalt mines, it is cobalt hydroxide. For nickel mines, it is typically nickel matte or mixed hydroxide precipitate. Upstream production is dominated by Australia for lithium, the DRC for cobalt, and Indonesia for nickel. Furthermore, foreign firms have a significant ownership presence across global mines, including in these countries.

[3] Most minerals are found in hard rock deposits, but South American lithium is typically found in liquid brine deposits.

Midstream. Refiners process mineral concentrates into battery-grade chemical products, including lithium hydroxide and carbonate, cobalt sulfate, and nickel sulfate. Battery manufacturers use these refined minerals to produce cathodes, which are key components of battery cells.[4] China plays a major role in both mineral refining and battery manufacturing, as Chinese firms account for 65 to 80% of global refined lithium, cobalt, and nickel production and roughly 70% of global battery production (IEA 2024, 2025).

Downstream. EV manufacturers use battery cells to produce vehicles for sale to final consumers. The EV manufacturing industry is global, with production concentrated in China, the United States, Europe, South Korea, and Japan. The largest EV producers in these markets are BYD, Tesla, Volkswagen, Hyundai, and Nissan, respectively. China is the largest EV market by volume, but adoption continues to rise across Europe and North America.

Battery technologies

Two battery technologies dominate: nickel manganese cobalt (NMC) and lithium iron phosphate (LFP). Both use lithium, which moves between the cathode and anode when charging and discharging. In NMC batteries, nickel provides energy density, and cobalt provides thermal stability. In LFP batteries, lithium iron phosphate plays both roles. These batteries are imperfect substitutes. NMC batteries have higher cost and higher energy density, which sustains a longer driving range. They are more common in the higher-end vehicles that prevail in Europe and North America. LFP batteries have lower cost and lower energy density, resulting in a shorter driving range. They are better suited for dense urban markets, which are common in China. At the same time, LFP adoption is on the rise. Between 2021 and 2023, the LFP share of the EV market grew from 52% to 67% in China, from 4% to 6% in Europe, and from 2% to 7% in the United States (IEA 2024).

We also consider nickel cobalt aluminum (NCA) batteries, which have experienced slower growth in demand. NCA batteries contain more nickel than NMC batteries, and it is nickel that drives the trade-offs between LFP and NMC. Thus,

[4] The cathode is the primary source of differentiation across battery technologies. Lithium, cobalt, and nickel are cathode inputs. Cathodes are paired with anodes, which are typically made from graphite. Other cell components include electrolytes and separators. Cells are assembled into packs with structural housing, thermal management, and electronics.

relative to NMC batteries, NCA batteries have even higher cost and even longer range. Together, LFP, NMC, and NCA batteries account for 98% of annual EV battery installations (BMI 2025).

2.2 Data sources

We compile detailed data on upstream mining and processing, drawing on three proprietary sources: BMI (2025), GlobalData (2025), and S&P Global (2025). We link these sources to construct a global panel dataset at the mine level, covering lithium, cobalt, and nickel from 2010 to 2024. Our sample includes 397 mines, which account for 92% of world production for lithium, 94% for cobalt, and 80% for nickel. We observe production quantities for both mined ore and processed minerals, as well as production capacity, ore grade, mine type, and mining method.[5] Ownership information allows us to account for ownership across mines by foreign firms. Cost data measure the average costs of production, which we will leverage in our empirical modeling.[6] We observe monthly mineral prices over the sample period.

We do not directly observe midstream refining or battery manufacturing activity. However, we do observe the chemistries that govern the use of minerals in battery manufacturing, which we refer to as battery recipes. We obtain these scientific recipes for each battery technology from Argonne National Laboratory (Knehr et al. 2024). We also observe upstream prices for processed minerals, as well as downstream prices for battery technologies. The gap between these prices is informative of midstream costs and markups, which we recover in estimating our empirical model.

Finally, we observe downstream EV manufacturing, which represents demand for batteries. EV sales data from BMI (2025) record battery installations by technology, make, model, country, and quarter. For example, we observe the number of NMC batteries installed in all Tesla Model S units produced in the United States in the second quarter of 2022. For each battery technology, we observe monthly battery prices from 2021 to 2024.

[5] Mine types include underground and aboveground. Mining methods include open pit, stoping, and block caving.

[6] Costs are broken down by category, including mining costs, processing costs, and royalties.

Figure 1: Market concentration in world production (2015-2024)

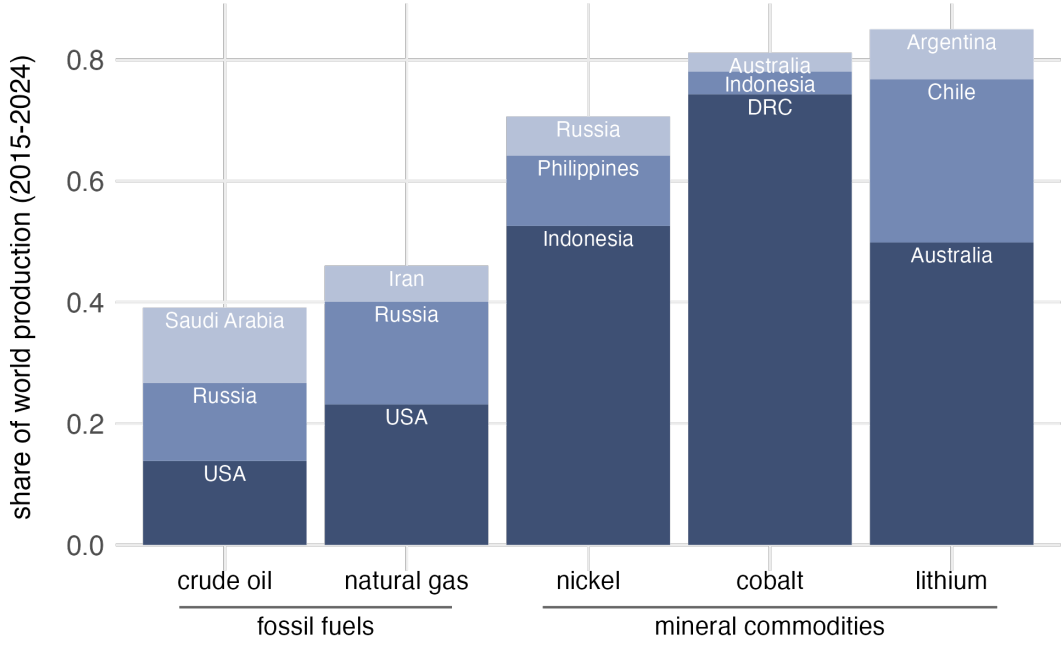

We aggregate production values to the country level from our mine-level data. Crude oil data come from EIA (2025) and include lease condensates. Natural gas data come from IEA (2023).

2.3 Stylized facts

Mineral endowments are geographically concentrated

For both minerals and fossil fuels, unequal geological endowments lead to concentrated production. Minerals markets are especially concentrated. Figure 1 shows that the market share of the top three producing countries ranges from 40 to 50% for fossil fuels, but a substantially higher 70 to 85% for mineral commodities. This pattern also holds when concentration is measured by reserves instead of production. Concentration is particularly high for lithium and cobalt, which are more nascent markets, and it remains high for nickel, which is a more mature market.

Mining costs are heterogeneous

Our mine-level data reveal substantial heterogeneity in production costs across mines. Geological fundamentals drive this heterogeneity. Across minerals, deposits vary in their accessibility. Lithium in South America exists in underground brine

Figure 2: Cobalt production across mines

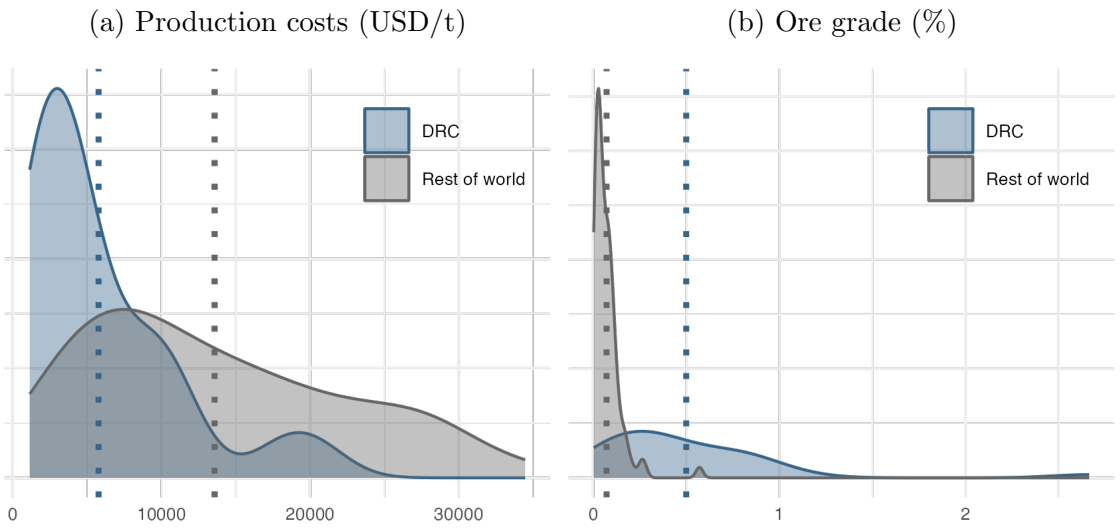

On the left, we plot the density of average production costs across mines. On the right, we plot the density of average ore grades across mines. In both cases, we average over time for each mine.

deposits, which require capital-intensive pumping to the surface, while cobalt in the DRC can be mined artisanally because it is already accessible at surface level. Within minerals, deposits vary in their purity. Figure 2 illustrates this fact for cobalt. Average ore grade in the DRC is five times as high as ore grade elsewhere (left panel), and higher ore grade translates directly into lower production costs (right panel).

Mineral recipes vary by battery technology

Battery technologies differ in their mineral inputs. Figure 3a shows battery recipes for LFP, NMC, and NCA batteries. These recipes capture interdependence across minerals: lithium, cobalt, and nickel are used jointly in producing NMC and NCA batteries, and this joint use creates complementarity. Figure 3b shows the growing adoption of LFP and NMC batteries over the last decade. LFP market share grew sharply from 5% to 60% between 2020 and 2024, as the expiration of key patents allowed for rapid global diffusion.[7] LFP overtook NMC as the top battery technology by 2024.

[7] The core LFP patents were largely developed in North America, with follow-on advances in China allowing for LFP production at commercial scale. These patents created frictions with licensing requirements outside of China. The core patents had largely expired by 2022, reducing the barriers to commercial adoption outside of China.

Figure 3: Battery technologies

(a) Battery recipes (kg/kWh)

	nickel	lithium	cobalt
LFP	-	0.09	-
NMC	0.47	0.11	0.18
NCA	0.65	0.1	0.12

(b) Global adoption (TWh)

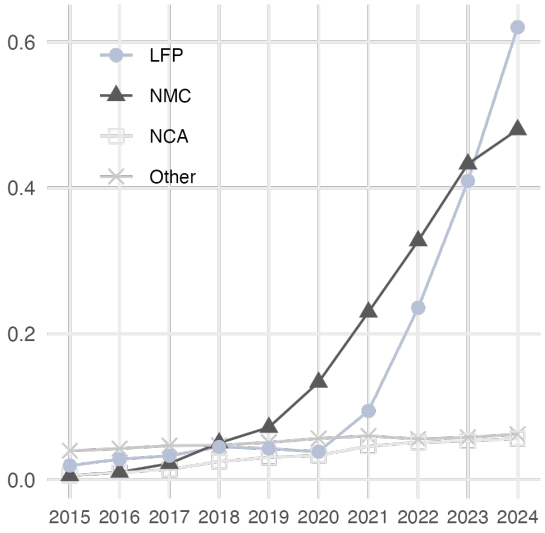

On the left, we plot the recipe matrix for selected battery technologies and minerals. NMC values average across all NMC varieties. Appendix A.1 shows the full recipe matrix. On the right, we plot newly installed capacity worldwide for each battery technology with data from BMI (2025).

3 Theory

We present a stylized model with two minerals, which serve as inputs for battery production. We highlight both substitution and complementarity across minerals. We characterize the impacts of market power at the country level, and we discuss implications for geopolitics and the green transition.

Minerals and batteries

We focus on lithium and nickel. Consider two mining countries $m \in \{\ell, n\}$ and two battery technologies $j \in \{L, N\}$. Country ℓ produces only lithium, and country n produces only nickel. These minerals become batteries based on scientific recipes. We define r_j^m as the units of mineral m required to produce one unit of battery j. Battery L uses only lithium with a recipe given by $r_L^\ell > 0$ and $r_L^n = 0$. Battery N uses both lithium and nickel with a recipe given by $r_N^\ell > 0$ and $r_N^n > 0$. We use superscripts for upstream minerals m and subscripts for downstream batteries j.

Supply and demand

Mining countries supply minerals $m \in \{\ell, n\}$. Country m supplies s^m units of its mineral as a function of world price p^m and mineral policy τ^m.

$$s^m(p^m - \tau^m), \quad \sigma^m = \frac{\mathrm{d}s^m}{\mathrm{d}(p^m - \tau^m)} > 0$$

We denominate the value of mineral policy in dollars. Restrictive policy τ^m reduces the quantity supplied with the goal of raising world prices.

Consumers demand batteries $j \in \{L, N\}$. Batteries are substitutes, and demand d_j is given by

$$d_j(p_L, p_N), \quad \delta_{jj} = \frac{\partial d_j}{\partial p_j} < 0, \quad \delta_{jj'} = \frac{\partial d_j}{\partial p_{j'}} > 0 \quad \text{for} \quad j \neq j'.$$

When consumers can switch across batteries, own-price demand effects are negative and cross-price effects are positive. We maintain the standard regularity condition that demand Jacobian $\boldsymbol{D} = [\delta_{jj'}]$ is negative semidefinite, such that own-price effects dominate cross-price effects in magnitude. Furthermore, when battery prices rise, total battery demand d falls.

$$d(p_L, p_N) = d_L(p_L, p_N) + d_N(p_L, p_N), \quad \delta_j = \frac{\partial d}{\partial p_j} < 0$$

In equilibrium, mineral prices clear mineral markets. We have specified mineral supply and battery demand, which yield mineral demand as follows.

$$d^\ell = d_L r_L^\ell + d_N r_N^\ell, \quad d^n = d_N r_N^n$$

Prices equalize supply and demand for each mineral.

$$s^m(p^m - \tau^m) = d^m(p^\ell, p^n) \quad \forall\, m \in \{\ell, n\} \tag{1}$$

If battery manufacturing is perfectly competitive, then mineral prices also pin down battery prices.

$$p_L = p^\ell r_L^\ell, \quad p_N = p^\ell r_N^\ell + p^n r_N^n \tag{2}$$

Substitution and complementarity

Mineral demand features substitution and complementarity, which generate spillovers across minerals. The cross-price effect of nickel prices on lithium demand is

$$\frac{\partial d^{\ell}}{\partial p^{n}} = \delta_{LN} r_{L}^{\ell} r_{N}^{n} + \delta_{NN} r_{N}^{\ell} r_{N}^{n}. \tag{3}$$

On one hand, $\delta_{LN} r_{L}^{\ell} r_{N}^{n} > 0$ captures *substitution*. Higher nickel prices raise the price of battery N, and so demand shifts toward battery L, which becomes cheaper in relative terms. Lithium demand rises as L expands and uses more lithium. On the other hand, $\delta_{NN} r_{N}^{\ell} r_{N}^{n} < 0$ captures *complementarity*. Higher nickel prices reduce demand for battery N, which also uses lithium. Lithium demand falls as N shrinks and uses less lithium. The net effect depends on whether substitution or complementarity dominates.

Definition 1 (Gross substitutes)**.** Lithium and nickel are gross substitutes when substitution dominates, such that $\frac{\partial d^{\ell}}{\partial p^{n}} > 0$.

Definition 2 (Gross complements)**.** Lithium and nickel are gross complements when complementarity dominates, such that $\frac{\partial d^{\ell}}{\partial p^{n}} < 0$.

Geopolitics

We study the distributional impacts of policy action τ^{n}, which country n implements unilaterally. Differentiating equilibrium condition 1 with respect to τ^{n}, we obtain the own- and cross-price effects of restrictive nickel policy.

$$\frac{\mathrm{d}p^{n}}{\mathrm{d}\tau^{n}} = \frac{\sigma^{n}\left(\sigma^{\ell} - \frac{\partial d^{\ell}}{\partial p^{\ell}}\right)}{\left(\sigma^{\ell} - \frac{\partial d^{\ell}}{\partial p^{\ell}}\right)\left(\sigma^{n} - \frac{\partial d^{n}}{\partial p^{n}}\right) - \frac{\partial d^{\ell}}{\partial p^{n}}\frac{\partial d^{n}}{\partial p^{\ell}}}, \qquad \frac{\mathrm{d}p^{\ell}}{\mathrm{d}\tau^{n}} = \left(\frac{\frac{\partial d^{\ell}}{\partial p^{n}}}{\sigma^{\ell} - \frac{\partial d^{\ell}}{\partial p^{\ell}}}\right)\frac{\mathrm{d}p^{n}}{\mathrm{d}\tau^{n}}.$$

Proposition 1 (Mineral effects)**.** Restrictive nickel policy raises nickel prices. It raises lithium prices if and only if lithium and nickel are gross substitutes, and it reduces lithium prices if and only if lithium and nickel are gross complements.

Proof. For nickel prices, we show that $\frac{\mathrm{d}p^{n}}{\mathrm{d}\tau^{n}} > 0$. The denominator is positive: it is the determinant of market-clearing matrix $\boldsymbol{M} = \mathrm{diag}(\sigma^{\ell}, \sigma^{n}) - \boldsymbol{R}'\boldsymbol{D}\boldsymbol{R}$ for recipe matrix $\boldsymbol{R} = [r_{j}^{m}]$, and negative semidefinite $\boldsymbol{D}$ implies negative semidefinite $\boldsymbol{R}'\boldsymbol{D}\boldsymbol{R}$ and

positive definite $\boldsymbol{M}$, such that $\det \boldsymbol{M} > 0$. The numerator is also positive: negative semidefinite $\boldsymbol{D}$ implies $\frac{\partial d^\ell}{\partial p^\ell} = (\boldsymbol{r}^\ell)' \boldsymbol{D} \boldsymbol{r}^\ell \leq 0$ for $\boldsymbol{r}^\ell = (r_L^\ell, r_N^\ell)'$, such that $\sigma^\ell - \frac{\partial d^\ell}{\partial p^\ell} > 0$. For lithium prices, given $\frac{\mathrm{d}p^n}{\mathrm{d}\tau^n} > 0$ and $\sigma^\ell - \frac{\partial d^\ell}{\partial p^\ell} > 0$ above, the sign of $\frac{\mathrm{d}p^\ell}{\mathrm{d}\tau^n}$ is determined by the sign of $\frac{\partial d^\ell}{\partial p^n}$. Definitions 1 and 2 then give the result. □

These policy effects have geopolitical implications. Country n exercises market power through τ^n, restricting nickel supply and raising nickel prices to increase its producer surplus. This action spills over to country ℓ through lithium prices. If lithium and nickel are gross substitutes, then country ℓ free rides: it gains producer surplus as lithium prices rise. The countries' interests are aligned. By contrast, if lithium and nickel are gross complements, then country ℓ suffers: it loses producer surplus as lithium prices fall. The countries' interests are in conflict.

Green adoption

We turn to aggregate implications for the green transition. Batteries L and N are equally effective at reducing carbon emissions, and so the green transition depends only on total adoption $d = d_L + d_N$. Differentiating with respect to τ^n, we obtain the effect of restrictive nickel policy.

$$\frac{\mathrm{d}d}{\mathrm{d}\tau^n} = \delta_L \frac{\mathrm{d}p_L}{\mathrm{d}\tau^n} + \delta_N \frac{\mathrm{d}p_N}{\mathrm{d}\tau^n}$$

Proposition 2 (Battery effects)**.** Under pure substitution ($r_N^\ell = 0$), restrictive nickel policy raises both battery prices and reduces total green adoption. Under pure complementarity ($\delta_{LN} = \delta_{NL} = 0$), it taxes battery N but subsidizes battery L with ambiguous effects on total green adoption.

Proof. Under pure substitution, $r_N^\ell = 0$ implies that $\frac{\partial d^\ell}{\partial p^n} = \delta_{LN} r_L^\ell r_N^n > 0$, and so $\frac{\mathrm{d}p^\ell}{\mathrm{d}\tau^n}, \frac{\mathrm{d}p^n}{\mathrm{d}\tau^n} > 0$ by proposition 1. It follows that $\frac{\mathrm{d}p_L}{\mathrm{d}\tau^n}, \frac{\mathrm{d}p_N}{\mathrm{d}\tau^n} > 0$ by equation 2. Thus, $\frac{\mathrm{d}d}{\mathrm{d}\tau^n} = \delta_L \frac{\mathrm{d}p_L}{\mathrm{d}\tau^n} + \delta_N \frac{\mathrm{d}p_N}{\mathrm{d}\tau^n} < 0$ for $\delta_L, \delta_N < 0$. Under pure complementarity, $\delta_{LN} = \delta_{NL} = 0$ implies $\frac{\partial d^\ell}{\partial p^n} = \delta_{NN} r_N^\ell r_N^n < 0$, and so $\frac{\mathrm{d}p^\ell}{\mathrm{d}\tau^n} < 0$ and $\frac{\mathrm{d}p^n}{\mathrm{d}\tau^n} > 0$ by proposition 1. It follows that $\frac{\mathrm{d}p_L}{\mathrm{d}\tau^n} = r_L^\ell \frac{\mathrm{d}p^\ell}{\mathrm{d}\tau^n} < 0$ and $\frac{\mathrm{d}p_N}{\mathrm{d}\tau^n} = r_N^n \frac{\mathrm{d}p^n}{\mathrm{d}\tau^n} [\sigma^\ell - \delta_{LL}(r_L^\ell)^2] / [\sigma^\ell - \frac{\partial d^\ell}{\partial p^\ell}] > 0$ by equation 2. Thus, $\frac{\mathrm{d}d_L}{\mathrm{d}\tau^n} > 0$ but $\frac{\mathrm{d}d_N}{\mathrm{d}\tau^n} < 0$, such that $\frac{\mathrm{d}d}{\mathrm{d}\tau^n}$ is ambiguous in sign. □

Substitution and complementarity have distinct implications for the green transition. At one extreme with $r_N^\ell = 0$, we eliminate joint mineral use in battery N.

Lithium and nickel are pure substitutes, and both battery prices rise. Restrictive nickel policy thus acts as a tax on both L and N, such that green adoption falls. At the other extreme with $\delta_{LN} = \delta_{NL} = 0$, we eliminate consumer substitution across L and N. Lithium and nickel are pure complements, and p_L falls while p_N rises. Restrictive nickel policy thus acts as a subsidy to L and a tax on N: L adoption rises, N adoption falls, and the effect on total green adoption is ambiguous. Our empirical analysis will resolve the ambiguity by decomposing and quantifying these forces in the data.

4 Demand

We model the demand for battery technologies with an almost ideal demand system (Deaton and Muellbauer 1980). We estimate the model with regional panel data on battery technologies used by EV manufacturers.

4.1 Model

In each region k and period t, EV manufacturers choose among battery technologies $j \in \mathcal{J}$. They do so by allocating a battery budget X_{kt} across batteries j, taking world prices p_{jt} as given. We specify their demand in product space with expenditure shares w_{jkt} given by

$$w_{jkt} = \alpha_{jt} + \beta_j \log\left(\frac{X_{kt}}{P_t}\right) + \sum_{j' \in \mathcal{J}} \gamma_{jj'} \log p_{j't} + \varepsilon_{jkt}, \tag{4}$$

where battery price index P_t is of translog form

$$\log P_t = \sum_{j \in \mathcal{J}} \alpha_{jt} \log p_{jt} + \frac{1}{2} \sum_{j,j' \in \mathcal{J}^2} \gamma_{jj'} \log p_{jt} \log p_{j't}. \tag{5}$$

Quality α_{jt} captures unobservables that vary across products and over time, but are common across regions. Income effects β_j allow expenditure shares to respond heterogeneously to changes in the battery budget, which may increase over time as battery technologies diffuse and adoption rates rise.[8] Price index P_t deflates the

[8] In our data, we observe variation in total battery spending X_{kt} that captures a range of factors affecting EV vehicle adoption among final consumers. These factors include fossil fuel prices, con-

battery budget, and semi-elasticities $\gamma_{jj'}$ govern own- and cross-price elasticities. We allow for unobserved regional demand shocks ε_{jkt}.

Our demand specification is in product space. The main advantage to modeling demand in this way is that we allow for flexible patterns of substitution across products. This approach is feasible because we study a small choice set of batteries. For a choice set $\mathcal{J}$ that includes J products, the full semi-elasticity matrix includes J^2 parameters. Estimation encounters a curse of dimensionality and becomes infeasible with many products. The more typical approach is to model demand in characteristic space, as in Berry (1994) and Berry et al. (1995). By projecting products onto their observed characteristics, this approach solves the dimensionality problem by reducing the number of parameters to be estimated. The cost is less flexibility in capturing the full range of cross-product substitution patterns.

We compute quantities demanded as follows. For battery prices $\boldsymbol{p}_t = \{p_{jt}\}_{j \in \mathcal{J}}$, the quantity demanded d_{jkt} of battery j in region k and period t is

$$d_{jkt}(\boldsymbol{p}_t) = \frac{X_{kt} w_{jkt}(\boldsymbol{p}_t)}{p_{jt}}$$

by definition of expenditure share w_{jkt}. World demand aggregates across regions.

$$d_{jt}(\boldsymbol{p}_t) = \sum_{k \in \mathcal{K}} d_{jkt}(\boldsymbol{p}_t) \tag{6}$$

4.2 Estimation

We estimate the model with iterated linear least squares (ILLS), as in Blundell and Robin (1999). The first estimation challenge is nonlinearity. Because of translog price index P_t, parameters α_{jt} and $\gamma_{jj'}$ enter nonlinearly in expenditure shares w_{jkt}. However, expenditure shares are linear in parameters when we condition on a fixed value of P_t. We exploit this linearity for estimation. First, we initialize estimation by computing P_t from data with the Stone price index, which is the expenditure-weighted average of prices p_{jt}. Second, we estimate expenditure share equation 4 by linear regression. Third, we use the resulting parameter estimates to recompute

sumer preferences, and government policies. We treat this battery spending by EV manufacturers as a broad summary measure of category demand for electric vehicles.

P_t with equation 5. Fourth, we iterate until convergence. Throughout, we impose the standard restrictions of adding-up, homogeneity, and symmetry. Appendix B provides additional details.

The second estimation challenge is price endogeneity: battery prices p_{jt} are correlated with demand shocks ε_{jkt}. We address this concern by instrumenting with cost shifters in the form of mineral prices. Relevance holds because minerals are inputs for batteries. Mineral prices shift battery costs and thus battery prices. Exclusion requires us to distinguish between two sets of minerals, $\mathcal{M}$ and $\mathcal{M}'$. The first set $\mathcal{M}$ consists of lithium, cobalt, and nickel, which are the main minerals of focus in this paper. EV batteries are a major part of global demand for these minerals, and so their prices are likely correlated with battery demand shocks. The second set $\mathcal{M}'$ consists of aluminum, phosphoric acid, iron, and manganese. These minerals are inputs in EV battery production, but EV battery production is less than 3% of global demand for these minerals. Their prices are therefore plausibly uncorrelated with battery demand shocks.

We therefore instrument for battery prices p_{jt} with mineral prices $\{p^{mt}\}^{m\in\mathcal{M}'}$ from the second set $\mathcal{M}'$. We do so by constructing predicted prices z_{jt}, which weight input minerals according to their use in battery recipes r_j^m.

$$z_{jt} = \sum_{m\in\mathcal{M}'} r_j^m p^{mt} \tag{7}$$

Figure 4 shows that mineral prices in the first set $\mathcal{M}$ move closely with battery prices, in contrast to those in the second set $\mathcal{M}'$. Our instrument leverages the second, more orthogonal source of variation. Figure 4 also shows that battery prices move closely with each other. Our instrument disentangles these battery prices with battery-specific recipes r_j^m, by which variation in mineral prices generates differential variation in battery prices.

Our estimation sample is a quarterly panel from 2021 Q1 to 2025 Q1 with four regions of EV manufacturers: China, North America, Europe, and Asia ex-China. Battery prices are in units of USD per MWh of capacity. The choice set includes two battery technologies that account for 98% of global EV battery installations: battery L is LFP, and battery N aggregates over NMC and NCA.[9] We combine NMC and

[9] We omit the remaining 2%, which consists of EV battery technologies with limited commercial

Figure 4: Battery and mineral prices

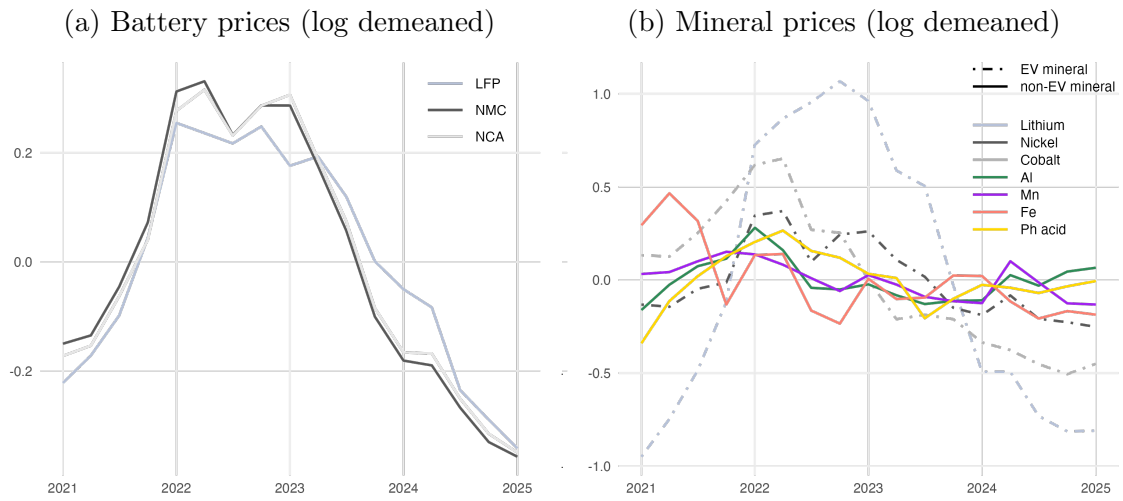

Mineral prices are from S&P Global (2025), and battery prices are from BMI (2025). All prices are averaged at the quarterly level and log demeaned.

NCA because their recipes and prices are particularly highly correlated, as seen in Figures 3 and 4. Furthermore, this grouping aligns to first order with the mineral endowments that drive our geopolitical impacts of interest. Mineral policy in nickel- and cobalt-endowed countries must consider substitution away from NMC and NCA batteries, which use nickel and cobalt, toward LFP batteries, which do not.

4.3 Estimates

Figure 5 reports our estimated demand parameters and elasticities. We begin with the estimates from our baseline IV specification, which the figure displays in black. Panel A shows lower estimated quality α_{jt} for battery L, rationalizing lower shares despite lower prices in the data. Panel B shows larger income effects β_j for L, reflecting that rising aggregate battery expenditure coincides with the rapid growth of LFP adoption during our sample period. Panel C shows that the own-price semi-elasticities γ_{jj} are negative, while the cross-price semi-elasticities $\gamma_{jj'}$ are positive. These parameter estimates determine the Marshallian elasticities in panel D, with similarly negative own-price elasticities (LL and NN) and positive cross-price elas-

deployment (e.g., sodium ion). The aggregate expenditure share sums over NMC and NCA. The aggregate recipe and price are expenditure-weighted averages.

Figure 5: Demand estimates

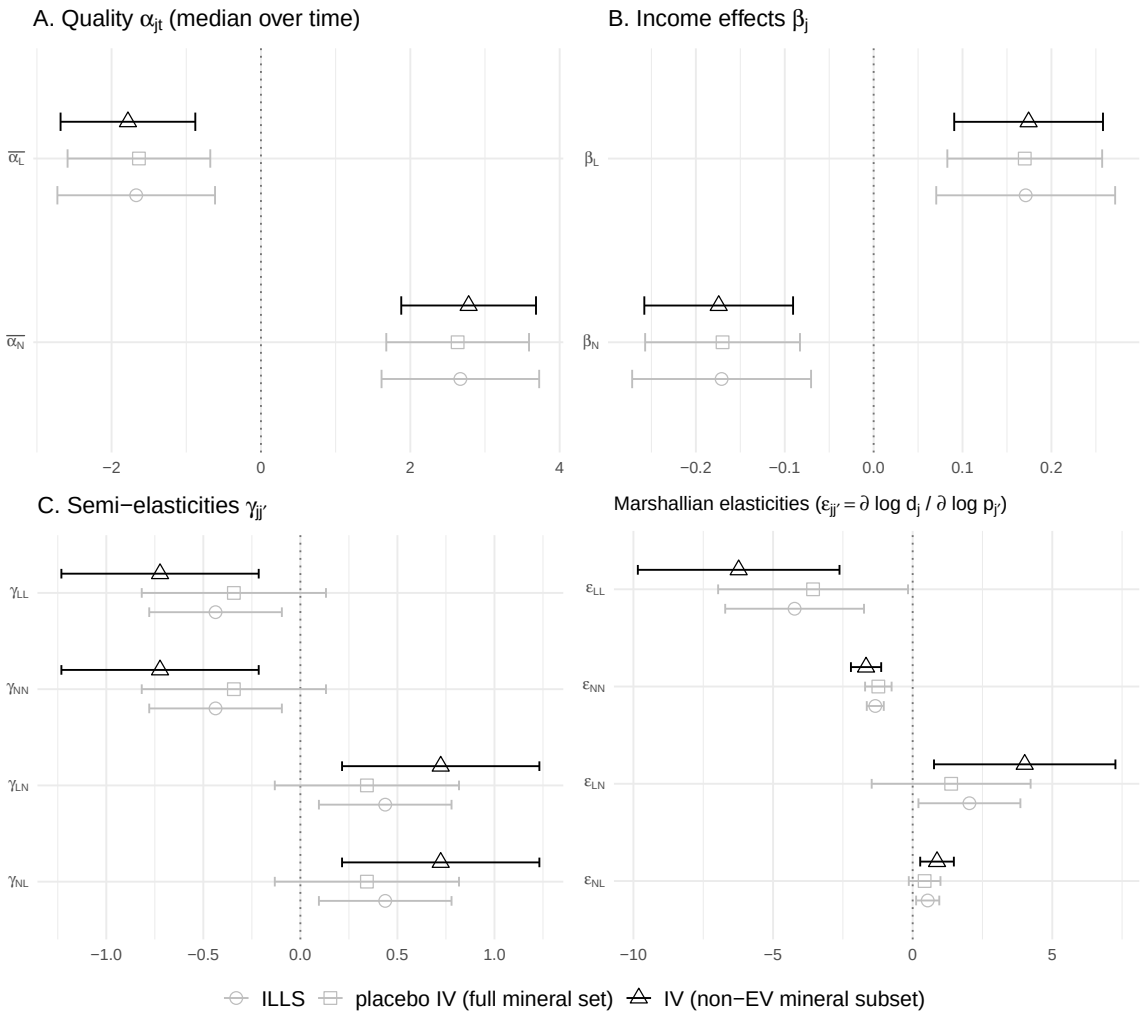

The vertical axis of each panel labels the subindices of each parameter. We compute Marshallian elasticities as $\frac{\partial \log d_j}{\partial \log p_{j'}} = \frac{1}{\bar{w}_j}(\gamma_{jj'} - \beta_j \bar{w}_{j'}) - \mathbb{1}[j = j']$, where $\bar{w}_j$ is the worldwide expenditure share averaged over time for battery j. We cluster standard errors by regional market. We compute Marshallian elasticity standard errors using the delta method. The first-stage F-statistic for our main IV specification is 7302.

ticities (LN and NL). We find that demand is elastic for both L and N, with elasticities that exceed one in magnitude, and that demand is particularly elastic for L. This result is consistent with the use of LFP batteries in lower-end vehicles, which have shorter driving ranges and serve more price-sensitive consumers.

We compare our baseline IV estimates to two additional benchmarks in Figure 5: estimation without instruments and a placebo IV that uses the full set of minerals

$\mathcal{M} \cup \mathcal{M}'$. Estimation without instruments suffers from price endogeneity, as demand shocks and battery prices are positively correlated. Panel D shows that the resulting estimates are biased toward zero. By comparison, our baseline IV estimates are roughly twice as large in magnitude. The placebo IV adds the first set of minerals $\mathcal{M}$ to the instruments. EV demand drives global demand for these minerals, and so their prices remain correlated with battery demand shocks ε_{jkt}. Instruments based on $\mathcal{M}$ thus fail the exclusion restriction, and panel D confirms that the placebo estimates remain similar to the estimates without instruments. Our baseline IV corrects this simultaneity bias by excluding $\mathcal{M}$ from the instrument set.

5 Supply

We model the supply of minerals with a dynamic model of extraction, and we estimate the model with mine-level panel data on production quantities and costs. The model delivers country-level supply curves that characterize the distribution of producer rents across mineral-endowed countries.

5.1 Cost data

We observe the average costs of extraction at the mine level, and we seek to leverage these data to construct supply curves. Figure 6 illustrates with global data for lithium, plotting costs against cumulative capacity with mines ordered by cost. This empirical cost curve represents a global supply curve under three restrictive assumptions: myopia, linear costs, and price taking. To characterize supply, we must consider marginal revenues and costs for each mine. Under myopia, we can restrict attention to current revenues and costs. Under linear costs, marginal costs equal average costs, and so we can read marginal costs from our average cost data. Under price taking, marginal revenues equal market prices, and so we can read marginal revenues from our price data. We thus obtain a supply curve directly from the data. It predicts that mines produce to capacity when world prices exceed their observed average costs.

Myopia and linear costs are strong assumptions. Myopia is in tension with the non-renewable nature of mineral resources, which are in finite supply. Linear costs of-

Figure 6: Lithium production costs across mines (USD/t)

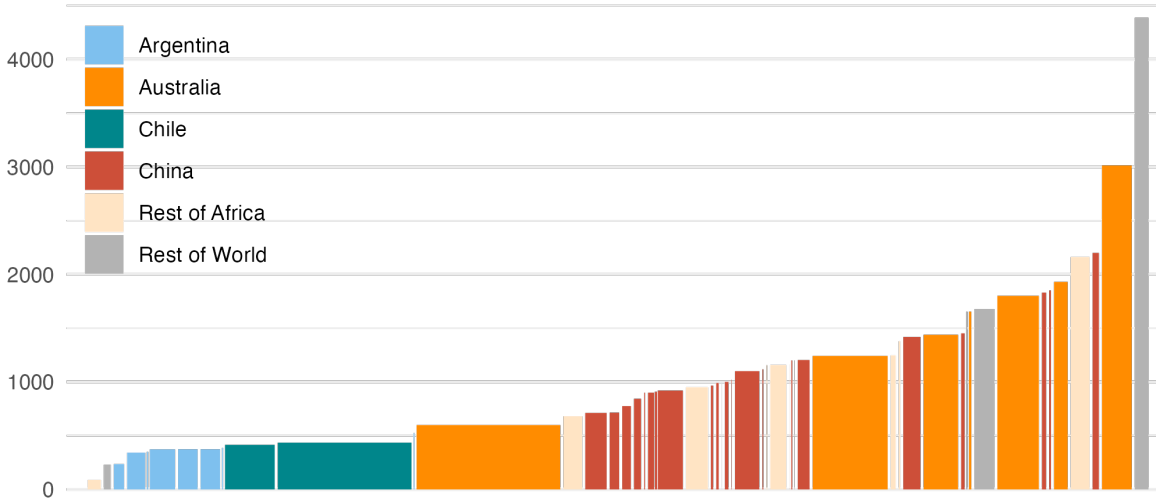

Mine-level cost data for 2024 come from BMI (2025). Each bar is a mine. Bar heights indicate realized average costs in USD per ton of lithium carbonate equivalent. Bar widths are proportional to production capacities.

fer a local approximation that is appropriate only for evaluating small policy changes. Linear costs also imply that mines produce either at full capacity or not at all, which is in conflict with the data. Our model relaxes both assumptions. We maintain price taking at the mine level, as is consistent with the large number of mines competing in these global commodity markets.

5.2 Model

We model forward-looking mines with finite reserves of minerals and convex costs of extraction, focusing our exposition on our baseline specification of costs. Appendix C.1 presents derivations, and Appendix C.2 shows that our approach generalizes to a broad class of cost specifications.

Extraction

In each period t, a mine i in country $k(i)$ extracts raw ore x^{it}. It processes this ore to produce mineral $m(i)$, which it sells in the world market. Production s^{it}

depends on the grade g^i of the extracted ore x^{it}.

$$s^{it} = g^i x^{it} \tag{8}$$

Marginal revenue mr^{it} is given by world price p^{mt} and ore grade g^i.

$$mr^{it} = p^{mt} g^i \tag{9}$$

We assume that mines take these world prices as given, noting that individual mines are small despite geographic concentration at the country level. We also impose constant grade g^i over time. In principle, ore quality may degrade with extraction if the highest-quality deposits are extracted first. Appendix C.3 shows how to accommodate the added dynamics from ore degradation, while also noting that ore degradation is limited in our particular empirical setting.

Cost convexity

We specify costs that are quadratic in extraction x^{it} and scale with capacity $\bar{x}^i$ through utilization rate u^{it}.

$$C^{it} = \left(\alpha^{km} + \frac{1}{2}\kappa^{km} u^{it} + \varepsilon^{it}\right) x^{it}, \quad u^{it} = \frac{x^{it}}{\bar{x}^i}, \tag{10}$$

where α^{km} and κ^{km} vary by country and mineral. Cost level α^{km} captures the linear component of costs, which scales proportionally with extraction. Cost convexity κ^{km} captures a quadratic component of costs, which rises as mines approach or potentially exceed their extraction capacity. Unobserved cost shock ε^{it} is additively separable and linear. Appendix C.2 shows that we are not restricted to quadratic convex costs, and that we require only separability and not linearity in our cost shocks. From equation 10, it follows directly that average and marginal costs are

$$ac^{it} = \alpha^{km} + \frac{1}{2}\kappa^{km} u^{it} + \varepsilon^{it}, \tag{11}$$

$$mc^{it} = \alpha^{km} + \kappa^{km} u^{it} + \varepsilon^{it}. \tag{12}$$

If $\kappa^{km} > 0$, we have cost convexity with marginal costs that exceed average costs. If $\kappa^{km} = 0$, we have cost linearity and $ac^{it} = mc^{it}$.

Dynamics

Mines extract from finite reserves with a stock R^{it} that follows a law of motion

$$R^{it+1} = R^{it} - x^{it} \geq 0.$$

The Euler equation equates the net payoff of extracting today to its discounted expected value tomorrow.

$$mr^{it} - mc^{it} = \beta \mathbb{E}^{it}\left[mr^{it+1} - mc^{it+1}\right] \tag{13}$$

This expression is the Hotelling (1931) condition. The net payoff is the shadow cost of reserves, where $\lambda^{it} = mr^{it} - mc^{it}$. This net payoff grows at rate $1/\beta$ for discount factor β. Substituting equations 11 and 12 and rearranging, utilization is given by

$$u^{it} = \frac{2}{\kappa^{km}}\left(mr^{it} - ac^{it} - \lambda^{it}\right), \tag{14}$$

reflecting observed average costs ac^{it}, convexity κ^{km}, and shadow costs λ^{it}. This expression nests the empirical cost curve that we describe above, in which mines operate at full capacity when marginal revenue exceeds average cost. With myopia ($\lambda^{it} = 0$), linear costs ($\kappa^{km} \to 0$), and binding capacity, we have $u^{it} = \mathbb{1}[mr^{it} > ac^{it}]$.

Production

We compute quantities supplied as follows. By equations 8 and 10, the quantity supplied s^{it} by a mine i in period t is given by production

$$s^{it}(\boldsymbol{p}^{mt}) = g^i \bar{x}^i u^{it}(\boldsymbol{p}^{mt}),$$

which depends on the mineral price path $\boldsymbol{p}^{mt} = \{p^{mt'}\}^{t' \geq t}$. For these forward-looking mines, production depends on both current and future prices. World supply aggregates across mines.

$$s^{mt}(\boldsymbol{p}^{mt}) = \sum_{i \in \mathcal{I}^m} s^{it}(\boldsymbol{p}^{mt}) \tag{15}$$

5.3 Estimation

We estimate cost parameters $(\alpha^{km}, \kappa^{km})$ with mine-level panel data. We show that our cost data allow us to address endogeneity concerns in estimation.

Euler estimation

We begin with the typical Euler approach, which derives an estimating equation from the Euler condition. We substitute equations 9 and 12 into Euler equation 13, and we take quasi-differences with notation $\Delta y^{it} = \beta y^{it} - y^{it-1}$.

$$\Delta \mathbb{E}^{it-1}\left[p^{mt} g^{i}\right] = \Delta \mathbb{E}^{it-1}\left[\alpha^{km} + \kappa^{km} u^{it} + \varepsilon^{it}\right]$$

In expectation, the change in marginal revenue on the left-hand side is equal to the change in marginal cost on the right-hand side. We obtain an estimating equation by rewriting expected values in terms of realized values and expectational errors.

$$\Delta p^{mt} g^{i} = \Delta \alpha^{km} + \kappa^{km} \Delta u^{it} + \Delta \varepsilon^{it} + \eta^{it} \tag{16}$$

Expectational errors η^{it} capture the difference between realized and expected values. For $\tilde{\Delta} y^{it} = \Delta y^{it} - \Delta \mathbb{E}^{it-1}[y^{it}]$, we have $\eta^{it} = \tilde{\Delta} p^{mt} g^{i} - \kappa^{km} \tilde{\Delta} u^{it} - \tilde{\Delta} \varepsilon^{it}$. Positive values represent positive surprises. Under rational expectations, these expectational errors are uncorrelated with any variables contained in information set $\mathcal{N}^{it-1}$ in period $t-1$, such that $\mathbb{E}^{it-1}[\eta^{it} \,|\, \mathcal{N}^{it-1}] = 0$ (Hall 1978).

Endogeneity

We highlight two endogeneity problems that arise from error terms $\Delta \varepsilon^{it}$ and η^{it}. First, cost shocks ε^{it} enter average costs by equation 11, and so they are correlated with utilization u^{it} by equation 14. A positive cost shock reduces utilization, and the result is downward bias in our estimate of κ^{km}. Second, by their definitions, difference Δu^{it} is increasing in u^{it}, while expectational error η^{it} is decreasing in u^{it}.[10] The result is again downward bias in our estimate of κ^{km}.

We address the first endogeneity problem by inverting equation 11, which allows

[10] Expanding yields $\Delta u^{it} = \beta u^{it} - u^{it-1}$ and $\eta^{it} = -\beta \kappa^{km} u^{it} + \beta \kappa^{km} \mathbb{E}^{it-1}[u^{it}] + \tilde{\Delta} p^{mt} g^{i} - \tilde{\Delta} \varepsilon^{it}$.

us to express the cost shocks as a function of our data.

$$\varepsilon^{it} = ac^{it} - \frac{1}{2}\kappa^{km}u^{it} - \alpha^{km}$$

The value of our cost data is that they provide this additional moment, which we substitute into equation 16 to eliminate cost shocks ε^{it}. We obtain our baseline specification for estimation.

$$\Delta p^{mt}g^{i} - \Delta ac^{it} = \frac{1}{2}\kappa^{km}\Delta u^{it} + \eta^{it} \tag{17}$$

We address the second endogeneity problem by instrumenting for differenced utilization Δu^{it} with lagged utilization $u^{it-1} \in \mathcal{N}^{it-1}$. The lagged value is correlated with Δu^{it} by definition and orthogonal to η^{it} by rational expectations.

To estimate equation 17, we construct the left-hand side terms from observed prices p^{mt}, grades g^{i}, and average costs ac^{it}, setting $\beta = 0.9$. We regress on observed difference Δu^{it} to obtain cost convexity $\hat{\kappa}^{km}$ with errors $\hat{\eta}^{it}$ as residuals. The regression does not identify cost levels α^{km}, which we recover using average cost equation 11.

$$ac^{it} - \frac{1}{2}\hat{\kappa}^{km}u^{it} = \alpha^{km} + \varepsilon^{it} \tag{18}$$

Country-mineral fixed effects isolate cost levels $\hat{\alpha}^{km}$ with shocks $\hat{\varepsilon}^{it}$ as residuals. Intuitively, the estimated $\frac{1}{2}\hat{\kappa}^{km}u^{it}$ term corrects for cost convexity in recovering α^{km} from our average cost data.[11] Appendix C.4 shows that mismeasured costs do not bias estimation.

Discussion

The average cost data allow us to use average cost equation 11 as an additional moment. We use this moment to eliminate cost shocks ε^{it} from estimating equation 17 and thus to address the endogeneity bias that arises from these unobserved cost shocks. Without this additional moment, we would instead require an instrument that is uncorrelated with the shocks. This approach applies both in static settings, where estimation is based on a first order condition, and in dynamic settings, where

[11] By contrast, the empirical cost curve approach of Section 5.1 fixes marginal costs at the average costs observed in the data. This approach is equivalent to imposing $\kappa^{km} = 0$ in equation 18.

estimation is based on an Euler equation.[12]

We draw a contrast to other approaches for leveraging cost data. Asker et al. (2019) and Clausing et al. (2025) assume constant marginal costs and treat average cost data as direct measures of marginal costs. This approach treats the empirical cost curve as a supply curve, interpreting observed average costs as marginal cost levels α^{km}. Our framework nests this special case with $\kappa^{km} = 0$ and thus $ac^{it} = mc^{it}$, but we do not impose such restrictions.

5.4 Estimates

Figure 7 reports our estimated supply parameters and elasticities. The top row shows our cost parameter estimates: top producers in each mineral market have higher cost convexity κ^{km} but lower cost levels α^{km} than the rest of the world. On one hand, lower cost levels are an advantage that reflects higher levels of productive efficiency in the top-producing countries. On the other hand, higher cost convexity is a disadvantage that reflects decreasing returns to scale given the already high levels of production in these countries. Our estimates of cost convexity κ^{km} are larger with IV than with OLS, as is consistent with the downward bias from cost shocks and expectational errors that we describe above.

The bottom row shows the implications of these cost estimates for our supply elasticities. Heterogeneity across and within countries highlights the value of building elasticities up from mine-level data. Top producers generally operate at lower utilization rates in the data, such that their cost advantage prevails. Spare capacity in the DRC and Australia allows them to respond quickly to market conditions for cobalt and lithium. Top producers also benefit from higher ore grades, as shown in Appendix Figure A4. For all top producers, high ore grades lead to large revenue increases when prices rise, amplifying their responsiveness to world prices. Higher cost convexity has the opposite effect, but lower utilization rates and higher ore grades dominate on net.

[12] Our dynamic setting introduces expectational error η^{it} and the need for lagged instruments. We highlight two further advantages of cost data in the dynamic setting. First, the cost moment allows us to separately identify cost shocks ε^{it} and expectational error η^{it}, as average cost equation 18 excludes the latter. We do not rely solely on Euler equation 16, which can only recover composite error $(\Delta\varepsilon^{it} + \eta^{it})$. Second, the cost moment incorporates cross-sectional variation in estimation. We do not rely solely on intertemporal Euler variation, which may be subject to noise in the exact timing of production decisions.

Figure 7: Supply estimates

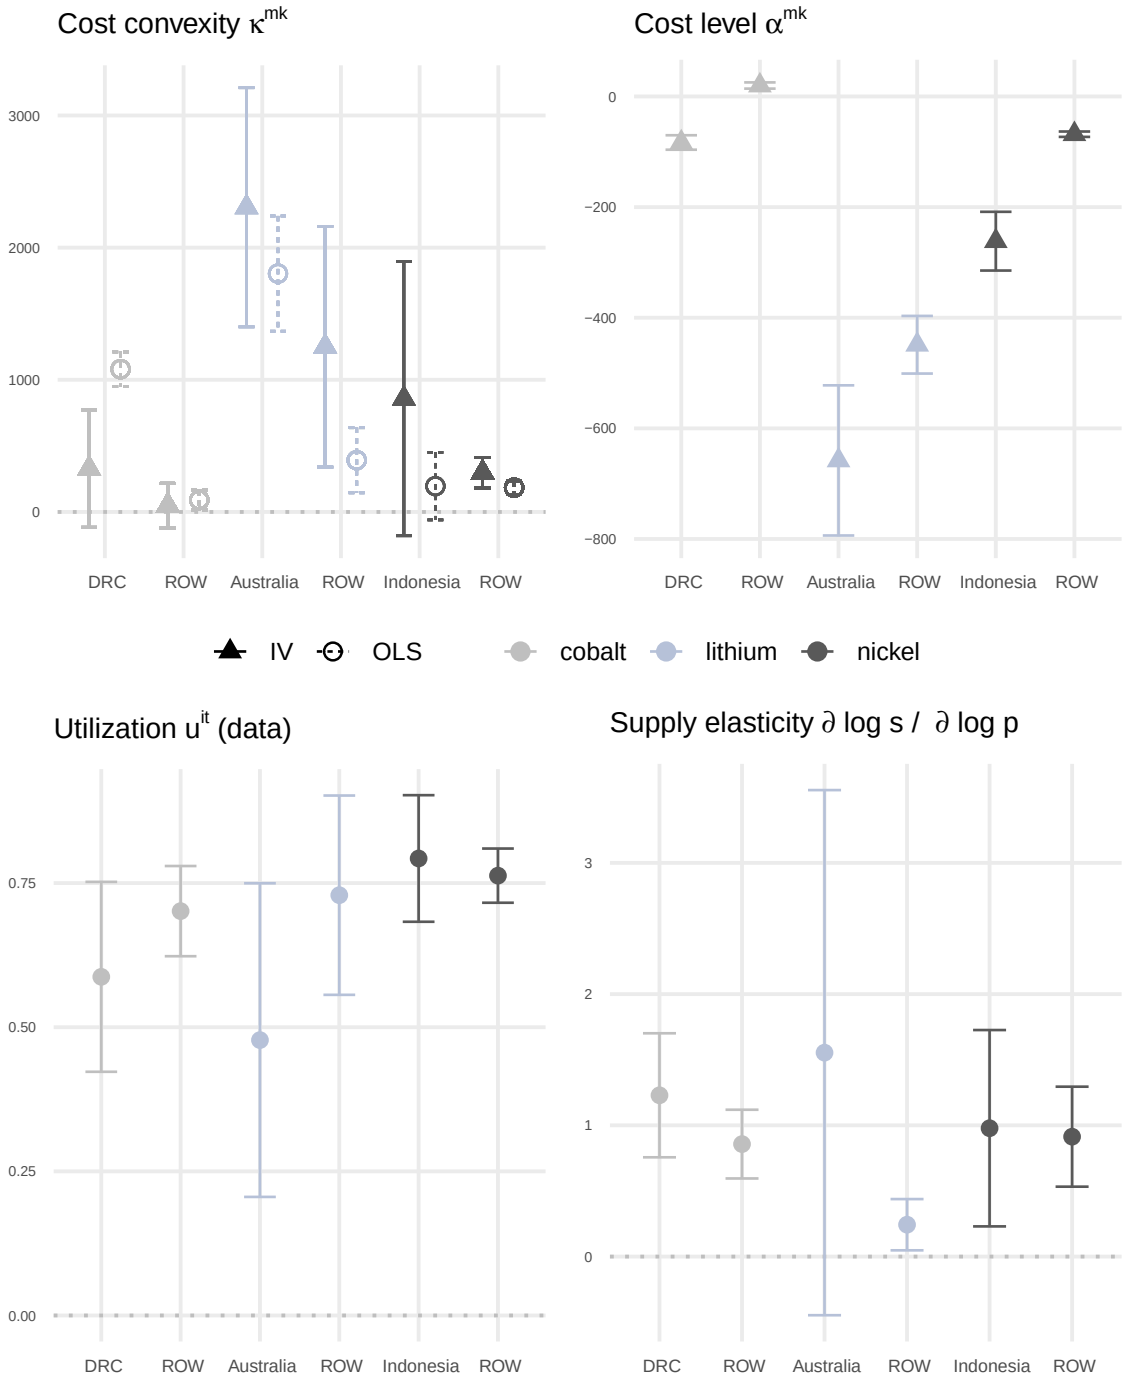

In the top row, we show point estimates and standard errors for cost convexity κ^{km} and cost level α^{km}. For each mineral, we separate top-producing countries and the rest of the world. In the bottom row, we construct utilization rates from observed utilization across mines and time, and we compute average supply elasticities as $\frac{\partial \log s^{it}}{\partial \log p^{mt}} = \frac{g^i p^{mt}}{\kappa^{km} u^{it}}$ for a transitory price change. Appendix C.1 derives this expression.

More broadly, these quantitative responses of top producers will shape the impacts of their mineral policy, which we explore in our counterfactual analysis.

6 Equilibrium

We close the model with a simple midstream layer that connects upstream mineral prices to downstream battery prices. We define an equilibrium in which mineral prices clear global mineral markets, and we describe how we solve the model for counterfactuals.

6.1 Supply chain

Recall the three broad stages of the supply chain: upstream mining and processing, midstream refining and battery manufacturing, and downstream EV manufacturing. Section 5 characterizes upstream mineral supply s^{mt} as a function of world mineral prices p^{mt}. Upstream mines produce processed minerals, which they sell to midstream refiners. Refiners produce refined minerals, which they sell to midstream battery manufacturers. Manufacturers produce batteries, which they sell to downstream EV manufacturers. Section 4 characterizes downstream battery demand d_{jt} as a function of world battery prices p_{jt}.

We now specify the midstream layer. Battery prices include mineral inputs and a midstream wedge.

$$p_{jt} = \sum_{m \in \mathcal{M}} r_j^m p^{mt} + \mu_{jt}, \tag{19}$$

where μ_{jt} captures mineral refining and battery manufacturing costs, as well as potential markups. We recover μ_{jt} from the data as a function of observed prices (p_{jt}, p^{mt}) and recipes r_j^m. For counterfactuals, changes in mineral prices pass through to battery prices, but we hold μ_{jt} fixed. Fixing μ_{jt} may bias our results in two ways. First, midstream costs may change. We simulate upstream supply restrictions, which reduce midstream quantities. If midstream costs are convex, then lower midstream quantities correspond to lower μ_{jt} and lower battery prices. Fixing μ_{jt} leads us to overstate counterfactual battery prices. Second, markups may change. We recover these markups as part of our midstream wedge, but we hold them fixed in counterfactuals. Although midstream production is concentrated among large Chinese firms,

we argue that markups are plausibly limited. For mineral refining, the dominant refiners are vertically integrated into battery manufacturing – either directly or through joint ventures – which limits their incentive to exercise market power on their own downstream operations.[13] For battery manufacturing, existing evidence suggests that pricing is largely competitive: Barwick et al. (2025) explicitly model market power in battery manufacturing and estimate limited markups of 11%.[14]

6.2 Market clearing

The equilibrium market clearing condition equates mineral supply and demand for each mineral m in each period t.

$$s^{mt}(p^{mt}, p^{mt+1}, \ldots, p^{mT}) = d^{mt}(\boldsymbol{p}^t) \quad \forall\ m, t, \tag{20}$$

where $\boldsymbol{p}^t = \{p^{mt}\}^{m \in \mathcal{M}}$ collects mineral prices. Mineral supply depends on prices over time because mines are forward-looking and have finite reserves. Mineral demand depends on prices across minerals because minerals are used jointly in battery technologies and because battery consumers can switch across technologies. Mineral demand $d^{mt}(\boldsymbol{p}^t)$ aggregates over battery demand $d_{jt}(\boldsymbol{p}_t)$, weighting by battery recipes r_j^m and allowing for non-EV demand ν^{mt}.

$$d^{mt}(\boldsymbol{p}^t) = \sum_{j \in \mathcal{J}} r_j^m d_{jt}(\boldsymbol{p}_t) + \nu^{mt},$$

where $\boldsymbol{p}_t = \{p_{jt}\}_{j \in \mathcal{J}}$ collects battery prices. Mineral supply s^{mt} is defined by equation 15, and battery demand d_{jt} by equation 6. In equilibrium, mineral prices $\boldsymbol{p}^t$ pin down battery prices $\boldsymbol{p}_t$ by equation 19.

[13] The top lithium refiner in China, Ganfeng Lithium, is vertically integrated through its subsidiary Ganfeng LiEnergy, which manufactures battery cells and energy storage systems. For cobalt, the leading midstream firms are Zhejiang Huayou Cobalt, GEM Co., and CNGR Advanced Material. They supply cathode materials directly to major battery manufacturers through joint ventures with LG Chem and POSCO. For nickel, Tsingshan Holding Group is the single largest producer at 25 to 30% of the global nickel market, and it owns the battery manufacturer REPT Battero.

[14] Barwick et al. (2025) report battery markups of $142 per kWh relative to baseline total production costs of $1,286 per kWh. The largest battery manufacturer reports average markups of $83 per kWh.

6.3 Solving the model

We solve for the path of mineral prices $\{\boldsymbol{p}^t\}_{t=2024}^{T}$ that satisfies equilibrium condition 20 in each year until final year T. We set $T = 2035$ in our reported counterfactuals, and we explore the robustness of our results to this choice in Appendix D.6. On the supply side, we solve by backward induction from T: each mine chooses its extraction path taking the full price path as given, subject to the reserves constraint. We impose no terminal scrap value for unextracted reserves after 2035. On the demand side, we construct regional battery expenditures X_{kt} from 2025 to 2035 using BMI (2025) forecasts. These forecasts allow us to incorporate current expectations over the evolution of battery technologies. For midstream costs μ_{jt}, we recover their 2024 values with equation 19. For non-EV mineral demand, we impose a linear demand function with a price elasticity of -10%.[15] We recover the intercept from 2024 data. We set demand and supply shocks $(\varepsilon_{jkt}, \varepsilon^{it}, \eta^{it})$ to zero.[16] Appendix D provides additional detail.

Solving over a finite horizon is consistent with our treatment of battery technologies, mine entry, and reserve discovery, which evolve exogenously in counterfactuals. Over the longer run, battery research may lead to technological innovation that reduces the demand for our minerals of interest beyond what is already captured by our forecast data. Similarly, the entry of new mines and discovery of new reserves may reduce the price pressures from exhausting existing reserves. At the same time, we note that these adjustments are likely to be slow in nature. For example, sodium-ion batteries remain in the early stages of EV deployment after decades of research, and the lead time to establish mine production is more than 10 years (Appendix A.3).

7 Counterfactuals

We use our model to perform three sets of counterfactual exercises. First, we quantify supply chain vulnerability by removing the largest producer of each mineral

[15] With finite reserves and a sufficiently long time horizon, elasticity in non-EV demand is required for equilibrium existence. If non-EV demand is perfectly inelastic, then supply is unable to satisfy demand at any price when reserves are eventually exhausted.

[16] Doing so yields a deterministic environment, such that mines face no uncertainty about the future evolution of demand or supply.

Table 2: Supply chain vulnerability

		Price effects (%)				
	Share (%)	Nickel	Lithium	Cobalt	Battery N	Battery L
Nickel (Indonesia)	34.5	89.1	−3.9	−6.2	5.1	−0.8
Lithium (Australia)	33.3	−2.8	95.7	−7.8	18.9	20.3
Cobalt (DRC)	83.8	−3.9	−6.7	276.7	9.3	−1.4

For each mineral, we remove the top producer and compute the resulting average mineral and battery price changes over our time horizon. Each is relative to our baseline equilibrium. We report the top producer's baseline share of global production. Battery L is LFP, and battery N is NMC and NCA.

and evaluating impacts on mineral and battery markets. Second, we simulate unilateral Indonesian nickel policy, and we assess its implications for geopolitics and green adoption. Third, we discuss the potential impacts of mineral cartels.

7.1 Supply chain vulnerability

Concentrated mineral production makes global supply chains vulnerable to natural disaster, political instability, and other forms of disruption. We evaluate this vulnerability with an extreme scenario: eliminating supply from each mineral's top producer. Indonesia accounts for 35% of global nickel supply, Australia for 33% of lithium, and the DRC for 84% of cobalt. Table 2 presents the equilibrium price effects of removing these countries from production. Nickel and lithium prices nearly double, while cobalt prices nearly quadruple. The effect for cobalt is especially pronounced, reflecting both the high baseline share of DRC production and the large ore-grade advantage of DRC mines (Figure 2). When the DRC exits, other mines have low ore grades and high costs of expanding production, and so prices rise sharply.

The cross-mineral price effects reveal strong complementarity. Nickel prices rise when we remove Indonesian nickel, but lithium and cobalt prices fall. As established in Section 3, this pattern is the empirical signal that lithium, cobalt, and nickel are gross complements. If these minerals were gross substitutes, then lithium and cobalt prices would rise – not fall – as rising nickel prices push consumers toward lithium and cobalt. If these minerals were perfect substitutes, then lithium and cobalt prices would rise as much as nickel prices rise. Measured against this benchmark, the declines we observe are modest in absolute terms but represent a large deviation in the opposite

direction. Lithium and cobalt prices fall as nickel prices rise, driven by the joint use of these minerals in NMC and NCA batteries. Analogous patterns hold for Australian lithium and DRC cobalt.

Battery price effects reflect battery recipes. When we remove Indonesian nickel or DRC cobalt, prices rise for battery N (NMC and NCA) and fall for battery L (LFP). Battery price effects are more muted than the rise in nickel or cobalt prices because consumers can substitute toward battery L, which uses neither nickel nor cobalt. Elastic demand dampens the price spike. Batteries also include non-mineral inputs that we hold fixed, further dampening the impact of mineral price shocks. When we remove Australian lithium, prices rise with larger magnitudes because lithium is a universal input. Higher lithium prices raise all battery prices, and so consumers cannot substitute away from lithium. Inelastic demand exacerbates the price spike.

7.2 Unilateral policy

We consider a scenario in which an individual country restricts its mineral supply to exercise market power, motivated by recent export restrictions imposed by Indonesia and the DRC on their nickel and cobalt producers.[17] Our main analysis simulates unilateral Indonesian nickel policy, which we model as a production tax set to maximize Indonesian welfare. We compute welfare as the sum of producer surplus and government revenue.[18] This analysis of restrictive nickel policy is the empirical analog of the theory in Section 3.

Figure 8 shows the geopolitical effects of unilateral Indonesian policy. In panels (a) and (b), nickel prices rise and quantities fall as Indonesian policy restricts world supply. For lithium and cobalt, prices and quantities both fall. The magnitudes are smaller than those in the previous section because Indonesian policy restricts production but does not eliminate it. Nonetheless, we still observe the empirical signal of complementarity: lithium and cobalt prices fall, rather than rise. In panels (c) to (e), we report welfare effects for the top producers of each mineral. Nickel producers all gain: interests are aligned among substitute minerals. Indonesia gains from its own policy, exercising market power by restricting nickel supply to raise

[17] The DRC government has explicitly stated its policy goal of raising cobalt prices (Cook 2025).
[18] With price-taking mines, the welfare-maximizing production tax implements the same allocation as a centralized Indonesian planner exercising market power.

Figure 8: Geopolitics

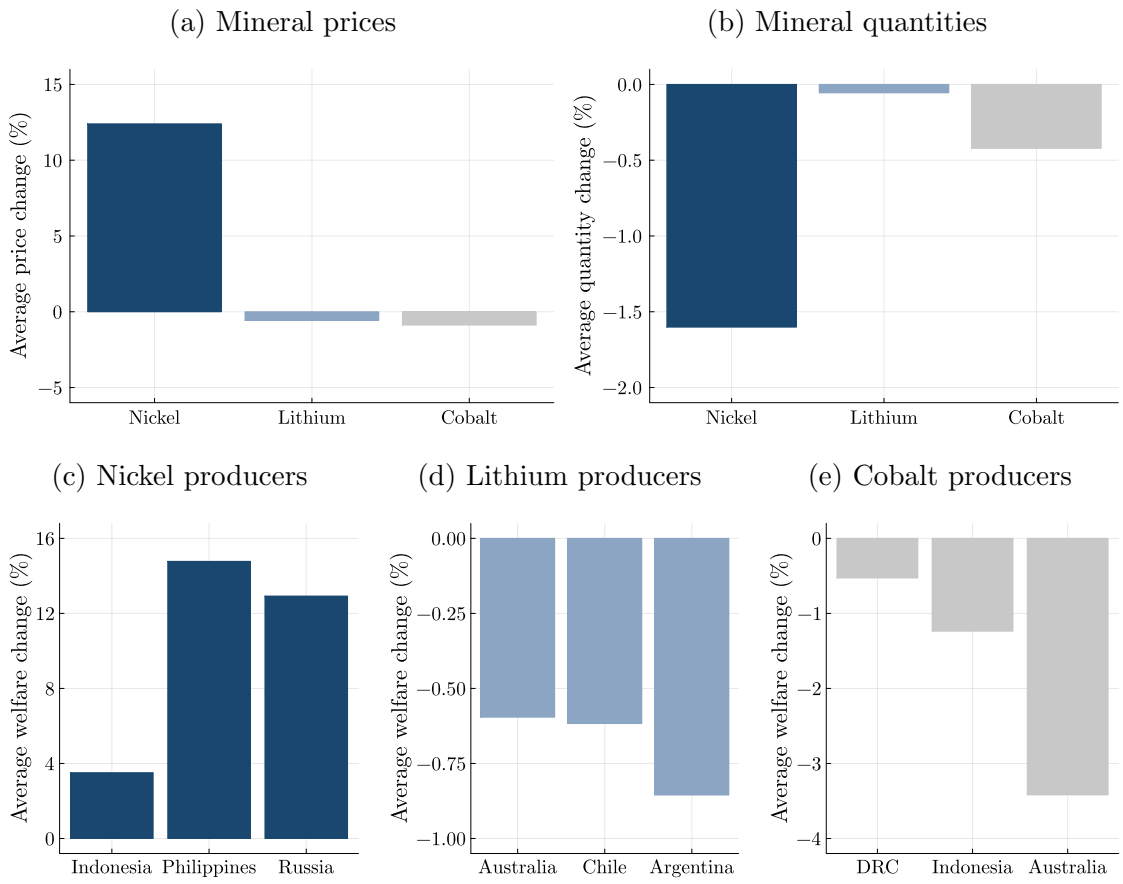

We simulate restrictive nickel policy by Indonesia in the form of a production tax set to maximize Indonesian welfare. Welfare is the discounted sum of producer surplus and government revenue over our time horizon. We report mineral price and quantity effects, as well as welfare effects for the top three producers of each mineral. Each is relative to our baseline equilibrium.

nickel prices. The Philippines and Russia gain from higher nickel prices, free-riding on Indonesian supply restrictions without the need to restrict their own supply. In percentage terms, they gain even more than Indonesia. By contrast, lithium and cobalt producers all lose: interests are in conflict across complementary minerals. In percentage terms, Australian cobalt losses nearly match Indonesian nickel gains.

Figure 9 turns to implications for green adoption. In panels (a) and (b), mineral price effects pass through to battery prices. Nickel-intensive battery N (NMC and NCA) becomes more expensive and adoption falls, while lithium-intensive battery L (LFP) becomes cheaper and adoption rises. On net, restrictive nickel policy reduces total green adoption. In panels (c) and (d), we decompose the contributions of substi-

Figure 9: Green adoption

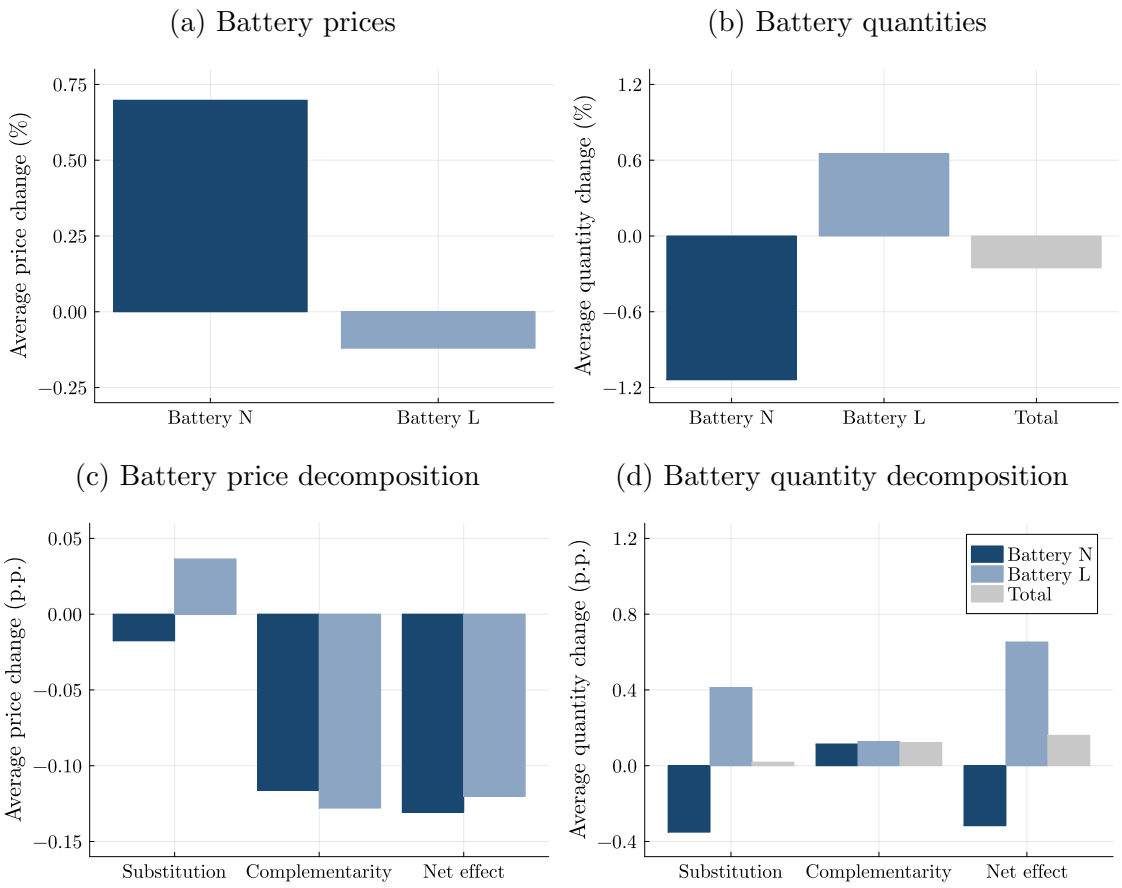

We simulate restrictive nickel policy by Indonesia in the form of a production tax set to maximize Indonesian welfare. Welfare is the discounted sum of producer surplus and government revenue over our time horizon. We report battery price and quantity effects, where battery L is LFP and battery N is NMC and NCA. Each is relative to our baseline equilibrium. We then decompose these price and quantity effects into the marginal contributions of substitution and complementarity. Each is relative to a scenario in which neither force operates, but the same production tax applies.

tution and complementarity. We do so by simulating a benchmark scenario in which the same production tax applies but neither force operates: no consumer switching across battery technologies and no joint mineral use. We shut down consumer switching by freezing battery expenditure shares at their baseline values, and we shut down joint mineral use by eliminating lithium from the battery N recipe.[19] Relative to this benchmark scenario, allowing for consumer switching isolates the marginal effect of substitution. Substitution leads to modest reallocation from battery N to battery L

[19] Appendix D.4 provides implementation details.

as consumers shift in response to relative price changes. The effect on total adoption is muted because increased L adoption is offset by decreased N adoption.

We emphasize the key role of complementarity. Relative to the benchmark scenario, restoring joint mineral use isolates the marginal effect of complementarity. Complementarity reduces price and raises adoption for both batteries. The reason is that complementarity places downward pressure on lithium and cobalt prices when Indonesia restricts its nickel supply. Lower lithium prices subsidize battery L, which uses only lithium. Lower lithium and cobalt prices subsidize battery N, which uses both alongside nickel. Battery prices fall, and quantities rise. Finally, allowing consumer switching and restoring joint use, we obtain the net effect of substitution and complementarity. Complementarity drives the battery price effects and the total adoption effect. We have established above that lithium, cobalt, and nickel are gross complements, and indeed complementarity dominates.

7.3 Cartel policy

We turn to policies set by countries that form a cartel. Single-mineral cartels coordinate among substitute producers within a given mineral, echoing the standard collusive concerns from OPEC and oil markets. Multi-mineral cartels coordinate among complementary producers across minerals, which is a novel possibility in our setting. Although coordination across countries and product markets is difficult in principle, common ownership by multinational firms – which hold stakes in lithium, nickel, and cobalt mines globally – creates scope for such coordination in practice.

We simulate single-mineral cartels that include the top three producers of each mineral. Cartel members impose a joint tax on their production to restrict supply and maximize cartel surplus, which we compute as producer surplus and tax revenue across cartel members. Figure 10 shows the effects of nickel cartel policy by Indonesia, the Philippines, and Russia, relative to unilateral nickel policy by Indonesia. Nickel prices rise by a further 37%. Indonesia enjoys large gains because coordinated cartel policy is more effective than unilateral Indonesian policy, which is subject to leakage. Because cartel policy operates at larger scale, it amplifies losses for lithium and cobalt producers. It also exacerbates declines in green adoption (Appendix Figure D3). At the same time, deviation incentives arise because the Philippines is worse off under nickel cartel policy than under unilateral Indonesian policy, having lost the ability to

Figure 10: Nickel cartel

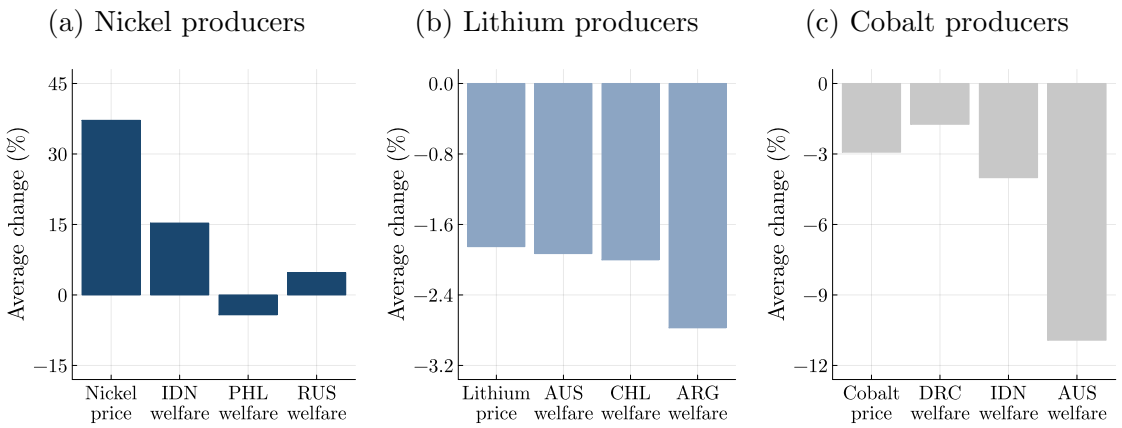

We simulate restrictive nickel policy by Indonesia in the form of a production tax set to maximize Indonesian welfare. Welfare is the discounted sum of producer surplus and government revenue over our time horizon. We then simulate restrictive cartel policy by the top three nickel producers – Indonesia, the Philippines, and Russia – in the form of a production tax set to maximize cartel welfare. We report the differences between the cartel and unilateral scenarios in world mineral prices and in welfare for the top three producers of each mineral.

free-ride on Indonesian restrictions. Cartel cohesion may therefore require internal transfers. Appendix Figures D1 and D2 show similar tensions for lithium and cobalt. For lithium, Argentina gains by exiting the coalition with Australia and Chile and free-riding on their supply restrictions. For cobalt, Indonesia and Australia gain substantially by exiting the coalition with the DRC.

We show that complementarity shapes the impacts of multi-mineral cartels. Consider a nickel cartel that includes Indonesia, the Philippines, and Russia and a lithium cartel that includes Australia, Chile, and Argentina. We first suppose that each cartel independently maximizes its own-cartel welfare. Their minerals are complements in battery demand, and so restrictive policy by one cartel destroys demand for the other cartel. Neither cartel internalizes these cross-mineral spillovers. We then suppose that the two cartels merge into a single nickel-lithium cartel that maximizes joint welfare.[20] The merged cartel fully internalizes the cross-mineral spillovers.

Figure 11 decomposes the differences between these two scenarios. We isolate the marginal contributions of substitution and complementarity by simulating bench-

[20] The gains from joint welfare maximization may not be distributed evenly. Appendix Figure D4 shows that nickel producers all have considerable gains, while lithium producers all have modest losses. The merged cartel may require nickel-to-lithium transfers to sustain cohesion.

Figure 11: Multi-mineral cartel

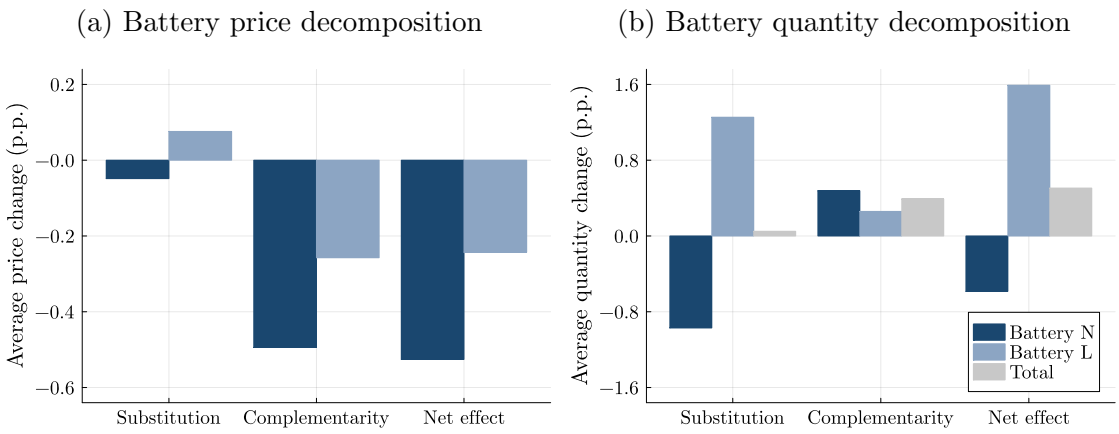

We simulate a two-cartel scenario in which a nickel cartel and a lithium cartel each set a production tax to maximize their individual welfare. Welfare is the discounted sum of producer surplus and government revenue over our time horizon. We then simulate a one-cartel scenario in which the two cartels merge and jointly set production taxes on nickel and lithium to maximize their total welfare. We compute the price and quantity differences between the two-cartel and one-cartel scenarios. We then decompose these effects into the marginal contributions of substitution and complementarity. Each is relative to a scenario in which neither force operates, but the same merger scenarios apply.

mark scenarios in which these cartels operate but without consumer switching across batteries or joint use across minerals. The isolated complementarity force shows that the merged cartel sets less aggressive joint policy than the two separate cartels. The merged cartel internalizes the complementarity, and both battery prices fall. The price decline is larger for battery N (NMC and NCA) than for battery L (LFP) because it is battery N that generates the complementarity, and so it is battery N that benefits most from a cartel that internalizes this complementarity. Lower battery prices then result in higher battery adoption. This result reflects the classic logic of Cournot complementarity (Cournot 1838), and it is akin to double marginalization in vertical mergers. When complementary suppliers merge, they remove double markups, lower prices, and expand quantities. Complementarity is pro-adoption.

8 Conclusion

We argue that the structure of critical mineral markets fundamentally shapes the pace and pattern of the green transition. Critical minerals are highly concentrated in a small number of countries, and these countries exercise market power by restricting

supply. We show that cross-mineral complementarity determines how mineral concentration affects battery adoption, both in sign and in magnitude. For unilateral nickel policy, complementarity causes lithium prices to fall as nickel supply contracts, subsidizing lithium-intensive batteries. For coordinated cartel policy, complementarity encourages a nickel-lithium cartel to reduce – rather than raise – both nickel and lithium prices, again supporting green adoption. Our climate progress will depend on these critical markets.

References

Abuin, Constanza. Power decarbonization in a global energy market: The climate effect of US LNG exports. 2025.

Administration, U.S. Energy Information. Petroleum and other liquids. 2025. URL https://www.eia.gov/international/data/world/petroleum-and-other-liquids/annual-petroleum-and-other-liquids-production.

Aguirregabiria, Victor and Andres Luengo. A microeconometric dynamic structural model of copper mining decisions. 2016.

Alfaro, Laura, Harald Fadinger, Jan Schymik, and Gede Virananda. Trade and industrial policy in supply chains: Directed technological change in rare earths. 2025.

Allcott, Hunt, Reigner Kane, Maximilian Maydanchik, Joseph Shapiro, and Felix Tintelnot. The effects of "Buy American": Electric vehicles and the Inflation Reduction Act. 2026.

Asker, John, Allan Collard-Wexler, and Jan De Loecker. (Mis)allocation, market power, and global oil extraction. *American Economic Review*, 109(4):1568–1615, 2019.

Asker, John, Allan Collard-Wexler, Charlotte De Canniere, Jan De Loecker, and Christopher Knittel. Two wrongs can sometimes make a right: The environmental benefits of market power in oil. 2024.

Bachmann, Ruediger, David Baqaee, Christian Bayer, Moritz Kuhn, Andreas Löschel, Benjamin Moll, Andreas Peichl, Karen Pittel, and Moritz Schularick. What if? The macroeconomic and distributional effects for Germany of a stop of energy imports

from Russia. *Economica*, 91(364):1157–1200, 2024.

Barwick, Panle Jia, Hyuk-Soo Kwon, Shanjun Li, and Nahim Zahur. Drive down the cost: Learning by doing and government policies in the global EV battery industry. 2025.

Barwick, Panle Jia, Hyuk-Soo Kwon, Shanjun Li, Yucheng Wang, and Nahim Zahur. Industrial policies and innovation: Evidence from the global automobile industry. *International Journal of Industrial Organization*, 104:103230, 2026.

Benchmark Mineral Intelligence. Supply chain data. 2025. URL https://www.benchmarkminerals.com.

Berry, Steven. Estimating discrete-choice models of product differentiation. *RAND Journal of Economics*, 25(2):242–262, 1994.

Berry, Steven, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

Blackwill, Robert and Jennifer Harris. *War by other means: Geoeconomics and statecraft*. Harvard University Press, 2016.

Blundell, Richard and Jean Marc Robin. Estimation in large and disaggregated demand systems: An estimator for conditionally linear systems. *Journal of Applied Econometrics*, 14(3):209–232, 1999.

Bonakdarpour, Mohsen, Frank Hoffman, and Keerti Rajan. Mine development times: The US in perspective. Technical report, S&P Global, 2024.

Bornstein, Gideon, Per Krusell, and Sergio Rebelo. A world equilibrium model of the oil market. *Review of Economic Studies*, 90(1):132–164, 2023.

Clausing, Kimberly, Jonathan Colmer, Allan Hsiao, and Catherine Wolfram. The global effects of carbon border adjustment mechanisms. 2025.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger. A framework for geoeconomics. *Econometrica*, 94(1):105–136, 2026a.

Clayton, Christopher, Matteo Maggiori, and Jesse Schreger. A theory of economic coercion and fragmentation. 2026b.

Cook, Alexander. Cobalt export quotas: DRC sets limits to rebalance global supply, 2025. URL https://www.fastmarkets.com/insights/drc-cobalt-export-

quotas-2025/.

Copeland, Brian and M. Scott Taylor. *Trade and the environment: Theory and evidence*. Princeton University Press, 2003.

Copeland, Brian, Joseph Shapiro, and M. Scott Taylor. Globalization and the environment. In Gopinath, Gita, Elhanan Helpman, and Kenneth Rogoff, editors, *Handbook of International Economics*, volume 5, pages 61–146. 2022.

Cournot, Antoine Augustin. *Recherches sur les principes mathématiques de la théorie des richesses*. L. Hachette, Paris, 1838.

De Cannière, Charlotte. Pump it? Market power and the energy transition in the global oil market. 2025.

Deaton, Angus and John Muellbauer. An almost ideal demand system. *American Economic Review*, 70(3):312–326, 1980.

Farrokhi, Farid. Global sourcing in oil markets. *Journal of International Economics*, 125:103323, 2020.

GlobalData. Mining data and insights. 2025. URL https://www.globaldata.com.

Hall, Robert. Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy*, 86(6):971–987, 1978.

Hassler, John, Per Krusell, and Conny Olovsson. Directed technical change as a response to natural resource scarcity. *Journal of Political Economy*, 129(11):3039–3072, 2021.

Head, Keith, Thierry Mayer, Marc Melitz, and Chenying Yang. Industrial policies for multi-stage production: The battle for battery-powered vehicles. 2026.

Hirschman, Albert. *National Power and the Structure of Foreign Trade*. University of California Press, 1945.

Hotelling, Harold. The economics of exhaustible resources. *Journal of Political Economy*, 39(2):137–175, 1931.

Hsiao, Allan. Coordination and commitment in international climate action: Evidence from palm oil. *Econometrica*, 94(1):1–33, 2026.

International Energy Agency. Natural gas supply. 2023. URL https://www.iea.org/

countries/russia/natural-gas.

International Energy Agency. Global EV outlook 2024. Technical report, 2024.

International Energy Agency. Global critical minerals outlook 2025. Technical report, 2025.

Johnson, Simon, Lukasz Rachel, and Catherine Wolfram. Design and implementation of the price cap on Russian oil exports. *Journal of Comparative Economics*, 51(4): 1244–1252, 2023.

Johnson, Simon, Lukasz Rachel, and Catherine Wolfram. A theory of price caps on non-renewable resources. 2025.

Kellogg, Ryan. The end of oil. 2024.

Knehr, Kevin, Joseph Kubal, Paul Nelson, and Shabbir Ahmed. Battery performance and cost modeling for electric vehicles: A manual for BatPaC v5.0. Technical Report ANL/CSE-22/1, Argonne National Laboratory, Lemont, IL, 2024. URL https://publications.anl.gov/anlpubs/2022/07/176234.pdf.

Kortum, Samuel and David Weisbach. The design of border adjustments for carbon prices. *National Tax Journal*, 70(2):421–446, 2017.

Kortum, Samuel and David Weisbach. Optimal unilateral carbon policy. 2024.

Kwon, Hyuk-Soo. Subsidies versus tradable credits for electric vehicles: The role of market power in the credit market. 2025.

Mohr, Cathrin and Christoph Trebesch. Geoeconomics. *Annual Review of Economics*, 17:563–587, 2025.

Moll, Benjamin, Moritz Schularick, and Georg Zachmann. The power of substitution: The great German gas debate in retrospect. *Brookings Papers on Economic Activity*, 54(2):395–481, 2023.

Natural Resources Canada. The Canadian critical minerals strategy. Technical report, Government of Canada, 2022. URL https://www.canada.ca/content/dam/nrcan-rncan/site/critical-minerals/Critical-minerals-strategyDec09.pdf.

Popp, David. Induced innovation and energy prices. *American Economic Review*, 92 (1):160–180, 2002.

Rauscher, Michael. *International Trade, Factor Movements, and the Environment*. Oxford University Press, 1997.

Remmy, Kevin. Adjustable product attributes, indirect network effects, and subsidy design: The case of electric vehicles. *American Economic Journal: Economic Policy*, 2026.

S&P Global. Metals and mining. 2025. URL https://www.spglobal.com/en/research-insights/market-insights/energy-commodities/metals-mining.

United States Geological Survey. What is a critical mineral? 2026. URL https://www.usgs.gov/faqs/what-a-critical-mineral.

Wachtmeister, Henrik and Mikael Höök. Investment and production dynamics of conventional oil and unconventional tight oil: Implications for oil markets and climate strategies. *Energy and Climate Change*, 1:100010, 2020.

APPENDIX

A Data

A.1 Battery recipes

Battery recipes come from BatPaC, developed by Argonne National Laboratory (Knehr et al. 2024). Their lithium-ion battery cost model includes the mineral intensities of commodities for a large number of EV battery technologies. Figure A1 summarizes the recipes for each technology. While many values are pulled directly from the BatPaC model, a few require additional explanation. Aluminum exists in batteries both because it may be part of the cathode active material (in NCA batteries) and because it is part of a current collector foil in all batteries. We use the sum of aluminum from both of these sources. We additionally convert iron phosphate ($FePO_4$) quantities to quantities of phosphoric acid (H_3PO_4), a precursor to iron phosphate, based on stoichiometric conversions. The chemical reaction to iron phosphate is $Fe_2O_3 + 2\,H_3PO_4 \longrightarrow 2\,FePO_4 + 3\,H_2O$, and since $FePO_4$ is 150.8 g/mol and H_3PO_4 is 98.0 g/mol, one tonne of iron phosphate requires 98.0/150.8 tonnes of phosphoric acid. Additionally, we aggregate NMC varieties to a representative category that we call NMC by averaging over the recipes, weighted by aggregate market shares.

A.2 Harmonization of mine-level datasets

We combine data from three proprietary sources – BMI (2025), GlobalData (2025), and S&P Global (2025) – to create a harmonized mine-level panel for lithium, nickel, and cobalt from 2010 to 2024. Given there are no common identifiers across datasets and mine names are not always standardized, this exercise consists of verifying that each individual mine is correctly matched across data sources. We do so with variables that appear across datasets: names, location, ownership information, production, and capacity volumes. Our final panel consists of 397 mines that produce lithium, cobalt, or nickel.

We perform several validation exercises to gauge the coverage and representativeness of our harmonized panel. We use USGS values for 2020 to 2024 as reference

Figure A1: Mineral intensities (kg/kWh) across battery technologies

	LFP	NCA	NMC-111	NMC-523	NMC-622	NMC-811
phosphorus	0.38	-	-	-	-	-
phosphoric acid	1.43	-	-	-	-	-
nickel	-	0.65	0.32	0.43	0.51	0.61
manganese	-	-	0.3	0.24	0.16	0.07
lithium	0.09	0.1	0.12	0.11	0.1	0.09
iron phosphate	2.2	-	-	-	-	-
iron	0.68	-	-	-	-	-
graphite	0.95	0.85	0.84	0.84	0.85	0.85
fluorine	0.09	0.06	0.07	0.06	0.06	0.06
copper	0.62	0.33	0.39	0.36	0.36	0.32
cobalt	-	0.12	0.32	0.17	0.17	0.08
aluminum	0.24	0.14	0.15	0.14	0.13	0.12

Mineral intensities are obtained from the battery technology data described in section 2.2.

values for world and country-level production. First, our mines cover 92% of world production for lithium, 94% for cobalt, and 80% for nickel. Second, we aggregate our mine-level data up to the country-level and compare production values to USGS (which does not report production below the country-level). Figure A2 shows our country-level aggregates are highly correlated with USGS values. Country-level correlation coefficients are above 0.93 for lithium, nickel, and cobalt.

A.3 Mine development

Time lags from first discovery to first production are measured in decades in the mining sector. Lead times in most countries range between 15 and 25 years (Bonakdarpour et al. 2024). The United States has an average lead time of 29 years,

Figure A2: Validation of mine-level panel against USGS production values

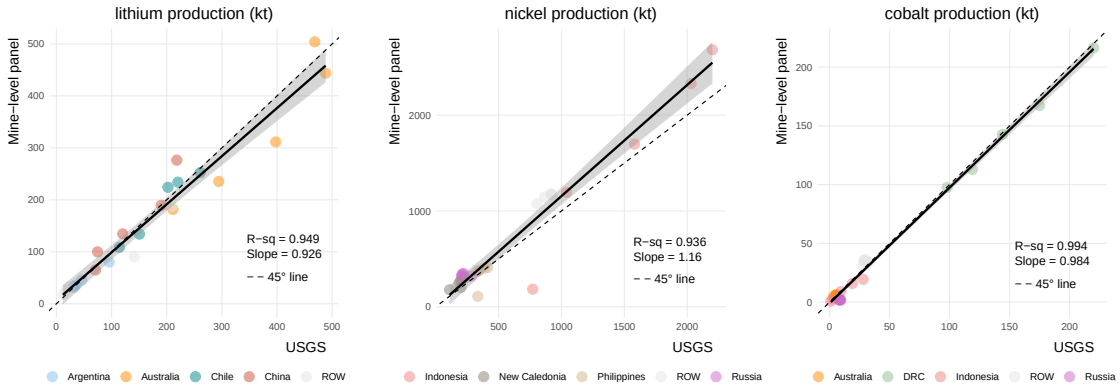

The vertical axes display country-level production values aggregated from our mine-level data. The horizontal axes show country-level production values reported by USGS for 2020 to 2024.

the world's second-longest after Zambia (at 34 years), partly due to uncoordinated and overlapping regulatory jurisdictions on the federal lands where most mining resources are located. By contrast, lead times for conventional oil wells are shorter, on the order of 5 years (Wachtmeister and Höök 2020). There are a variety of reasons why mining projects have notoriously long development times. Mining is inherently more environmentally disruptive than other extractive activities given it requires stripping open entire swathes of land, in contrast to drilling operations that can take place mostly underground or entirely offshore. As a consequence, mining projects are subject to extensive exploration and feasibility studies, lengthy permitting processes, and higher litigation risk from local opposition. Nevertheless, there has been significant growth in mining production over the past decade, especially minerals whose demand primarily stems from the battery sector, such as lithium and cobalt (Figure A3).

A.4 Mine heterogeneity

Ore grade. Ore grade is a major determinant of a mine's profitability and explains why some countries dominate specific mineral markets. Figure A4 shows the distribution of ore grades across all the mines in our data, separately for those located in the top producing country of each mineral and in the rest of the world. While Indonesian nickel grades are higher than in the rest of the world, the gap is not as large as the case of DRC and cobalt. The DRC's significantly higher-grade cobalt deposits are a

Figure A3: World production of mineral commodities (2010-2024)

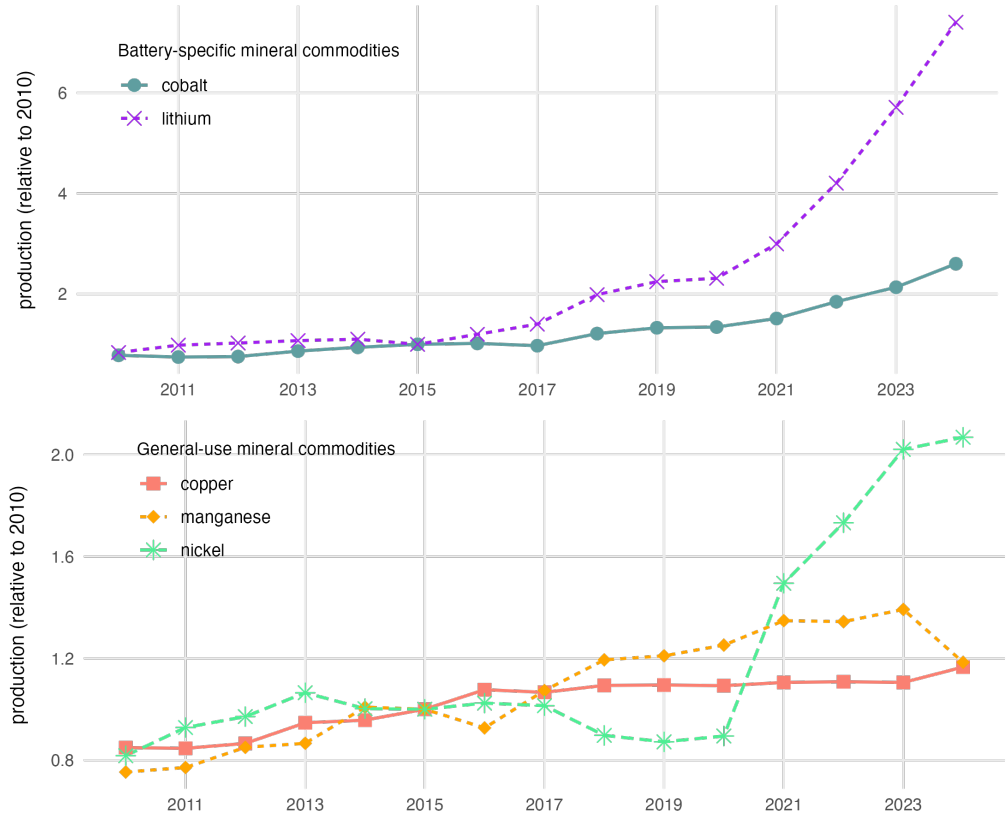

Values are constructed from mine-level production data described in section 2.2. Categorization of mineral commodities into battery-specific versus general-use minerals is based on demand composition values from Table 1.

major reason why it dominates the global cobalt market.

Capacity and utilization rates. Figure A5 shows how mine-level utilization rates change over time, and how they correlate with world prices.

A.5 Recent market trends

A.5.1 Mineral markets

Between 2017-2022 demand from the energy sector was the main factor behind a tripling in global production of lithium and a 70% jump for cobalt. Short-term pressures weakened with the contraction in demand caused by Covid-19. After 2021 production of all minerals grew rapidly as the global economy recovered and many

Figure A4: Distribution of ore grade across mines

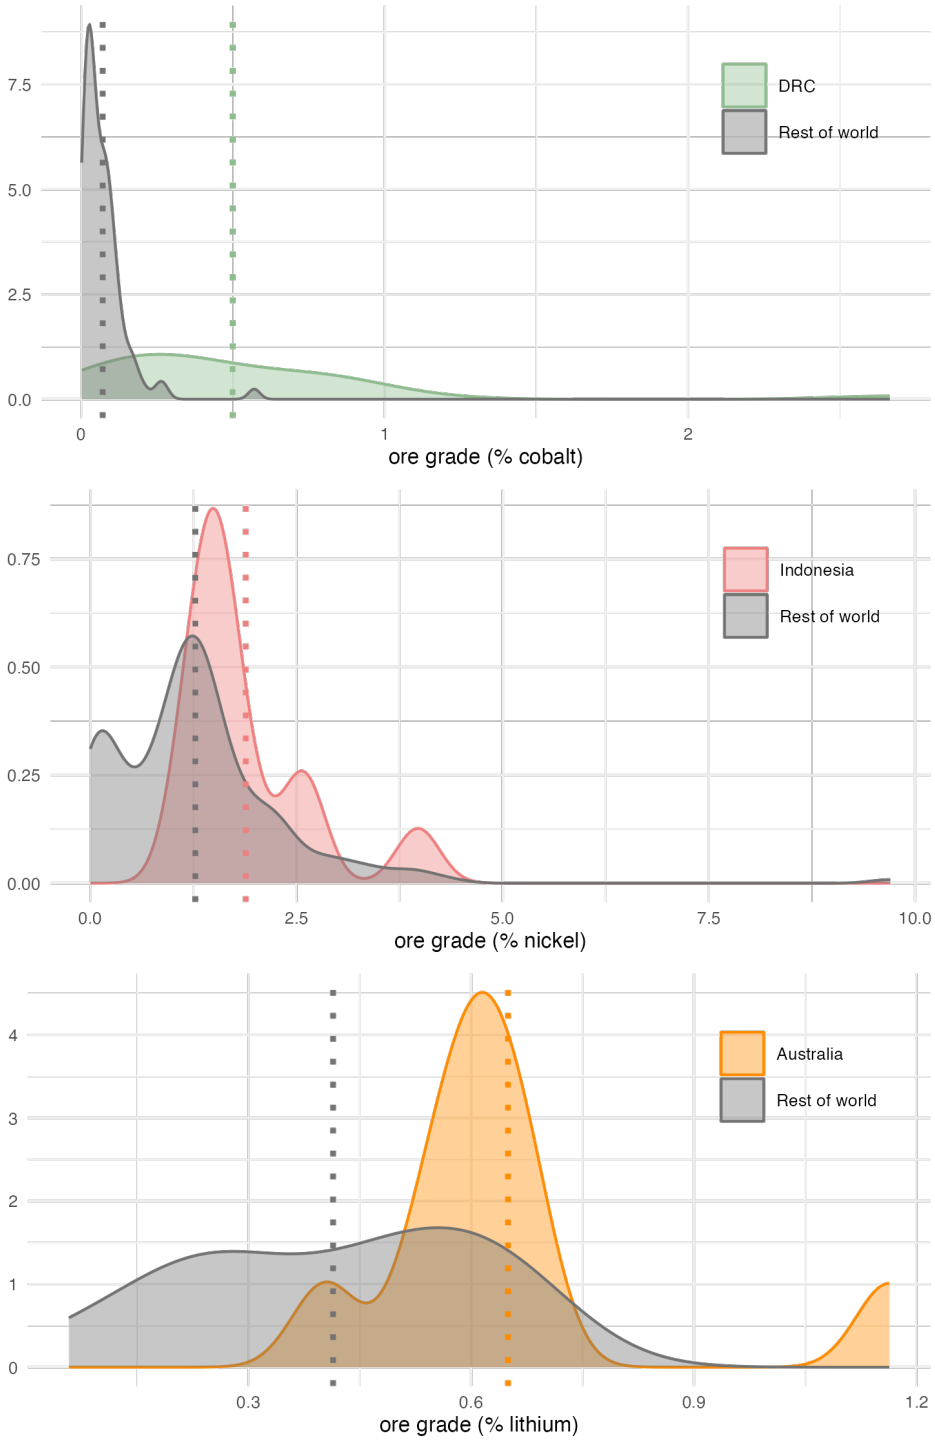

The figure above shows density plots constructed from mine-level ore grade data described in section 2.2. The distribution of mines from the top producer of each mineral is compared to mines in the rest of the world. Vertical lines show the average ore grade within each region.

Figure A5: Mine-level utilization rates and price fluctuations.

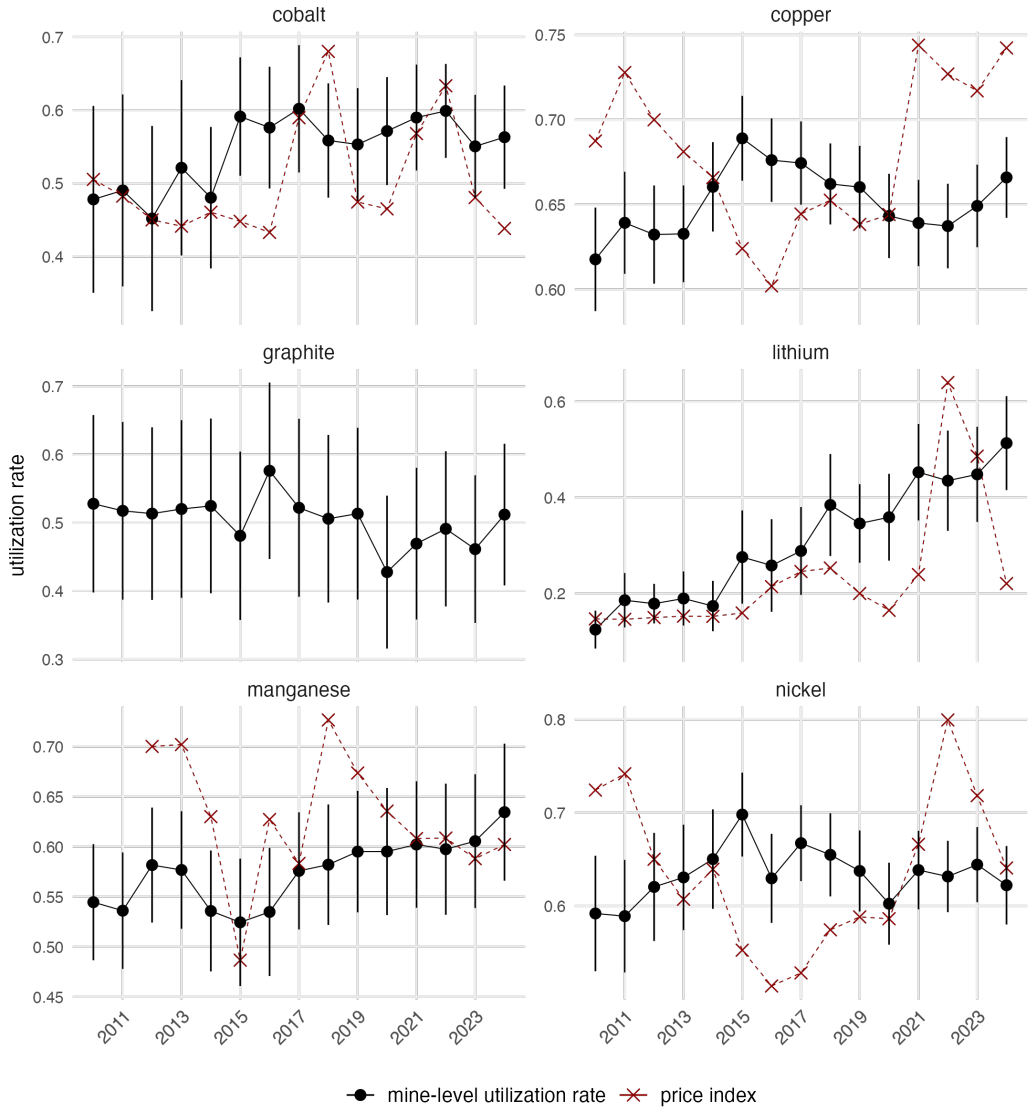

All values are constructed from mine-level production and capacity data described in section 2.2. The distribution of mine-level utilization is displayed in black print as a mean value with a 95% confidence interval. World commodity prices are scaled for visualization, and hence reported as a unit-less index.

countries boosted efforts to accelerate the energy transition. Between 2017-2022, the EV share of mineral demand for lithium, cobalt, and nickel grew from 15% to 60%, 15% to 30%, and 2% to 10% (IEA 2024).

Cobalt market. Between 2015-2022, the EV and grid storage share of cobalt demand

Figure A6: Mineral price trends and cycles (2010-2024)

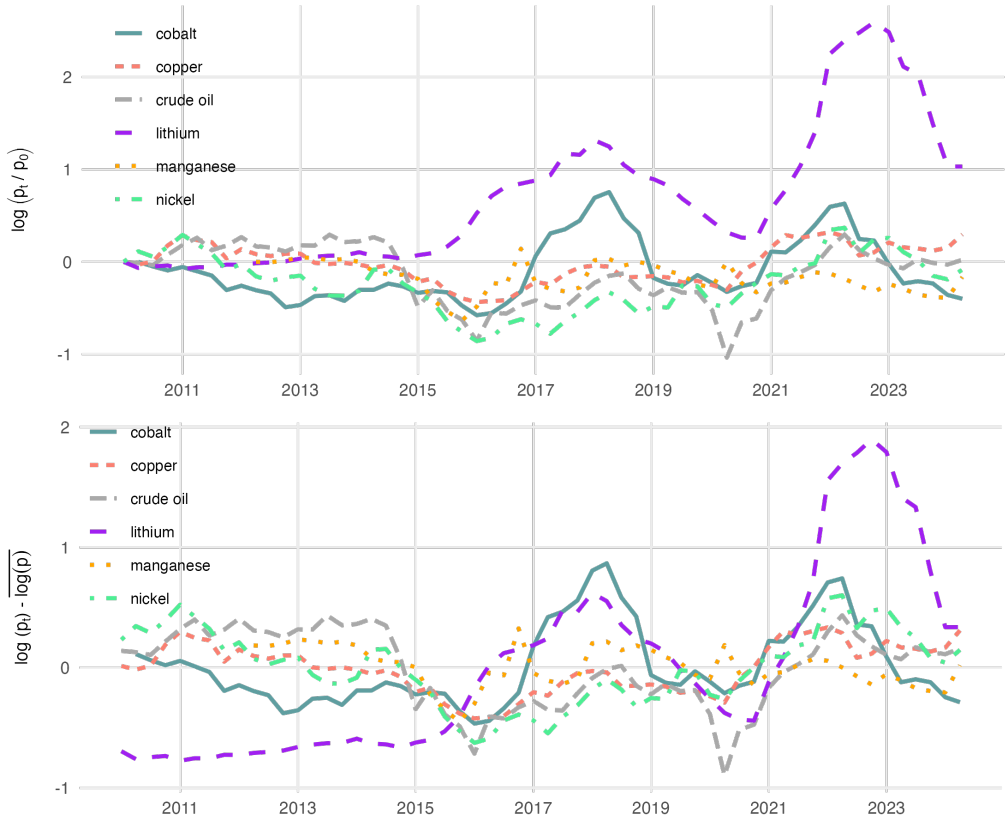

The two panels show the price of each commodity normalized by (i) its price at the start of the sample, and (ii) its average value across the sample period. Nominal values from which this figure is constructed are quarterly-averaged prices from S&P Global (2025). Cobalt, nickel, and copper prices are based on over-the-counter (OTC) and contract for difference (CFD) financial instruments, primarily from the LME. Crude oil prices are WTI futures. Lithium prices are spot prices for lithium carbonate (battery grade, 99.5%) traded in China. Manganese prices are based on Tianjin port prices.

increased from 5 percent to almost 25 percent. Such rapid growth partly explains the rise in cobalt prices between 2016 and early 2018. Disruptions by major suppliers also fueled fear of material security. Manufacturers began shifting to batteries with lower cobalt intensities in response to cobalt price spikes and ethical-sourcing concerns (IEA 2024). Such shifts contributed to the price decline after 2018.

Nickel market. In 2014 Indonesia implemented a ban on unprocessed nickel ore exports to encourage domestic smelting. Prices initially rose but gradually fell in 2015-2016 as alternative production sources and Indonesian domestic smelters took

Figure A7: Mineral, cathode material, and battery prices (2021-2024)

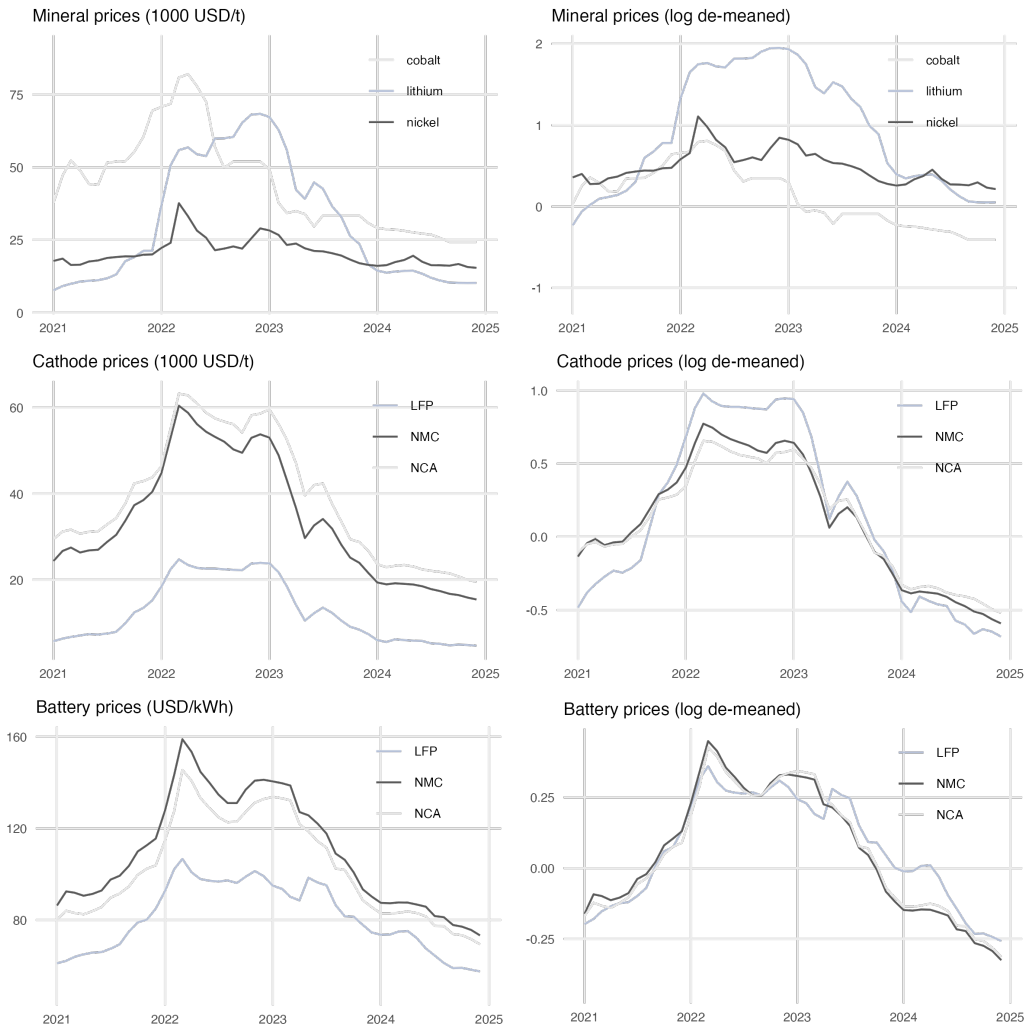

Mineral prices are monthly-averaged prices from S&P Global (2025). Cathode and battery prices are at monthly frequency from BMI (2025).

off. Following Russia's invasion of Ukraine in February 2022, fears of sanctions on Russian nickel led to a price increase. Additionally, in March 2022 a short squeeze on the London Metal Exchange caused nickel prices to surge (Reuters). As a result, in 2022 the price of nickel reached a peak twice as high as its 2015-2020 average. This added incentives to use battery technologies that were less reliant on nickel, such as LFP, despite their lower energy density. Hence, demand substitution contributed to nickel prices dropping from their 2022 peaks.

Figure A8: Mineral prices and battery technology adoption

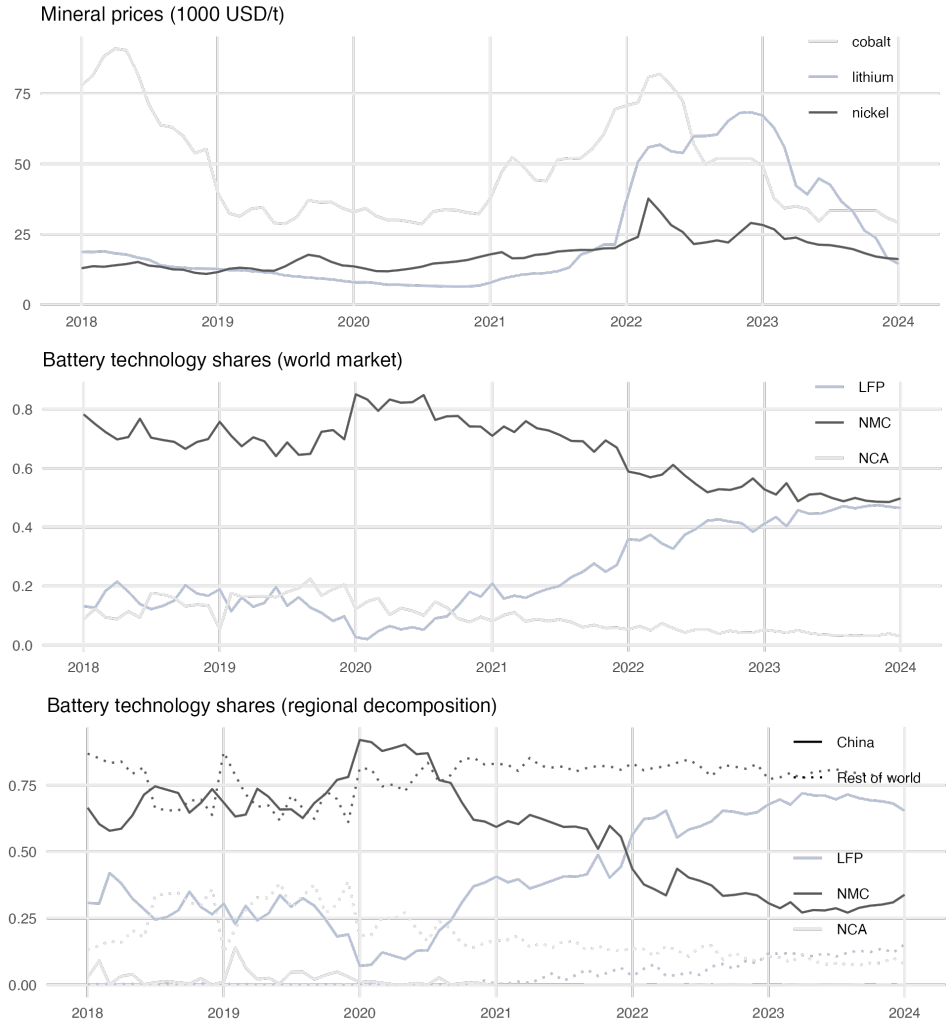

Mineral prices are monthly-averaged prices from S&P Global (2025). Battery technology market shares constructed from units sold, using monthly EV sales data from BMI (2025).

Lithium market. The increase in EV demand drove lithium prices to rise towards 2017. In response to high prices, new lithium mines in Australia, Argentina, and Chile expanded production, alleviating upward price pressure around 2019. Yet by 2022 lithium prices stood six times above their average over the 2015-2020 period, with supply lagging behind the post-2021 demand surge. Supply eventually did respond in Australia, Chile, Argentina, and China, so that by late 2023 lithium prices receded 20 percent, returning to their 2021 levels. Additionally, many automakers shifted to LFP batteries, which require less lithium (IEA 2024).

Figure A9: Regional market shares of EV battery technologies

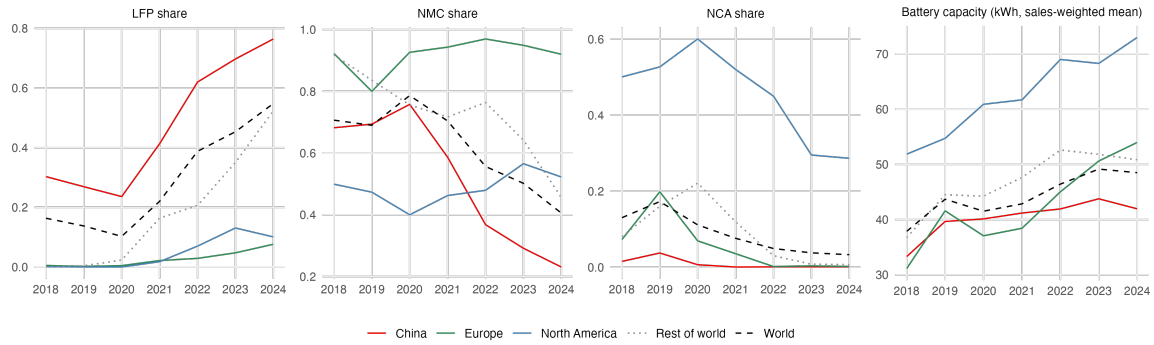

Market shares constructed from units sold, using EV sales data from BMI (2025).

A.5.2 Battery markets

Figure A9 shows the evolution of EV battery technologies worldwide and across four consumer regions: China, Europe, North America and the rest of the world. As of 2024, consumers in China accounted for over 60% of global EV sales, while Europe and North America represented 20% and 10% of the market. Beginning in 2018, every region had positive NMC and NCA shares, while LFP was less than 1% everywhere except China. Since then, LFP has increased its market share everywhere, but especially in China and the rest of Asia, primarily at the expense of NMC. In Europe and North America, NMC remains the main battery technology. NCA has fallen to near zero market share in all regions except North America, where it is still highly demanded for its higher energy density in premium brands. The far right panel of Figure A9 plots the average battery capacity of EV sales in each region.

B Demand

We apply the standard adding-up, homogeneity, and symmetry restrictions in estimating demand. Adding-up imposes $\sum_j w_{jkt} = 1$ for all k, t, and it implies that $\sum_j \alpha_{jt} = 1$, $\sum_j \beta_j = 0$, and $\sum_j \gamma_{jj'} = 0$ for all j'. Homogeneity imposes $\sum_{j'} \gamma_{jj'} = 0$ for all j. Symmetry imposes $\gamma_{jj'} = \gamma_{j'j}$ for all j, j'. The adding-up restriction results in collinearity among the J expenditure share equations, and so we drop one of these equations in estimation. With two batteries L and N, we drop the expenditure share equation for battery N, and estimation proceeds with a single equation for battery L.

We further impose that quality α_{Lt} varies over time only through a common linear time trend, such that $\alpha_{Lt} = \alpha_{L0} + \alpha_{L1}t$. Applying these restrictions, we obtain our estimating equation.

$$w_{Lkt} = \alpha_{L0} + \alpha_{L1}t + \beta_L \log\left(\frac{X_{kt}}{P_t}\right) + \gamma_{LL} \log\left(\frac{p_{Lt}}{p_{Nt}}\right) + \varepsilon_{Lkt}$$

Relative prices $\frac{p_{Lt}}{p_{Nt}}$ appear on the right-hand side because, by homogeneity, we have $\gamma_{LL} = -\gamma_{LN}$. The first stage is

$$\log\left(\frac{p_{Lt}}{p_{Nt}}\right) = \pi_0 + \pi_1 t + \pi_L \log z_{Lt} + \pi_N \log z_{Nt} + \pi_X \log\left(\frac{X_{kt}}{P_t}\right) + \varepsilon_{kt},$$

where instruments z_{jt} are as defined in equation 7.

C Supply

We present key derivations, and we show that our framework accommodates a more general cost structure. We extend the model to capture dynamics from ore degradation, as well as measurement error in our cost data.

C.1 Derivations

Euler equation. Mine i chooses extraction path $\{x^{it}\}^{t\geq 0}$ to maximize the present value of expected profits, subject to the reserves constraint. The Lagrangian is

$$\mathcal{L}^{it} = \sum_{t'=t}^{\infty} \beta^{t'-t}\, \mathbb{E}^{it}\left[mr^{it'}\, x^{it'} - C^{it'} + \lambda^{it'}\left(R^{it'} - x^{it'} - R^{it'+1}\right)\right].$$

The first-order conditions with respect to x^{it} and R^{it+1} are

$$\lambda^{it} = mr^{it} - mc^{it}, \quad \lambda^{it} = \beta \mathbb{E}^{it}\left[\lambda^{it+1}\right].$$

Substitution gives Euler equation 13.

$$mr^{it} - mc^{it} = \beta \mathbb{E}^{it}\left[mr^{it+1} - mc^{it+1}\right]$$

Utilization. Substituting equation 12 into $\lambda^{it} = mr^{it} - mc^{it}$,

$$\lambda^{it} = mr^{it} - \alpha^{km} - \kappa^{km} u^{it} - \varepsilon^{it}.$$

We apply equation 11, and we rearrange to obtain equation 14 for utilization.

$$u^{it} = \frac{2}{\kappa^{km}} \left(mr^{it} - ac^{it} - \lambda^{it} \right)$$

Supply elasticity. Substituting equations 9 and 11, we rewrite utilization as

$$u^{it} = \frac{1}{\kappa^{km}} \left(p^{mt} g^{i} - \alpha^{km} - \varepsilon^{it} - \lambda^{it} \right).$$

Ore grade g^i and cost parameters $(\alpha^{km}, \kappa^{km}, \varepsilon^{it})$ are fundamentals, while λ^{it} represents the shadow value of the reserves constraint. Consider the response of current production s^{it} to a change in current price p^{mt}. Differentiating utilization with respect to price and applying the chain rule to $s^{it} = g^i \bar{x}^i u^{it}$, we obtain the supply elasticity

$$\frac{\partial s^{it}}{\partial p^{mt}} \frac{p^{mt}}{s^{it}} = \frac{p^{mt}}{\kappa^{km} u^{it}} \left(g^i - \frac{\partial \lambda^{it}}{\partial p^{mt}} \right).$$

We have $\frac{\partial \lambda^{it}}{\partial p^{mt}} > 0$ for two reasons. First, for both transitory and permanent price changes, a higher current price encourages current extraction, depleting reserves and raising their shadow value. Second, for permanent price changes, higher future prices increase future revenues and thus the value of conserving reserves. To compute elasticities for Figure 7, we set $\frac{\partial \lambda^{it}}{\partial p^{mt}} = 0$ as if reserves are not binding. This approximation is more accurate for a transitory shock, where only the reserves channel operates. It is less accurate for a permanent shock, where the future-revenue channel further dampens the current response. This approximation does not affect our counterfactuals, which solve the full dynamic supply problem in equation 15.

Myopia and linear costs. We consider the special case with $\beta = 0$ and $\kappa^{km} = 0$. The Euler equation specializes to $mr^{it} = ac^{it}$, which yields a corner solution. Mines produce at capacity if $mr^{it} > ac^{it}$, and they produce zero otherwise.

C.2 Generalized cost structure

We show that our approach accommodates the more general cost structure

$$C^{it} = F(x^{it}; \theta) + H(x^{it}, \varepsilon^{it})$$

for supply decisions x^{it}, cost parameters θ, and cost shocks ε^{it}. Function F captures the parametric components of costs, and function H captures unobserved shocks. We allow shocks ε^{it} to interact with supply decisions x^{it}. Average and marginal costs are

$$ac^{it} = f(x^{it}; \theta) + h(x^{it}, \varepsilon^{it}), \tag{21}$$

$$mc^{it} = f(x^{it}; \theta) + f'(x^{it}; \theta)\, x^{it} + h(x^{it}, \varepsilon^{it}) + h'(x^{it}, \varepsilon^{it})\, x^{it}, \tag{22}$$

where $f = F/x$, $h = H/x$, and primes denote derivatives with respect to x. Our approach is to invert cost shocks from average cost data and to substitute into the Euler equation, eliminating ε^{it}. Identification requires three conditions.

Separability. Costs separate into components F and H, as written above.

Invertibility. The unobserved component takes the form

$$h(x^{it}, \varepsilon^{it}) = b(\varepsilon^{it})\, v(x^{it}).$$

By equation 21, we can invert to express $b(\varepsilon^{it})$ as a function of observed (ac^{it}, x^{it}).

$$b(\varepsilon^{it}) = \frac{ac^{it} - f(x^{it}; \theta)}{v(x^{it})}$$

Differentiating and substituting,

$$h'(x^{it}, \varepsilon^{it}) = \frac{v'(x^{it})}{v(x^{it})} \left[ac^{it} - f(x^{it}; \theta)\right].$$

Equation 22 becomes

$$mc^{it} = f(x^{it}; \theta) + f'(x^{it}; \theta)x^{it} + \underbrace{ac^{it} - f(x^{it}; \theta)}_{h(x^{it}, \varepsilon^{it})} + \underbrace{\tfrac{v'(x^{it})}{v(x^{it})}\, h(x^{it}, \varepsilon^{it})}_{h'(x^{it}, \varepsilon^{it})}\, x^{it}.$$

Collecting terms, we express marginal costs as a function of terms that do not depend on ε^{it}. We thus eliminate ε^{it} from the Euler estimating equation.

$$mc^{it} = ac^{it} + \frac{v'(x^{it})}{v(x^{it})}\, ac^{it}\, x^{it} + \left[f'(x^{it};\theta) - f(x^{it};\theta)\, \frac{v'(x^{it})}{v(x^{it})} \right] x^{it}. \tag{23}$$

Nondegeneracy. Only the last term of equation 23 depends on θ. Thus, identification of θ requires

$$f(x^{it};\theta) \neq a(\theta)\, v(x^{it}) \quad \forall\, x^{it}.$$

If this condition fails, we have

$$f'(x^{it};\theta) = f(x^{it};\theta)\, \frac{v'(x^{it})}{v(x^{it})},$$

which eliminates θ from equation 23 and thus from the Euler estimating equation. Intuitively, if f and v are proportional, then variation in x^{it} cannot separately identify $(\theta, \varepsilon^{it})$ because $ac^{it} = [a(\theta) + b(\varepsilon^{it})]\, v(x^{it})$.

Our baseline specification satisfies all three conditions. Suppressing superscripts, the main text has parameters $\theta = \{\alpha, \kappa\}$ and functions

$$F(x;\theta) = \left(\alpha + \frac{\kappa}{2\bar{x}} x \right) x, \quad H(x,\varepsilon) = \varepsilon x.$$

Separability holds with $f(x;\theta) = \alpha + \frac{\kappa}{2\bar{x}} x$ and $h(x,\varepsilon) = \varepsilon$. Invertibility holds because $b(\varepsilon) = \varepsilon$ and $v(x) = 1$ yield $\varepsilon = ac - f(x;\theta)$. Nondegeneracy holds because f depends on x, while v does not.

C.3 Ore degradation

We can allow ore quality to degrade with extraction. In the spirit of Aguirregabiria and Luengo (2016), we specify future grade as a function of current grade and current production. Current grade acts as a state variable that summarizes cumulative extraction.

$$\log g^{it+1} = \log g^{it} - \delta \log s^{it} + \varphi^{i} + \psi^{t} + \epsilon^{it}, \tag{24}$$

where δ is the elasticity of degradation, φ^{i} and ψ^{t} are mine and year fixed effects, and ϵ^{it} is the error term. We estimate this specification separately for each mineral,

Figure C1: Nickel ore grade over time

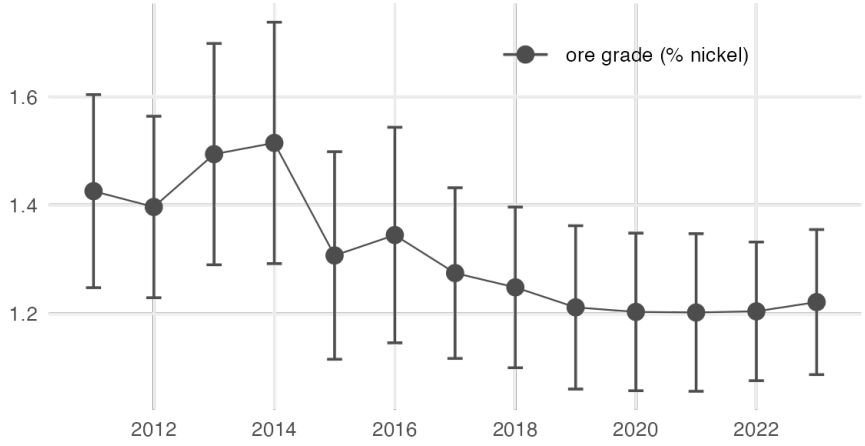

We plot world average nickel ore grades with 95% confidence intervals. We average across mines, weighting by output.

and we obtain $\hat{\delta}$ estimates of 0.02 for lithium, 0.06 for nickel, and 0.09 for cobalt. Each is small in magnitude. Figure C1 illustrates ore grade decay over time for nickel, again revealing limited degradation in the data. Our baseline model therefore imposes constant ore grade g^i.

It is conceptually straightforward to accommodate ore degradation, by which current extraction reduces future marginal revenues. We simply augment the Euler equation to capture this additional dynamic cost.

$$mr^{it} - mc^{it} - \beta\, \mathbb{E}^{it}\left[mc_d^{it+1}\right] = \beta\, \mathbb{E}^{it}\left[mr^{it+1} - mc^{it+1} - \beta\, mc_d^{it+2}\right],$$

where mc_d^{it+1} is the marginal degradation cost. It captures the impact of current extraction on future revenue through the deterioration of ore grade.

To construct mc_d^{it+1}, we substitute production $s^{it} = g^{it}x^{it}$ into grade evolution equation (24), and we differentiate to obtain the elasticity of future grade g^{it+1} with respect to current extraction x^{it}.

$$\frac{\partial \log g^{it+1}}{\partial \log x^{it}} = -\delta < 0, \quad \frac{\partial g^{it+1}}{\partial x^{it}} = -\delta \frac{g^{it+1}}{x^{it}} < 0$$

The cost of degradation is the resulting reduction of future extraction revenue. Since

revenue $p^{mt+1}g^{it+1}x^{it+1}$ is linear in grade, each unit of grade lost reduces revenue by $p^{mt+1}x^{it+1}$.

$$mc_d^{it+1} = \delta \frac{g^{it+1}}{x^{it}} p^{mt+1}x^{it+1} > 0.$$

We use this expression to construct mc_d^{it+1} from our estimated $\hat{\delta}$ and our data on grades g^{it}, prices p^{mt}, and extraction x^{it}. Estimating equation 17 becomes

$$\Delta p^{mt} g^i - \Delta ac^{it} - \beta \, \Delta mc_d^{it+1} = \frac{1}{2}\kappa^{km}\Delta u^{it} + \eta^{it}.$$

C.4 Measurement error

Suppose we observe average costs with classical measurement error ν^{it}, such that $\hat{ac}^{it} = ac^{it} + \nu^{it}$ with $\mathbb{E}[\nu^{it}] = 0$. Estimating equations 17 and 18 become

$$\Delta p^{mt} g^i - \Delta \hat{ac}^{it} = \frac{1}{2}\kappa^{km}\Delta u^{it} + \eta^{it} - \Delta\nu^{it},$$
$$\hat{ac}^{it} - \frac{1}{2}\hat{\kappa}^{km} u^{it} = \alpha^{km} + \varepsilon^{it} + \nu^{it}.$$

Measurement error adds to the error terms but does not bias our estimates, as the composite error terms $(\eta^{it} - \Delta\nu^{it})$ and $(\varepsilon^{it} + \nu^{it})$ remain uncorrelated with our regressors. Intuitively, it is straightforward to accommodate mismeasurement of average costs because they enter our estimating equations on the left-hand sides, and not on the right-hand sides. Without measurement error, we recover expectational error η^{it} as residuals in equation 17, and we recover cost shocks ε^{it} as residuals in equation 18. With measurement error, the residuals instead recover the composite error terms.

D Counterfactuals

We describe how we solve the model for our counterfactual analysis, and we present additional results for multinational mineral policy.

D.1 Projected battery expenditures

Our counterfactuals require a path for regional battery expenditures X_{kt} beyond the estimation sample. We construct this path using forecasts of battery quantities

by cathode chemistry from BMI (2025) and observed 2024 battery prices. We proceed in three steps.

First, we begin from observed 2024Q1 battery expenditures in each region. Let $q_{jk,2024Q1}$ denote observed battery demand in MWh for battery j and region k in 2024Q1, and let $p_{jk,2024}$ denote the corresponding 2024 battery price. We aggregate observed 2024Q1 expenditures across batteries to obtain base-year regional expenditure levels.

$$X_{k,2024Q1} = \sum_j p_{jk,2024} q_{jk,2024Q1}$$

Second, we use BMI forecasts of global battery demand by cathode chemistry. BMI reports forecasted future quantities in MWh at the chemistry-year level. We map these chemistries into our battery technologies, and we aggregate to obtain projected global quantities Q_{jt} for each battery j and year t. To distribute these projected global quantities across regions, we hold fixed each region's 2024Q1 share within a battery technology.

$$\omega_{jk} = \frac{q_{jk,2024Q1}}{\sum_{k'} q_{jk',2024Q1}}$$

Projected regional annual quantities are then

$$\hat{q}_{jkt} = \omega_{jk} Q_{jt}.$$

Third, we convert projected quantities into annual expenditures using 2024 prices.

$$\hat{X}_{kt} = \sum_j p_{jk,2024} \hat{q}_{jkt}$$

The expenditure path is driven by projected growth in battery quantities, while holding fixed both regional composition within a battery technology and the battery prices used to value those quantities.

D.2 Non-EV mineral demand

Our counterfactuals require a path of non-EV mineral demand ν^{mt} for each mineral m. We impose a linear demand function, calibrated from 2024 data and held

fixed over time.

$$\nu^{mt}(p^{mt}) = a^m + b^m p^{mt}$$

We recover 2024 non-EV demand $\bar{\nu}^m$ as the residual between observed 2024 mineral output and observed 2024 EV battery demand, such that $\bar{\nu}^m = s^{m,2024} - \sum_{j \in \mathcal{J}} r_j^m d_{j,2024}$. We calibrate the slope so that the elasticity of non-EV demand at the 2024 price and quantity is -0.1, which yields

$$b^m = -\frac{0.1\,\bar{\nu}^m}{p^{m,2024}}, \quad a^m = 1.1\,\bar{\nu}^m.$$

D.3 Mine-level reserves

While we do not require reserves data for the purpose of estimating mine costs in Section 5, we do require these reserves at the mine level for computing counterfactuals in Section 7. Our mining data has incomplete coverage of reserves, and so we allocate world reserves across mines to impute missing mine-level reserves. We measure world reserves with data from USGS. For each mineral, we allocate world reserves across mines in proportion to the maximum production that we observe in our sample for each mine. USGS world reserves are measured in units of mineral content, rather than crude ore extracted as in Section 5, and so we divide by ore grade g^i to convert. Reserves $R^{i,2024}$ for each mine i are

$$R^{i,2024} = \frac{\bar{s}^i}{\sum_{i' \in \mathcal{I}^{m(i)}} \bar{s}^{i'}} \frac{\bar{R}^{m(i),2024}}{g^i},$$

where $\bar{R}^{m,2024}$ is world reserves of mineral m in 2024, and $\bar{s}^i$ is the maximum observed production of mine i in our sample.

D.4 Decomposing substitution and complementarity

Substitution

We isolate the role of substitution across battery technologies with a counterfactual that freezes expenditure shares at their baseline values. As in equation 4, the

unrestricted share is

$$w_{jkt} = \alpha_{jt} + \beta_j \log\left(\frac{X_{kt}}{P_t}\right) + \sum_{j' \in \mathcal{J}} \gamma_{jj'} \log p_{j't} + \varepsilon_{jkt}.$$

We hold each component of this equation at its baseline value. Quality parameter α_{jt} and demand residual ε_{jkt} remain at their estimated values. We shut down the income channel by fixing battery expenditures and the price index at their baseline values $(\bar{X}_{kt}, \bar{P}_t)$. We shut down the price channel by evaluating $\gamma_{jj'} \log p_{j't}$ at baseline prices $\bar{p}_{j't}$. The result is equal to the baseline expenditure share by construction.

$$\bar{w}_{jkt} = \alpha_{jt} + \beta_j \log\left(\frac{\bar{X}_{kt}}{\bar{P}_t}\right) + \sum_{j' \in \mathcal{J}} \gamma_{jj'} \log \bar{p}_{j't} + \varepsilon_{jkt}$$

Quantities demanded nonetheless respond to prices through the budget constraint.

$$d_{jkt}(p_{jt}) = \frac{\bar{X}_{kt} \bar{w}_{jkt}}{p_{jt}},$$

where d_{jkt} decreases in p_{jt} for fixed $(\bar{X}_{kt}, \bar{w}_{jkt})$. Demand responds to own prices but not to cross prices, shutting down substitution in response to relative price changes.

Complementarity

We isolate the role of complementarity through battery recipes with a counterfactual that sets $r_N^\ell = 0$. We thus eliminate the joint use of nickel and lithium: nickel is used only in battery N, and lithium is used only in battery L.[21] When setting $r_N^\ell = 0$, we adjust midstream costs μ_{jt} and non-EV demand ν^{mt} such that baseline mineral prices still clear markets for each mineral in each period.[22] The adjusted values $(\hat{\mu}_{jt}, \hat{\nu}^{mt})$ satisfy

$$s^{mt}(\bar{p}^{mt}, \bar{p}^{mt+1}, \ldots, \bar{p}^{mT}) = \sum_{j \in \mathcal{J}} \hat{r}_j^m d_{jt}(\hat{\boldsymbol{p}}_t) + \hat{\nu}^{mt}, \quad \hat{p}_{jt} = \sum_{m \in \mathcal{M}} \hat{r}_j^m \bar{p}^{mt} + \hat{\mu}_{jt}$$

[21] Cobalt continues to be used in battery N but not L. There remains complementarity between nickel and cobalt, but not between lithium and the bundle of nickel and cobalt.

[22] Both objects are recovered as wedges between observed values and a value implied by the recipes, and so this procedure is equivalent to recalculating these objects based on the adjusted recipes.

for baseline mineral prices $\bar{p}^{mt}$ and our adjusted recipes $\hat{r}_j^m$. Battery N consumption no longer affects lithium consumption. Without complementarity, lithium and nickel interact only through substitution across battery technologies.

Substitution and complementarity

We obtain the net effect by freezing expenditure shares and changing battery recipes as described above. Lithium and nickel become independent markets.

D.5 Recovered parameters

Table D1: Recovered parameters

(a) Midstream costs

Battery	μ_j
N	46,679
L	46,155

(b) Non-EV mineral demand

	$\nu^m = a^m + b^m p^m$		Baseline value	
Mineral	a^m	b^m	Price	Quantity
Nickel	4,956,620	-70.677	6,964	4,506,018
Lithium	1,024,736	-18.109	13,405	931,578
Cobalt	221,541	-3.438	11,522	201,401

We report μ_j in USD/MWh, a^m in tons, b^m in tons per USD/ton. The baseline prices are in USD/ton and the baseline quantities are in tons.

D.6 Sensitivity to terminal year

Table D2: Baseline equilibrium by terminal year

Mineral	Outcome	End 2035	End 2040	End 2045
Cobalt	Price	11,522	11,241	11,121
Cobalt	Quantity	267,813	263,828	260,841
Lithium	Price	13,405	22,002	28,389
Lithium	Quantity	1,232,569	1,043,734	906,086
Nickel	Price	6,964	7,228	7,529
Nickel	Quantity	4,819,887	4,780,783	4,745,452

Prices are in USD per ton, and quantities are in tons. Each column solves for the baseline equilibrium with the indicated terminal year. Rows report 2024 prices and quantities for each mineral.

D.7 Multinational mineral policy

Figure D1: Lithium cartel

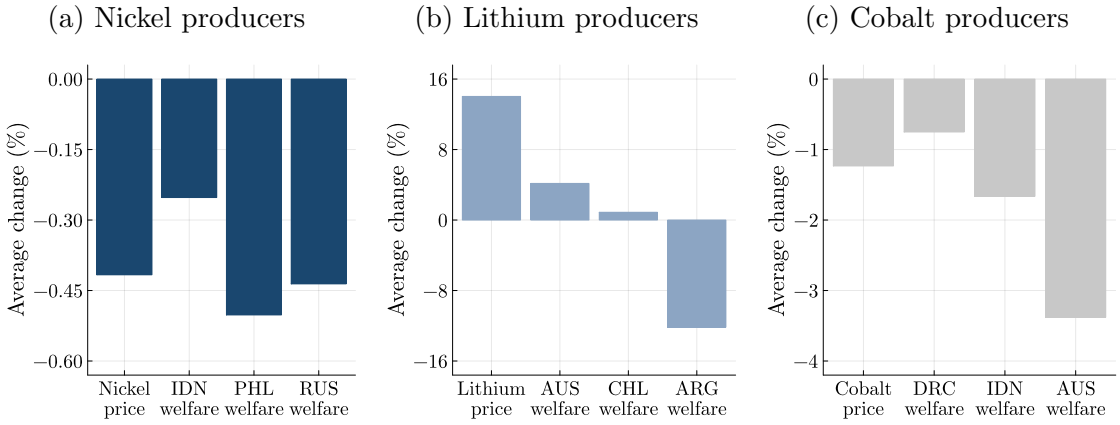

We simulate restrictive lithium policy by Australia in the form of a production tax set to maximize Australian welfare. Welfare is the discounted sum of producer surplus and government revenue over our time horizon. We then simulate restrictive cartel policy by the top three lithium producers – Australia, Chile, and Argentina – in the form of a production tax set to maximize cartel welfare. We report the differences between the cartel and unilateral scenarios in world mineral prices and in welfare for the top three producers of each mineral.

Figure D2: Cobalt cartel

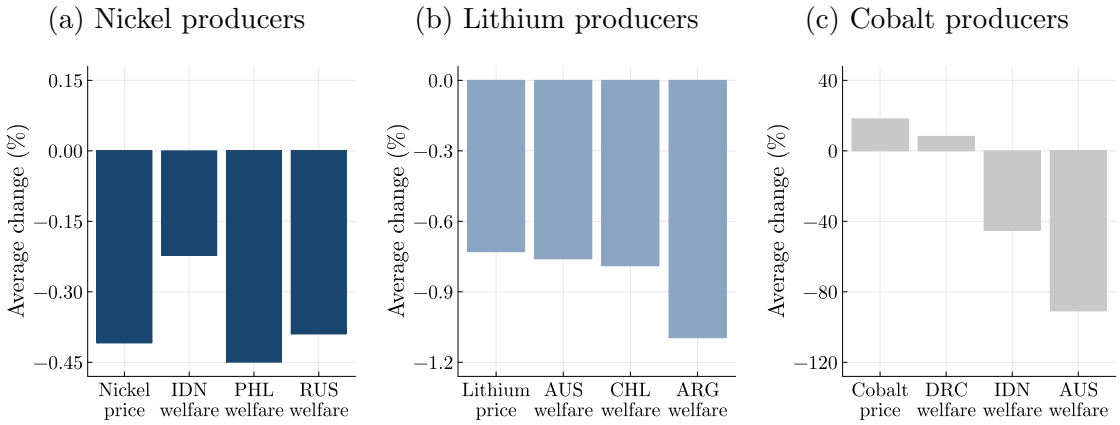

We simulate restrictive cobalt policy by the DRC in the form of a production tax set to maximize DRC welfare. Welfare is the discounted sum of producer surplus and government revenue over our time horizon. We then simulate restrictive cartel policy by the top three cobalt producers – the DRC, Indonesia, and Australia – in the form of a production tax set to maximize cartel welfare. We report the differences between the cartel and unilateral scenarios in world mineral prices and in welfare for the top three producers of each mineral.

Figure D3: Nickel cartel, green adoption

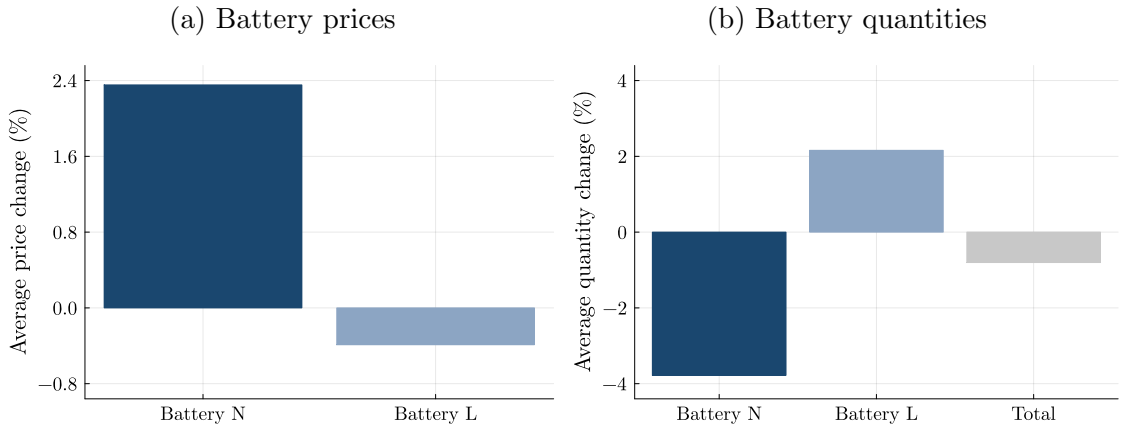

We simulate restrictive nickel policy by Indonesia in the form of a production tax set to maximize Indonesian welfare. Welfare is the discounted sum of producer surplus and government revenue over our time horizon. We then simulate restrictive cartel policy by the top three nickel producers – Indonesia, the Philippines, and Russia – in the form of a production tax set to maximize cartel welfare. We report the battery price and quantity differences between the cartel and unilateral scenarios, where battery L is LFP and battery N is NMC and NCA.

Figure D4: Multi-mineral cartel, geopolitics

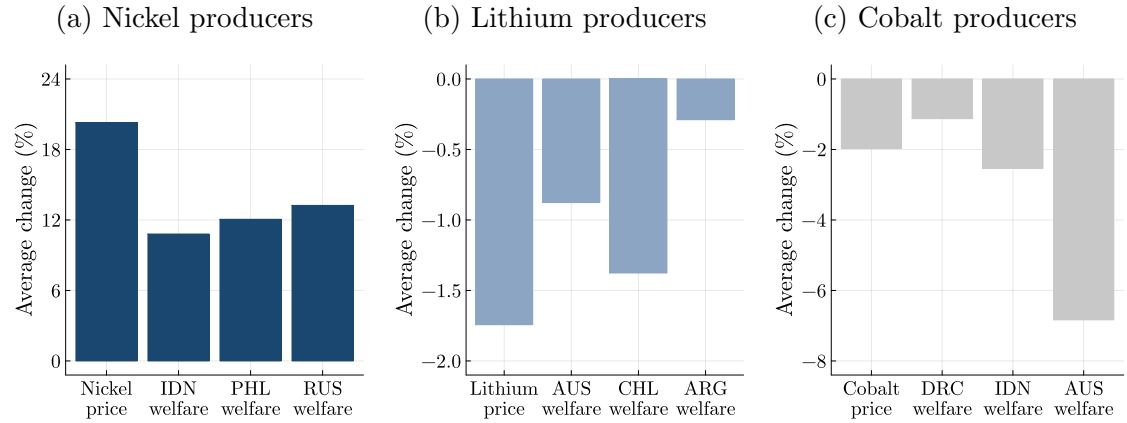

We simulate a two-cartel scenario in which a nickel cartel and a lithium cartel each set a production tax to maximize their individual welfare. Welfare is the discounted sum of producer surplus and government revenue over our time horizon. We then simulate a one-cartel scenario in which the two cartels merge and jointly set production taxes on nickel and lithium to maximize their total welfare. We report the differences between the two-cartel and one-cartel scenarios in world mineral prices and in welfare for the top three producers of each mineral.

E Notation

Symbol	Description
i	indexes mines
j	indexes battery technologies
k	indexes regions
m	indexes minerals
t	indexes years
$\mathcal{I}^m$	set of mines producing mineral m
$\mathcal{J}$	set of battery technologies
$\mathcal{K}$	set of regions
$\mathcal{M}$	set of minerals
p_{jt}	price per MWh of battery j in year t
P_t	battery price index in year t
X_{kt}	battery expenditure of EV manufacturers in region k and year t
w_{jkt}	expenditure share of battery j in region k and year t
d_{jkt}	demand for battery j in region k and year t
α_{jt}	battery technology quality in demand
β_j	battery-specific expenditure coefficient in demand
$\gamma_{jj'}$	battery own- and cross-price semi-elasticity in demand
ε_{jkt}	battery demand shock
r_j^m	tons of mineral m per MWh of battery j
μ_{jt}	refining and non-mineral manufacturing cost of battery j
p^{mt}	price per ton of mineral m in year t
d^{mt}	world demand for mineral m in year t
s^{mt}	world supply of mineral m in year t
ν^{mt}	non-EV demand for mineral m in year t
x^{it}	ore extraction by mine i in year t
s^{it}	mineral output of mine i in year t
g^i	ore grade of mine i
$\bar{x}^i$	extraction capacity of mine i
u^{it}	utilization rate of mine i in year t
R^{it}	remaining reserves of mine i in year t
C^{it}	total cost of mine i in year t
α^{km}	intercept of marginal cost for country k and mineral m
κ^{km}	slope of marginal cost for country k and mineral m
ε^{it}	mine-level cost shock
λ^{it}	shadow cost of reserves for mine i in year t
β	discount factor