Measuring Critical Raw Material Dependence

Job Market Paper

Lishu Zhang*
with Peter Tankov and Olivier David Zerbib

November 13, 2026

[Latest version]

Abstract

The supply of critical raw materials is increasingly exposed to geopolitical disruption, yet little is known about which firms ultimately bear these shocks or how they adjust. We measure structural material dependence by tracing material requirements through production paths and firms' product portfolios and translating them into cost exposure. The measure covers 12,481 firms across 55 materials, revealing a new dimension of firm heterogeneity within narrow industries. Independent benchmarks support the production links, and firms' cost responses to material-price changes closely match the magnitude predicted by the measure. Greater material-price pressure is followed by lower operating cash flow and greater external financing. These responses are stronger for price shocks identified from supply-tightening events. Dependence also predicts return volatility and downside risk, especially for firms exposed to more volatile materials. Across firms, material supply shocks can generate substantial cash-flow and financing needs, and much of this exposure comes from uses that receive little market attention.

*Job market candidate and corresponding author: lishu.zhang@ensae.fr. All authors are at CREST, ENSAE, Institut Polytechnique de Paris. Lishu Zhang is also at Tilburg University.

1 Introduction

Rising geopolitical tensions have drawn attention to supply constraints in critical raw materials (CRMs), inputs that underpin the energy and digital transitions (Adam et al., 2026; Boer et al., 2024). Policy frameworks such as the European Union's identify CRMs by combining economic importance and supply risk (European Parliament and Council of the European Union, 2024). Yet these national-level assessments do not reveal how material supply constraints reach individual firms or which firms ultimately bear them.

Whereas many supply disruptions propagate along supplier–customer links (Carvalho et al., 2021), a material supply shock can reach firms across value chains through their common reliance on the material. For many CRMs, geographic concentration and limited substitutability can turn a country-specific shock into a material-wide disruption that switching suppliers does not resolve (Graedel et al., 2015). Material supply risk is also not simply commodity-price risk. Financial derivatives can hedge price changes but not physical access, and for most CRMs even price hedging is limited by thin or absent derivatives markets (Bartram et al., 2011; International Energy Agency, 2024). Firms can instead change production recipes or invest in alternative technologies, but these adjustments are costly and take time. Material dependence therefore creates short-run rigidity, leaving firms with few immediate adjustment margins.

In this paper, we study firms' material exposures to CRM-related supply risks. To characterize this exposure, we translate the total material requirement embedded in firms' production into a proxy for the elasticity of costs with respect to material prices, which we term *structural material dependence*. This dependence captures a distinct dimension of firms' economic fundamentals beyond their industry classification and accounting characteristics, rooted in their production activities.

Such a measure is economically important in two ways. First, material shocks propagate unevenly along specific supply-chain routes. This measure helps isolate material relationships from general production linkages, identifying where economic vulnerability is concentrated during a material-specific shock. Second, CRMs are simultaneously among the inputs most exposed to geopolitical tensions and at the center of economic transitions. Measuring this dependence provides a tool to examine the impacts of many major events at the firm level.

Measuring this dependence, however, encounters two challenges. First, standard economic data do not record material requirements embedded in purchased goods. Existing measures each capture only a certain slice of these relationships through direct use and imports (Alonso et al., 2023b; Panon et al., 2026) or disclosed supplier links (Barrot and Sauvagnat, 2016). Second, existing measures do not link these re-

quirements to firms' product portfolios. Firms within the same industry can produce goods with different input requirements, which industry-level assessments cannot distinguish (Nassar et al., 2020, 2025; Alfaro et al., 2025).

We address both difficulties by tracing material requirements through production paths derived from technical evidence and linking them to firms' product portfolios. The resulting data cover 12,481 firms and 55 raw materials, comprising 53 materials on the current U.S. or EU critical-material lists and two comparison materials.

As the first step of measurement, we construct material-specific production networks, using large language models (LLMs) to assist in collecting and organizing technical evidence. Each network traces the physical input requirements connecting the raw material to intermediate and finished goods, identifying the material-specific linkages. We then combine these requirements with trade prices and industry markups to calibrate input cost shares at each stage. Propagating these shares along all production paths yields 6,140 material–product dependence estimates.

Economically, our measure proxies a product's first-order marginal-cost elasticity. Our measure is similar to but technically different from the Leontief total-requirement coefficients in Alfaro et al. (2025, 2026) in two respects. First, our measure propagates cost effects downstream from a material-price change, while the total-requirement coefficients map final demand to output from the raw-material sector. Second, we use identified material-specific product links without carrying over inter-industry relationships unrelated to the material. We show that the same measure also governs first-order output losses under pro-rata rationing.

This first-order approximation underlying our measure is most accurate for small price changes. For larger shocks, we derive bounds that translate initial dependence into a range of possible cost responses. Within these bounds, limited substitution produces greater second-order cost amplification for a given initial dependence. We therefore complement dependence with product-level substitution scores.

To obtain the firm-level measure, we map firms' product lines, as identified from their Refinitiv revenue-segment data, to the catalog of material-specific goods from our constructed networks. Weighting the product-level dependence by revenue yields the final CRM dependence measures for 75,510 firm-years. Our firm-level measure informs how much a firm's costs increase relative to its revenue following a rise in an individual material's price. Aggregating across materials provides a summary proxy for overall material dependence across measured 55 CRMs.

To remark, this construction assumes a common technology. That is, we assign the same average, static product-level technology coefficient to firms selling the same product. Because these coefficients are constructed independently of firm-specific material disclosures, realized costs, and financial outcomes, those elements cannot mechanically generate cross-firm variation in our measure.

We validate our construction of the material-specific product networks using two independent text-based benchmarks. The first comprises 119 specialist reports that were not used in network construction. Our networks recover 82.8 percent of the products they describe. The second comes from companies' material-related disclosures in annual filings with the U.S. Securities and Exchange Commission (SEC). At the firm–material–year level, product-line matches recover 74.9 percent of disclosed relationships, and the quantified dependence measure is positive for 53.7 percent. These comparisons support the physical links underlying the measure.

To put the magnitude in concrete terms, consider semiconductor and electronic-component manufacturers. Their average total dependence in 2019 is 0.094: a common 1 percent increase in material prices predicts an increase in production costs equal to 0.094 percent of identified revenue, holding output fixed. To see whether this magnitude is borne out in firms' accounts, we map each material-price change into predicted cost pressure and compare it with the subsequent change in cost of goods sold. At four quarters, a measure-implied cost increase equal to 1 percent of identified revenue corresponds to a 1.02 percent increase in cost of goods sold relative to initial sales, close to one-for-one. Because accounting outcomes do not enter the construction, this close correspondence provides a complementary quantitative validation of the measure.

All material-price variation is useful for assessing how changing input prices translate into firms' costs, but it combines demand- and supply-driven movements. To isolate the supply-driven price changes, we construct a dated list of 28 such events over 2017–2025, screened from trade-policy records and mining news, including 12 export-policy acts and 16 mine disruptions. We then decompose each material's price change into an event-driven component and the remaining variation. Specifically, for each event in the list, we define the event-driven component as the actual post-announcement price change relative to a no-event path estimated from seasonality and a common commodity factor outside event windows. We find that, around supply tightenings, operating cash flow falls and external financing rises, with larger responses to the event-driven component. For a predicted cost increase equal to 1 percent of identified revenue, operating cash flow over the following four quarters falls by 0.39 percent of initial assets when the price change is event-driven and by 0.13 percent otherwise. Net financing rises by 0.69 percent of initial assets in the first case and by 0.09 percent in the second.

Structural material dependence is also associated with market risk. Within six-digit industries and years, moving from the median firm, which has no measured dependence, to the 90th percentile of dependence ranks is associated with 1.79 percentage points higher annualized return volatility and a 0.15 percentage point larger expected shortfall. The composition of firms' material portfolios helps account for

these differences: among firms with the same dependence rank, those relying on more volatile materials have higher return volatility and downside risk, and the price-risk association becomes stronger as dependence rises. This interaction links material dependence to operating leverage: rigid input requirements can transmit material-price risk to equity holders.

The underlying exposure is economically large, and material shocks can simultaneously affect firms across conventional industry boundaries. Nearly one-third of exposed firms have at least one material dependence as large as U.S. manufacturers' exposure to natural gas (Ganapati et al., 2020). About half of the variation in dependence ranks lies within six-digit industry–year groups, and observed firm characteristics explain only 3.7 percent of the variation remaining after six-digit industry and country effects. For 25 of the 55 materials, there exists one product blocking every mapped route to downstream products that together account for at least 90 percent of the material's measured dependence.

Moreover, as material dependence is widespread across firms, the firm-level responses can add up to substantial cash-flow and financing exposure. Applying the quarter-4 estimates to firms' 2024–2025 exposures, we replay the 28 events at their observed event-driven price changes. Together, the events put $20 billion of operating cash flow at risk and imply $35.4 billion in external financing, with as much as $8 billion of cash flow at risk from a single event.

Public information captures only part of this exposure. In their respective matched samples, filings document relationships carrying 38.3 percent of measured dependence, and specialist reports cover 41.0 percent among the 30 non-investment materials they discuss. The undocumented share of cash flow at risk exceeds 80 percent for the 2023 and 2025 export controls on gallium, germanium, bismuth, indium, and tungsten. Firms are more likely to disclose materials that account for a larger share of their dependence, are harder to replace, or have recently faced export restrictions.

Our first contribution is to bring firm heterogeneity into the study of critical materials and production networks. Existing research assesses vulnerability and aggregate losses from mineral supply constraints (Nassar et al., 2020, 2025; Bogmans et al., 2026), potential firm-level output losses (Panon et al., 2026), and longer-run innovation and trade responses (Alfaro et al., 2025; Bastianin et al., 2026). Material-flow and input–output studies provide the physical foundation for this work by tracing minerals from extraction to final uses and quantifying indirect reliance across industries and countries (Sun et al., 2018; Alonso et al., 2023a,b; Baldwin et al., 2023). We construct firm-level dependence independently of firms' disclosure choices and the salience of particular material shocks.

The work of Alfaro et al. (2025, 2026) is closest to our measurement framework. Both approaches embed physical input requirements in production linkages to mea-

sure downstream exposure, but they differ along two dimensions. First, Alfaro et al. extend industrial input–output relationships, whereas we build material-specific product networks and aggregate them through firms' product portfolios across 55 raw materials. This bottom-up construction separates material-specific links from unrelated inter-industry expenditure relationships and reveals heterogeneity within narrowly defined industries. Second, the adjustment horizon differs. Their economic analysis emphasizes longer-run innovation and trade responses, while our elasticity measure captures short-run cost exposure under firms' existing technologies and product portfolios. Miranda-Pinto et al. (2024) instead construct sectoral exposure to broad mining and primary-metals industries to study inflation.

We also contribute to the literature on how corporate finance interacts with production-network risk. We connect a real input exposure to firms' cash flow, external financing, and market risk. Energy-intensive producers pass much of higher energy costs to customers (Ganapati et al., 2020), while mineral-price increases can affect investment and external borrowing (Ho et al., 2026). We document a short-run financing margin that output and accrual earnings alone do not reveal: material-price pressure can lower operating cash flow while firms raise external finance. The return-risk results connect this operating exposure to work on production-network risk and operating flexibility (Herskovic, 2018; Gu et al., 2018). The closest paper on financial-market outcomes is Ostry (2026), who constructs a news-based rare-earth supply-shock index and studies equity responses around rare-earth supply news. Our measure captures firms' pre-existing structural exposure across 55 materials and relates market risk to both the level of dependence and the price volatility of the materials behind it.

Finally, we contribute to research that measures economic exposure from disclosed information by separating physical dependence from its public visibility. Studies use customer–supplier links to examine information transmission (Cohen and Frazzini, 2008; Menzly and Ozbas, 2010) and corporate text to measure supply-chain risk and critical-mineral uncertainty (Ersahin et al., 2024; Nguyen-Tien, 2026).[1] Müller et al. (2016), for example, identify exposed rare-earth users from mentions in corporate filings. We construct production requirements independently of whether firms or analysts discuss them, then use the reverse comparison to study which economically important relationships become publicly visible. Our method also differs from work that uses LLMs to elicit and classify potential links among Harmonized System products (Fetzer et al., 2024; Klimek et al., 2025). Following Dagdelen et al. (2024), we extract material-specific relationships from scientific and technical evidence and tie each retained production-trunk relationship to a source passage. Product codes enter

[1] Related text-based measures capture political risk (Hassan et al., 2019), climate exposure (Sautner et al., 2023), and product-market relationships (Hoberg and Phillips, 2016).

only afterward to connect the physical network to trade data and firms.

The rest of the paper proceeds as follows. Section 2 derives the measure, and Section 3 constructs material–product dependence and maps it to firms. Section 4 benchmarks the measure against specialist research and corporate disclosures. Section 5 describes the firm data and sample. Section 6 documents the distribution of material dependence and its coverage in public reporting. Section 7 relates dependence to firm characteristics, operations and financing, and market outcomes. Section 8 concludes.

2 Structural Material Dependence

We characterize material dependence in a production network. The model links physical input requirements to product-level cost responses. We then aggregate exposure across firms' product portfolios, and establish when it is informative about losses from supply disruptions. The baseline assumes constant returns and fixed markups, isolating the propagation of material-price changes through unit costs. Appendix B.2 embeds this price-cost system in a quantity equilibrium with final demand and intermediate-goods market clearing.

2.1 Environment

Consider a static economy with a finite set of products $\mathcal{J}$. Material m enters the network in a traded benchmark form, denoted by $o \in \mathcal{J}$, whose price is p_m. A link $i \to j$ means that product i is used to produce j. An illustrative path is

$$o \longrightarrow i \longrightarrow \cdots \longrightarrow j.$$

Paths may branch and converge, and we focus on an acyclic production network.

Product j is made at a production stage that combines material-dependent inputs $\mathbf{x}_j = (x_{ij})_{i \in \mathcal{I}_j}$ with other inputs $\mathbf{z}_j$, including capital and labor. For each $j \neq o$, the constant-returns technology is

$$q_j = F_j(\mathbf{z}_j, \mathbf{x}_j).$$

Because the object of interest is unit cost, we normalize $q_j = 1$ for all $j \neq o$. The vectors $\mathbf{z}_j$ and $\mathbf{x}_j$ can then be read as per-unit input requirements, while equilibrium output remains unrestricted.[2] Under constant returns, unit cost equals marginal cost.

[2]Section 2.4 introduces constant-elasticity-of-substitution (CES) production to characterize larger shocks. The elasticity σ_k governs how input proportions respond to relative prices.

At input prices $\mathbf{w}_j$ and $\mathbf{p}_{\mathcal{I}_j} = (p_i)_{i \in \mathcal{I}_j}$, product j's marginal cost is

$$MC_j(\mathbf{w}_j, \mathbf{p}_{\mathcal{I}_j}) = \min_{\mathbf{z},\mathbf{x} \geq 0} \left\{ \mathbf{w}_j \cdot \mathbf{z} + \sum_{i \in \mathcal{I}_j} p_i x_i \ : \ F_j(\mathbf{z}, \mathbf{x}) \geq 1 \right\}. \tag{1}$$

We assume MC_j is positive and differentiable at the initial price vector.

At stage j, the producer takes input prices as given and charges a fixed markup over marginal cost:

$$p_j = \mu_j \, MC_j, \qquad \mu_j \geq 1.$$

This pricing rule passes marginal-cost changes proportionally into selling prices.[3] We vary p_m while holding technologies and the other-input prices $\mathbf{w}_j$ fixed.

Let $d_{ij} = x^*_{ij}$ denote the cost-minimizing quantity of input i needed for one unit of product j, for example, metric tons of ferrochrome per metric ton of stainless steel. At the initial price vector, Shephard's lemma equates the unit-cost elasticity to the input cost share:

$$s_{ij} \equiv \frac{p_i d_{ij}}{MC_j} = \frac{\partial \ln MC_j}{\partial \ln p_i}.$$

2.2 Dependence as a Cost Elasticity

Let $\beta_{mj} \equiv \partial \ln MC_j / \partial \ln p_m$ denote product j's dependence on material m. For a small log price change $z = \Delta \ln p_m$, the first-order response is

$$\Delta \ln MC_j = \beta_{mj} z + o(z),$$

where the remainder $o(z)$ becomes negligible relative to z as $z \to 0$. A downstream producer inherits its suppliers' price increases in proportion to their shares of its costs. Following this transmission through the network multiplies shares along each path and adds the contributions of different paths.

Lemma 1 (Cost exposure)**.** *Material dependence satisfies*

$$\beta_{mo} = 1, \qquad \beta_{mj} = \sum_{i \in \mathcal{I}_j} s_{ij} \beta_{mi}, \quad j \neq o. \tag{2}$$

Equivalently,

$$\beta_{mj} = \sum_{\ell \colon o \rightsquigarrow j} \prod_{k=1}^{h_\ell} s_{k,\ell}, \tag{3}$$

[3] One foundation is monopolistic competition with CES demand across differentiated products. Producer j faces isoelastic demand $q_j = A_j p_j^{-\eta_j}$, where $\eta_j > 1$, and takes the demand shifter A_j as given. Profit maximization yields the fixed markup $\mu_j = \eta_j / (\eta_j - 1)$. Variable markups can change pass-through (Atkeson and Burstein, 2008). Ganapati et al. (2020) document incomplete pass-through of energy-driven cost changes in U.S. manufacturing.

for $j \neq o$, where ℓ indexes paths from o to j, h_ℓ is the number of links, and $s_{k,\ell}$ is the input cost share at step k of path ℓ.

β_{mj} is the local elasticity of marginal cost to the material price, evaluated at the initial price vector under the existing production technology. It therefore provides a natural measure of exposure to small shocks under short-term production rigidity. For finite shocks, the exact response also depends on within-technology substitution. Section 2.4 bounds both its second-order coefficient and its full nonlinear response over the range from unit elasticity to fixed proportions.

Labor and other unexposed inputs dilute dependence along a chain, whereas multiple paths add exposure even when they share upstream stages. Equation (3) is a finite forward path sum on this material-specific acyclic network. It has the same matrix-series form as a Leontief inverse (Baqaee and Farhi, 2020, Section II.B). Our first-order result allows flexible input proportions and therefore does not impose a standard Leontief production system.

Physical requirements and cost exposure. Producing one unit of product j requires d_{ij} units of input i, costing $p_i d_{ij}$. Because $MC_j = p_j/\mu_j$, input i's share of marginal cost s_{ij} can be rewritten as $s_{ij} = p_i d_{ij} \mu_j / p_j$.

Proposition 1 (Cost decomposition)**.** *The cost elasticity can be written as*

$$\beta_{mj} = \frac{p_m}{p_j} \sum_{\ell:\, o \rightsquigarrow j} \prod_{k=1}^{h_\ell} d_{k,\ell} \mu_{k,\ell}, \tag{4}$$

for $j \neq o$, where $d_{k,\ell}$ is the physical unit requirement and $\mu_{k,\ell}$ the producing stage's markup at step k of path ℓ. The final step produces j, so its markup is μ_j.

Intermediate prices cancel along each production path. Dependence can therefore be recovered from physical requirements, stage markups, and the material's price relative to the final product's price, without observing every intermediate price.

A downstream product cannot be more exposed than its most exposed input: input cost shares sum to at most one, so $\beta_{mj} \leq \max_i \beta_{mi} \leq 1$. Its material content provides a lower bound. Let c_j denote retained material content per unit of product j, and c_o content per unit of the benchmark good. Then $\hat{c}_j = c_j/c_o$ expresses this content in benchmark units. When this content is traced to the origin, mass conservation implies $\beta_{mj} \geq p_m \hat{c}_j / p_j$. Processing losses and downstream markups raise dependence above this content-value floor. Dependence can also arise without retained material content: a platinum catalyst enters refining costs without remaining in the refined product.

Goods upstream of the benchmark. The recursion starts at the benchmark form, so its cost elasticity does not cover the ores, concentrates and intermediates from which

that form is made. A change in the benchmark price moves what these goods pay and what they are paid together. We therefore set their prices from the benchmark price using the netback terms on which concentrates and intermediates trade. The price of such a good is the benchmark value of the material it delivers, less the charges for the processing that remains and plus the value of its other contents. Let $\hat{k}_j$ denote the material charged into one unit of good j in benchmark units: $\hat{k}_j = \sum_{i \in \mathcal{I}_j} d_{ij} \hat{c}_i$ for a good with inputs and $\hat{k}_j = \hat{c}_j$ for a good without inputs. Writing $p_j = p_m \hat{k}_j + R_j$, where R_j collects the terms that do not move with p_m,

$$\frac{\partial \ln p_j}{\partial \ln p_m} = \frac{p_m \hat{k}_j}{p_j}.$$

This is a revenue-side elasticity for the producer of j. The cost share is a separate object. The elasticity exceeds one when processing charges make the good's price a leveraged claim on the benchmark price, and it falls below one when other contents account for part of the price. For a good sold as a commodity form of the material, $p_m \hat{c}_j / p_j$ is the share of its price represented by the material. We cap this share at one for the benchmark and the forms made from it. For goods upstream of the benchmark, it remains uncapped. We record the share beside the cost elasticity as the producer-side measure. Section 3.2 and Appendix D.1 describe both.

2.3 From Products to Firms

Firms inherit material exposure through their product portfolios. To express the resulting cost change relative to firm revenue, each product's elasticity is weighted by its production cost relative to firm revenue. We recover this weight from the line's sales share and markup. Let firm f sell product lines r, with revenues Y_r, total revenue $Y_f = \sum_r Y_r$, and revenue shares $\omega_r = Y_r / Y_f$. Each line has product cost elasticity β_{mr} and markup μ_r.

Proposition 2 (Firm dependence)**.** *Suppose each line uses its product's representative technology. Define*

$$D_{mf} \equiv \sum_r \omega_r \frac{\beta_{mr}}{\mu_r}. \tag{5}$$

At fixed output, the firm's variable cost C_f responds to a material-price change according to

$$\frac{dC_f}{Y_f} = D_{mf}\, d \ln p_m. \tag{6}$$

To first order, a one percent increase in p_m raises the firm's variable costs by D_{mf} percent of baseline revenue at fixed output. The weight ω_r / μ_r is line r's variable cost relative to total firm revenue. Dividing by μ_r removes the final producer's

own markup from Equation (4), while retaining upstream markups. Economically, a higher downstream markup reduces production costs relative to sales, buffering the cost shock as a share of revenue. Fixed upstream markups instead pass marginal-cost changes into intermediate-input prices and remain part of the dollar exposure transmitted along the chain.

Firm dependence also has a profit interpretation. Suppose the firm chooses its product-line prices optimally against a fixed demand system, and let q_r denote initial sales of line r. With unconstrained input deliveries, the envelope theorem gives

$$d\Pi_f = -\sum_r q_r \, dMC_r, \qquad \frac{d\Pi_f}{Y_f} = -D_{mf} \, d\ln p_m,$$

where Y_f is initial revenue. Thus, a one percent increase in p_m lowers operating profit by D_{mf} percent of initial revenue. Any induced adjustment in prices and sales has no additional first-order effect because product prices were initially chosen optimally. A shock that directly shifts product demand would generate a separate profit effect. Appendix B.1 gives the derivation.

Applying common product technologies isolates differences in firms' product mixes while abstracting from differences in their recipes and efficiency. With several material-price changes, first-order cost responses add: $dC_f/Y_f = \sum_m D_{mf} \, d\ln p_m$.

2.4 Supply Disruptions

A supply contraction can reduce downstream production through input linkages (Barrot and Sauvagnat, 2016; Boehm et al., 2019; Carvalho et al., 2021). Its incidence depends on material dependence and on how scarce inputs are allocated. We focus on two cases. The first is price allocation, in which the material price adjusts until the market clears, and the second is proportional rationing. Under rationing, q_j denotes supply-constrained output and equals realized output when the rationed input binds production.

Proposition 3 (Supply disruptions)**.** *Let supply of the benchmark material fall by $\delta > 0$ in logs.*

(i) ***Price allocation.*** *If a common price or tradable quotas allocate the material and downstream product markets clear, let $\bar{\varepsilon}_m > 0$ denote the local price elasticity of aggregate demand for material m. To first order, $\Delta \ln p_m = \delta/\bar{\varepsilon}_m$ and*

$$\Delta \ln MC_j = \beta_{mj} \frac{\delta}{\bar{\varepsilon}_m}. \tag{7}$$

(ii) ***Proportional rationing.*** *The origin's deliveries fall by $\Delta \ln q_o = -\delta$. Each downstream buyer receives the same proportional cut as its supplier's output, while other*

input quantities remain fixed. With differentiable downstream technologies, output changes to first order by

$$\Delta \ln q_j = -\beta_{mj}\delta. \tag{8}$$

For a common material and supply shock, β_{mj} ranks local cost increases under price allocation and local output losses under proportional rationing. Under rationing, cost shares are also output elasticities at the initial cost minimum, so the same network weights transmit quantity cuts.

Under price allocation, a locally downward-sloping aggregate demand schedule converts the supply contraction into a common material-price increase, which β_{mj} translates into product-specific cost responses. Under constant returns and fixed markups, allowing gross output to adjust through final demand and intermediate-goods market clearing affects $\bar{\varepsilon}_m$ but leaves the cost recursion unchanged (Appendix B.2). The rationing result instead conditions on proportional cuts to buyer-level deliveries.

Second-order responses. Initial dependence governs the first-order response to a small shock. For larger shocks, substitution possibilities affect how cost shares and output responses change. These possibilities are central to material-criticality assessments (Graedel et al., 2015). We specify constant-returns CES production to examine how limited substitution amplifies cost increases and output losses beyond their first-order predictions.

At stage k, a producer combines material-dependent inputs $i \in \mathcal{I}_k$ with an unexposed composite of all other inputs, including labor and capital, indexed by u. The corresponding marginal cost is

$$MC_k = \left[\sum_{i\in\mathcal{I}_k} \alpha_{ik}^{\sigma_k} p_i^{1-\sigma_k} + \alpha_{uk}^{\sigma_k} p_{uk}^{1-\sigma_k}\right]^{1/(1-\sigma_k)}, \qquad 0 < \sigma_k < 1.$$

The coefficients $\alpha_{ik} > 0$ are technology weights, and σ_k is the elasticity of substitution among the material-dependent inputs and the unexposed composite. Empirical estimates commonly find limited substitution among intermediate inputs over short horizons (Atalay, 2017; Peter and Ruane, 2025). We use $\sigma_k < 1$ to characterize this limited-substitution case. The composite's price p_{uk} does not respond to p_m.

As the material price rises, cost shares shift toward more material-dependent inputs because producers cannot readily substitute away from them. The product therefore becomes more sensitive to further material-price increases, amplifying the first-order cost response.

Given a downstream product j, at each non-origin stage k on paths $o \rightsquigarrow j$, β_{mk} is the cost-share-weighted mean of its inputs' dependence. Using initial shares s_{ik}, their

dispersion is

$$\mathrm{Var}_k^s(\beta_{mi}) = \sum_{i \in \mathcal{I}_k \cup \{u\}} s_{ik}(\beta_{mi} - \beta_{mk})^2.$$

The sum includes the unexposed composite with $\beta_{mu} = 0$, so the cost shares sum to one. To track how stage-k curvature passes downstream to product j, we use β_{kj} analogously for product j's dependence on upstream product k, with $\beta_{jj} = 1$.

Proposition 4 (Second-order amplification)**.** *Under the CES technologies above, with positive initial input quantities,*

(i) ***Material-price increase.*** *For $z = \Delta \ln p_m > 0$, holding unexposed-input prices fixed,*

$$\begin{aligned} \Delta \ln MC_j &= \beta_{mj} z + \frac{1}{2} A_{mj} z^2 + o(z^2), \\ A_{mj} &= \sum_{k \neq o} \beta_{kj}(1 - \sigma_k) \, \mathrm{Var}_k^s(\beta_{mi}) \geq 0. \end{aligned} \tag{9}$$

(ii) ***Material-supply cut.*** *For a log supply cut $\delta > 0$ under proportional rationing as in Proposition 3(ii),*

$$\begin{aligned} -\Delta \ln q_j &= \beta_{mj} \delta + \frac{1}{2} B_{mj} \delta^2 + o(\delta^2), \\ B_{mj} &= \sum_{k \neq o} \beta_{kj} \frac{1 - \sigma_k}{\sigma_k} \, \mathrm{Var}_k^s(\beta_{mi}) \geq 0. \end{aligned} \tag{10}$$

Amplification is stronger when inputs differ more in their material dependence and substitution is more difficult. For example, combining a material-dependent input with unexposed labor creates a positive variance term. The weights β_{kj} transmit each stage's contribution downstream. A stage with equally exposed inputs adds no amplification of its own, although it can inherit amplification from suppliers.[4] Appendix B.3 gives the derivation.

Finite price increases. Exact nonlinear responses depend on the stage elasticities σ_k. To bound them without estimating these elasticities, we hold initial cost shares fixed and consider two limiting technologies: unit substitution elasticity at every stage ($\sigma_k \to 1$) and fixed input proportions at every stage ($\sigma_k \to 0$).

Corollary 1 (Finite-shock cost bound)**.** *Under the CES technologies above, for any material-price increase $z = \Delta \ln p_m \geq 0$,*

(i) ***Cost response.*** *For each downstream product j,*

$$\beta_{mj} z \leq \Delta \ln MC_j \leq \ln\left[1 + \beta_{mj}(e^z - 1)\right]. \tag{11}$$

[4]These variance terms provide a product-level counterpart to Baqaee and Farhi (2019), who show how production-network structure and substitution elasticities govern nonlinear aggregate responses beyond first-order predictions.

The second-order coefficient satisfies $0 \leq A_{mj} \leq \beta_{mj}(1 - \beta_{mj})$.

(ii) ***Ranking across products.*** *For* $z > 0$*, consider two products* j *and* j' *with* $\beta_{mj} > \beta_{mj'}$*. If the lower bound for* j *exceeds the upper bound for* j'*,*

$$\beta_{mj} z > \ln\left[1 + \beta_{mj'}(e^z - 1)\right],$$

then $\Delta \ln MC_j > \Delta \ln MC_{j'}$ *for any substitution elasticities satisfying* $0 < \sigma_k < 1$*. This sufficient condition compares the least amplified response of* j *with the most amplified response of* j'*. If the bounds overlap, initial dependence alone does not rank their finite-shock responses.*

The lower bound corresponds to unit elasticity at every stage ($\sigma_k \to 1$): producers adjust input quantities so that cost shares remain at their initial values. The product's material-price elasticity therefore remains β_{mj} throughout the price change. The upper bound corresponds to fixed input proportions ($\sigma_k \to 0$): producers cannot reduce the physical material required per unit of output, so its cost share rises with its price. Substitution elasticities determine where the exact response lies within these bounds. Products with the same initial dependence can therefore have different finite-shock responses.

Even under fixed input proportions, the first-order approximation remains close for moderate price increases. For an exposed product, the error bound is 4.7 percent of the exact log-cost increase for a 10 percent material-price increase, and 8.8 percent for a 20 percent increase. Appendix B.3 derives these bounds and extends them to firm product portfolios.

Output losses, bottlenecks, and allocation. For fixed positive substitution elasticities, β_{mj} ranks output losses for sufficiently small supply cuts. At second order, the ranking also depends on B_{mj}: a product with lower initial dependence can have the larger predicted loss if weaker substitution and greater dispersion in input dependence make its B_{mj} sufficiently larger.

The output correction is particularly sensitive to limited substitution. The factor $(1 - \sigma_k)/\sigma_k$ in Equation (10) grows without bound as σ_k approaches zero, and the second-order approximation breaks down. At the Leontief limit ($\sigma_k = 0$), no substitution is possible and output is determined by the most constrained input. Under the same proportional-rationing rule, the origin's supply cut passes one-for-one through every exposed downstream stage, regardless of cost shares.[5] Thus $\Delta \ln q_j = -\delta$ for every product with $\beta_{mj} > 0$.

[5] Bachmann et al. (2024, Appendix B.2) illustrate the same one-for-one propagation in a fixed-proportions chain, while emphasizing that substitution across production processes, firms, and imports can attenuate aggregate losses. Proportional rationing imposes equal percentage cuts across buyers. Appendix B.4 shows that dependence continues to rank average losses under alternative allocation rules when buyer protection does not vary systematically with dependence.

Contrasting this limit with Corollary 1 highlights an asymmetry between price-mediated shocks and physical shortages. For any finite price increase, initial dependence bounds the downstream cost response even at fixed input proportions: producers can maintain output by paying more for the unchanged input recipe. Under physical rationing, by contrast, even a small positive cost share offers no protection at the Leontief limit. Thus, any required material input can become the binding bottleneck.

2.5 Financial Implications

Suppose the material-price shock affects the firm's production costs while product demand remains unchanged. The firm-level result above then implies that a small, mean-zero innovation u_m in the log material price, with variance v_m, changes operating profit relative to initial revenue by

$$\widetilde{\pi}_{mf} = -D_{mf}u_m, \qquad \mathrm{Var}(\widetilde{\pi}_{mf}) = D_{mf}^2 v_m.$$

For firms facing the same material-price process, greater dependence implies more volatile operating profits and larger losses when material prices rise. In the upper tail of profit losses, the expected shortfall of $-\widetilde{\pi}_{mf}$ equals D_{mf} times the expected shortfall of u_m at the same tail probability. To first order, exposure magnifies these losses without making the profit distribution more asymmetric.[6] Beyond first order, input complementarity makes the cost and rationed-output responses to adverse shocks convex, as Proposition 4 shows. These nonlinear operating responses provide a firm-level mechanism related to production-network tail risk (Baqaee and Farhi, 2019; Dew-Becker, 2023). Demand, pricing, and input allocation determine how this operating convexity translates into profit and return asymmetry.

Equity values reflect news about expected future cash flows and discount rates. For firms with comparable demand responses and revenue-to-value ratios, news that material prices will remain higher produces a larger negative cash-flow revision when dependence is greater. Greater dependence should therefore be associated with higher return volatility and more severe downside returns when material-price shocks are an important source of cash-flow news. These implications motivate our examination of how structural material dependence relates to stock-return volatility, downside tail risk, and valuation in the cross section.

[6] With several materials, the first-order profit component is $-\mathbf{D}_f^\top \mathbf{u}$, with variance $\mathbf{D}_f^\top \Sigma \mathbf{D}_f$, where $\mathbf{D}_f = (D_{mf})_m$ and Σ is the covariance matrix of material-price innovations. Risk therefore depends on shock variances and covariances as well as dependence. The sum $\sum_m D_{mf}$ is not itself a risk measure.

3 Measuring Material Dependence

Standard economic datasets do not report the material–product links or input cost shares needed to implement our model. Therefore, we construct both main components ourselves. We first construct production networks that link each material to the goods that contain or use it. We then value these links with physical requirements, prices, and markups to obtain each product's dependence on the material. Finally, after obtaining the product-level measures, we compute the firm-level dependence using firms' product-line revenues.[7]

3.1 Constructing Production Networks

Neither input–output tables nor supplier–customer data identify the physical route by which a specific raw material reaches a finished product. We construct material-specific production networks that connect physical goods (*nodes*) through production relationships involving that material (*edges*), tracing how each material passes through intermediate goods into finished products. A chromium-bearing alloy used in a component, for example, links chromium to the final good even when the final producer never purchases chromium directly.

We construct networks for 55 materials drawn primarily from the EU and U.S. critical-material lists,[8] representing rare-earth elements and platinum-group metals individually because their uses and prices can differ. We define each node as a single physical good at one point of production in one material's network. When a source description names several goods, we split it into one node per good. It is also common that the same good carries several materials. In this case, we include it in each relevant network and link those records so that we price and value it only once. The resulting graph approximates the material-specific production network in Section 2.1. Figure 1 summarizes the four-step construction process, which combines scientific information extraction with the physical tracing used in material-flow research (Alonso

[7]For more details, see the Internet Appendix. It provides a concrete account of the source hierarchy, LLM and human-review procedures, and rules for estimating input requirements, matching goods to prices, and mapping product lines to firms. It also presents an industry-based extension with wider firm coverage but less variation among industry peers.

[8]The materials are aluminum, antimony, baryte, bauxite, beryllium, bismuth, boron, cerium, chromium, cobalt, copper, dysprosium, erbium, europium, feldspar, fluorspar, gadolinium, gallium, germanium, gold, hafnium, holmium, indium, iridium, lanthanum, lead, lithium, lutetium, magnesite, magnesium, manganese, neodymium, nickel, niobium, osmium, palladium, phosphate rock, platinum, potash, praseodymium, rhodium, ruthenium, samarium, scandium, silicon, silver, tantalum, terbium, thulium, tin, tungsten, vanadium, ytterbium, yttrium, and zinc. Of these, 53 appear on the EU list incorporated in the 2024 Critical Raw Materials Act or the final 2025 U.S. list, covering 27 of the EU's 34 categories and 50 of the 60 U.S. minerals (European Parliament and Council of the European Union, 2024; U.S. Geological Survey, 2025a). Magnesite and gold are on neither list. Magnesite, which left the EU list in 2017 (European Commission, 2017), and gold are the two comparison materials.

et al., 2023b,a).

[Insert Figure 1 about here]

Extracting goods and production relationships. We begin with a fixed collection of engineering handbooks, industrial-chemistry encyclopedias, process journals, and government supply-chain studies. This step corresponds to named entity recognition and relation extraction from scientific text (Dagdelen et al., 2024). The LLM proceeds in three steps, with each step building on the previous one. It first constructs an ordered sequence of production stages, then identifies the physical goods within those stages, and finally extracts the directed production links between the goods.

The extraction stage produces a draft set of goods and links, which we ground separately so that every retained element has a documented technical basis. The grounding review uses engineering handbooks, industrial-chemistry encyclopedias, process journals, and, as a last resort, patents, including sources located beyond the initial collection. Every retained link carries at least one verbatim quotation, together with its source and a page, figure, or table locator where available. We replace government-study citations when a technical source supports the same relationship and retain them only when no such source can be found. An independent source-page recheck confirms 99 percent of the quotations in the reviewed trunk.

Constructing the production trunk. We call the core network assembled from these verified goods and links the *production trunk*. We construct it from the three outputs recovered above: the ordered production stages, the physical goods within those stages, and the verified direction of the input–output relationships. The resulting graph begins with the material's mined and refined forms, following each accepted input into the output it helps produce. Because technical sources typically focus on industrial processes, the trunk's endpoints mainly reflect industrial applications. We thus add an extension discovery procedure afterwards, searching for additional downstream uses and finished assemblies that technical sources may cover less fully.

Unlike the trade-classification-based approach of Fetzer et al. (2024) and Klimek et al. (2025), our discovery procedure starts with free-form descriptions of physical goods in scientific and technical sources. Following Waugh (2026), we then use LLMs to assign six-digit Harmonized System (HS6) codes to the goods in our catalog. The codes link the network to economic data but remain a separate classification layer. We preserve heterogeneous goods as distinct nodes even when they share a code.[9]

Extending coverage. We collect candidate goods from two main sources. The foundation is public agency work, including product concordances of the European Com-

[9] An ambiguous description can admit several possible codes, while some goods, such as phases within alloys, have no standalone code.

mission's Joint Research Centre, commodity and end-use tables of the U.S. Geological Survey (USGS), and customs descriptions. To complement these sources, we also perform LLM elicitation following Fetzer et al. (2024), discovering finished goods built from components at the trunk's endpoints.

These discovery sources only identify possible omitted goods. In the next step, we ground their proposed positions in the network separately and admit the new nodes under standards similar to those used for the trunk. More specifically, we place each candidate relative to the trunk, and then search technical sources for a passage that states the directed physical relationship. As this group of nodes and edges is less well documented in technical sources, grounding passages cannot always be found. In such cases, we judge the link from the good's standard recipe or established production practice and record the weaker evidence status.

Integrating the evidence. For each material, trunk and additional goods form a network, and we grade each link by the evidence behind it (Appendix C). Recycling flows that return upstream remain separate from the forward production network, preserving the acyclic structure of Section 2.1.

Across all 55 materials, the networks contain 6,848 nodes and 9,138 edges, with counts by material reported in Appendix Table C.1.[10]

3.2 Material–Product Dependence

The network shows how a material reaches a product. To measure dependence, we first estimate how much of each input a representative recipe uses. We combine those quantities with prices and markups to obtain β_{mj}, the elasticity of marginal cost with respect to the benchmark material price (Proposition 1).

Estimating input requirements. For each link from input i to product j, we need the quantity of input used per kilogram of output. When an input supplies material to the finished product, we work back from the product's material content. A recipe can draw on several sources of the same material, so we first separate their contributions. The amount contributed by i is the product's material content times the *share supplied by input i*, the fraction of that material that comes from i.

To allow for losses, we divide this contribution by the *processing yield*, the fraction of incoming material that survives the step. The input can contain other materials, so we use the *material content of input i*, measured per kilogram of input, to convert the material needed before processing to the weight of the input itself. The *physical input*

[10] Appendix C outlines the construction steps.

requirement is

$$d_{ij} = \frac{\text{material content of product } j \times \text{share supplied by input } i}{\text{processing yield for input } i \times \text{material content of input } i}.$$

Panel A of Figure 2 shows an example recipe using scrap and refined copper. Of the 200 g of copper in one kilogram of product, scrap supplies 60 percent (120 g). With the yield and copper content shown in the panel, supplying those 120 g takes 300 g of scrap.[11]

Other inputs, such as processing reagents, are used up without becoming part of the product. Their requirement is the *net quantity consumed* per kilogram of output.[12]

For each good, we obtain model-based estimates of material content, input shares, processing yields, and quantities consumed together in one LLM response. The prompt provides the good's name and physical form, its established inputs, and their production links from the network. These network descriptions were extracted from technical references, where available, from the construction in Section 3.1. We also include the corresponding customs descriptions.

For material content, the model uses a composition or chemical formula specified in the good's description. If the description leaves the composition open, it estimates a value typical of the good in that form. For example, when the description of bauxite gives no grade, the model estimates the usual weight share of alumina (aluminium oxide) and uses alumina's formula to calculate the ore's aluminium content.

The model estimates shares from the usual recipe for inputs used together, as in the example in Panel A. The input list can also contain alternatives used in different production routes. In that case, the model estimates the shares across a representative mix of routes. The same formula gives input requirements per kilogram of output for that mix. We require the shares of co-inputs sum to one, allowing for rounding. Yields and quantities consumed reflect how the model expects the inputs to be processed.

We use follow-up prompts to address calculation errors, wide estimated ranges, and inconsistencies in material content, and to distinguish material lost as waste or consumed from material recovered and reused.

To supplement these estimates with published numerical evidence, we supply original passages from the engineering references assembled for network construction in a separate LLM pass. The model extracts any reported *recovery rate*, the fraction of incoming material that reaches a particular product, together with a supporting quotation and source reference. The final yield uses a reported rate when it covers

[11] In the notation of Section 2.2, c_j counts the copper from both inputs. The d_{ij} in Proposition 1 is for a single link, here the final step on the scrap path, with scrap as input i and the finished good as product j.

[12] For process consumables, we count only the net quantity purchased by the producer. For example, for a catalyst that is regenerated and reused, we count the amount replaced during production.

the same input, product, and process, and when no material is diverted to a useful co-product. If the process also produces a useful co-product, the reported rate for the named product understates total material recovery, so we use the model-based yield. In other words, when some material reaches another useful output, as in Panel B of Figure 2, that material still counts toward the yield.

[Insert Figure 2 about here]

Benchmarking the material. For each material, we choose a traded form, usually a refined metal or separated oxide, whose quoted price serves as the benchmark p_m in the model. We normalize dependence to one at this starting point and trace the price change through the goods that use it, directly or through other inputs.

Ores and concentrates used to make the benchmark form lie upstream of this starting point. We measure their selling-price response separately. To obtain the revenue-side elasticity, we value the material entering one unit of the good at the benchmark price and divide by its selling price (Section 2.2).

Some routes bypass the benchmark form, such as alloys made directly from concentrate or scrap without first converting the material into refined metal. For these paths, we value the physical amount of material contained in the concentrate or scrap at the benchmark price and use this embodied value for downstream exposure.

Valuing inputs. Multiplying the physical requirements by their prices gives input expenditures per kilogram of output. For scrap in Panel A of Figure 2, this is the cost of buying 300 g. The cost share is this expenditure divided by the marginal cost of a kilogram of output.

For product prices, we use trade unit values, calculated as trade value divided by quantity. Our main sources are world trade unit values by HS6 code from CEPII's BACI database and U.S. import unit values by ten-digit tariff line over 2017–2019. CEPII's Trade Unit Values database, published commodity prices, and import declarations for specialized goods supplement these series.[13]

We approximate marginal cost by dividing the product price by its industry markup. We estimate markups following De Loecker et al. (2020).[14] We map each product to its producing North American Industry Classification System (NAICS) industry, starting from the Pierce and Schott (2012) concordance. Each product then takes its industry's median markup, estimated from Compustat North America and

[13]The window begins in 2017, the first year BACI reports trade under the 2017 revision of the Harmonized System, and ends before the COVID-19 pandemic disrupted trade in 2020. By default, a code's price is the three-year median of its annual value-weighted median unit value.

[14]A firm's markup is its ratio of sales to cost of goods sold, scaled by its sector's output elasticity. Following De Loecker and Eeckhout (2018), we apply U.S. sectoral output elasticities to listed firms globally.

Global over 1999–2016.[15] Appendix D.3 reports sensitivity to alternative valuation assumptions.[16]

The cost shares multiply along a path, and adding the path contributions gives β_{mj}. Intermediate prices cancel within each path, so we calculate its contribution from the physical requirements and markups along the path and the benchmark and product prices (Appendix D.1).

Product substitutability. Dependence measures the initial cost response. The substitution score captures how readily a producer can replace the material. We score each material–end-use pair using the *Substitutes* passages and end-use tables in the USGS *Mineral Commodity Summaries 2025* (U.S. Geological Survey, 2025b), supplemented by technical demand studies for rare-earth elements and platinum-group metals. We map network goods to end uses from recorded applications and downstream products. For unmatched goods, an LLM chooses from the material's sourced end-use list using the good's name, physical form, customs codes, production stage, and neighboring network products. The substitution score itself still comes from the source evidence.

The scale runs from no substitute (0) and experimental alternatives (1), through alternatives with major (2), moderate (3), and minor (4) performance or cost penalties, to several freely available substitutes with no significant penalty (5). For goods serving several end uses, we retain both the lowest score and a mean weighted by reported consumption shares. The minimum captures the use in which replacement is most difficult, and we take the minimum again when several goods share an HS6 code. The weighted mean summarizes the mix of uses served by the good. We report the scores separately and do not use them to adjust dependence.

3.3 Firm Material Dependence

Listed companies generally report revenues by broad business or product lines. To retain the most granular available information on firms' product mix, we use product-line revenue breakdowns from LSEG Refinitiv, even though this choice reduces firm coverage. We implement the firm-level measurement by matching these lines to network goods and weighting their dependence by revenue.

Our measure captures differences in product mix among industry peers using a common technology for each product. It covers the material embodied in a firm's

[15]The window ends in 2016, the final year of the published output elasticities of De Loecker et al. (2020). We use six-digit NAICS industries, moving to five digits when fewer than ten firms support the finer estimate. Missing markups and estimates below one are set to one.

[16]Producers of the same good can use different recipes. With a common fixed markup, cost shares calculated from total input use per unit of total output and the average selling price equal producers' sales-weighted average cost shares at the initial prices (Appendix B.1).

products and the material consumed in producing them along the recorded network paths. For instance, packaging is included when it is identified as part of the product. Generic product lines whose packaging format is unknown carry no inferred packaging requirement. Durable equipment and installed capital are outside the measure, which follows material embodied in or consumed per unit of current output.

Revenue data and sample. Refinitiv reports firms' product lines, each with its revenue in the fiscal years the firm reports it. Because Refinitiv does not flag subtotals, we keep only firm-periods whose product-line revenues approximately sum to total sales. Specifically, they must reconcile within 0.01 percent or fall short by no more than two percent. Appendix Table D.2 reports the firm-period sample in Panel A, the product lines in Panel B, and their final classification in Panel C.

Identifying products. We identify what each reported line sells in two stages. First, two LLMs classify it using its name, a short description, the firm's other product lines, and official descriptions of the relevant customs codes. A third resolves disagreements, and low-confidence decisions receive another review. This screen assigns service, software, and financial lines zero dependence and removes geographic lines that name no product. Second, an LLM reads each remaining line with positive revenue in the context of the whole company. It identifies the products sold and whether the firm makes or resells them, processes them for customers, provides a service, or has an unclear role. Each product is then linked to one or more of the 3,975 physical goods in our networks. If the evidence does not support a link, the LLM records whether the product lies outside the catalog or remains unidentified.

We revisit lines with unidentified products, conflicting classifications, or low-confidence matches carrying large potential exposure. Each flagged line is settled by the most reliable answer available for it, taken from a web-search reading and, for the economically largest cases, a model-based expert review. Appendix D.2 states the rules and their priority order. For 42,246 of the 53,727 lines with positive revenue, the original reading is retained unchanged.

The main measure includes goods the firm makes, goods it processes for customers, lines whose role remains unclear, and merchant lines that physically trade the material itself. Other resale lines are excluded from the main measure and enter a robustness measure.[17]

Aggregating dependence. Following Proposition 2, we divide each product's marginal-cost elasticity β_{mj} by its industry markup μ_j to express dependence per dollar of sales. We then aggregate these values from network goods to reported product lines. Because firms do not report product-specific sales within a line that

[17] Appendix Table D.2 reports classification counts.

carries several differentiated products, such as "copper wire and cable," each of the line's products then receives an equal share of its line revenue. A product matched to several network goods is valued at the mean of β_{mj}/μ_j across the matched goods that carry material m, and a product identified as outside the catalog contributes zero.[18] We exclude a product that cannot be identified from the measure.

Let $\bar{D}_{mp}$ denote the dependence assigned to product p for material m and N_r the number of products in line r. Firm f's dependence on material m in period t is then

$$D_{mft} = \frac{\sum_{r \in R_{ft}} Y_{rft}\, b_{mr}}{\sum_{r \in R_{ft}} Y_{rft} - U_{mft}}, \qquad b_{mr} = \frac{1}{N_r} \sum_{p \in \mathcal{V}_{mr}} \bar{D}_{mp}, \tag{12}$$

Here, $\mathcal{V}_{mr}$ denotes line r's product set with assigned dependence on m, and R_{ft} contains measurable lines, including non-goods assigned zero. Y_{rft} is line revenue, and U_{mft} represents revenue from products that cannot be identified or whose dependence on m cannot be quantified. The identified part of the total revenue, represented by $\sum_{r \in R_{ft}} Y_{rft} - U_{mft}$, covers on average 94.8 percent of company revenue across measured firm-periods.

As a proxy for total material dependence D_{ft}, we first sum each good's dependence estimates over the 55 materials, counting bauxite with aluminum and magnesite with magnesium once and capping the good's total at one. These totals then pass through the same lines and weights, so D_{ft} lies at or below the sum of the D_{mft}. A plain sum essentially gives equal weight to dependence on each material, which corresponds to assuming that all material prices increase by the same proportion. We acknowledge that this abstracts from actual price movements and use it only as a summary proxy.

We also record producer-side exposure P_{mft}. At or downstream of the benchmark, it measures the share of firm revenue attributable to material m sold in commodity form. It is built on the same lines, products, and identified revenue as D_{mft}. For firms selling goods upstream of the benchmark, we record the revenue-side elasticity of those goods.

3.4 Coverage and the Baseline Sample

The product panel from Section 3.2 initially contains 6,332 material–product pairs across 55 materials. We quantify 6,140 of these pairs: 5,381 cost-side dependence estimates and 759 upstream revenue-side elasticities. Appendix Table A.2 reports

[18] To expand coverage in eligible cases, we fill a good's missing material-dependence estimate following a hierarchical approximation. We first use the mean estimate of other goods under the same customs code. If no such estimate is available, we then use the mean among other goods made from the same input and, lastly, among goods in the same industry. The estimate remains missing if none of these groups provides an approximate estimate.

this product coverage and the subsequent firm sample. In more detail, for the 759 pairs upstream of the material's benchmark form, we report how each good's selling price responds to the benchmark price, given its material content. The downstream part contains 5,573 pairs at or downstream of the benchmark, of which we estimate dependence for 5,381.[19]

Following the procedure in Section 3.3, we construct a firm panel over fiscal years 2017–2025. Of the 31,086 firms with usable business-segment revenue, 14,365 report product lines, and we obtain firm-level dependence for 81,862 firm-years at 13,347 firms. A measured zero means no dependence within the materials and production activities we cover.

We further restrict to firms with identified revenue for at least 75 percent of total firm revenue. The concern is that, with lower coverage, an apparent zero can arise from unidentified products. This coverage condition reduces the sample to 75,476 firm-years at 12,478 firms and forms our ***baseline sample***.

A material-price increase can raise revenue as well as costs for firms that sell the material itself. We therefore separate upstream firms from downstream users. A firm-year is considered *upstream* when at least 25 percent of identified revenue comes from producing or reselling the material in a refined or near-pure form, or in an earlier form such as ore or concentrate. Within upstream firms, we further distinguish *producers*, i.e., firms that produce these goods, and *merchants*, i.e., those that only resell them, including trading, royalty, and streaming companies. To identify which materials a producer supplies, we require these forms to account for at least half of its revenue from goods carrying the material. All other firms are downstream *users*. For them, we measure their material dependence primarily through cost exposure.

This baseline sample is measured holding out firms' annual reports, so we use annual regulatory filings as an independent source for validation and comparison in Section 4.2. After validation, we integrate this information and refine the identification of product lines, especially where the reported business lines were ambiguous or uninformative. Section 5.1 describes the resulting ***full-information sample***, which serves as our main empirical sample throughout the analysis.

4 Comparison Benchmarks and Validation

We use two independently coded sources as benchmarks. The first is specialist research describing material applications and production relationships, often organized

[19] Among these estimates, median dependence is 0.018 per dollar of product sales and 0.001 for final goods in HS chapters 84–96. Appendix Table A.3 provides the summary statistics. The remaining 192 are not valued mostly because no usable price is available. Besides, we report 723 estimates as bounds when the underlying inputs imply cost shares above one. Appendix D explains the valuation rules and price-consistency checks.

around particular topics. The second uses annual filings that companies submit to the U.S. Securities and Exchange Commission (SEC). These filings contain the companies' own disclosures. Both sources describe material relationships qualitatively, allowing us to validate the links in our networks.[20]

We compute two recovery rates. *Recovery* is the share of reported relationships that our measure captures. *Reverse recovery* is the share of constructed relationships captured by each alternative source.

For both benchmarks, we code relationships from the surrounding context, following Gentzkow et al. (2019), and each coded relationship keeps its supporting passages. We form the comparisons only after applying the full classification and review procedure.

4.1 Comparison with Specialist Research

For specialist research, we first ask whether the networks recover the products described in the reports. We then compare material–HS6 relationships in both directions, obtaining recovery and reverse recovery.

Data and construction. We build the independent benchmark from 119 specialist documents not used in network construction: 95 from financial-analyst publishers and 24 from policy or public-research institutions. Most reports discuss materials in connection with green investment, commodity prices, battery and power demand, trade restrictions, or supply-chain bottlenecks, and few systematically examine one material's supply and uses.

To build the benchmark, we let LLMs apply common coding rules to extract claims connecting a material to a product, industry, or firm through production, use, demand, or substitution. Targeted human review then checks the resulting source-linked relationships.

Benchmark statistics. The primary benchmark contains 535 distinct material–product pairs. Each pairs a material with a normalized product description for which a report asserts an exposure or effect, with repeated descriptions counted once within each material. We compare these descriptions with network goods linked to the same material, including goods without an HS code. Appendix E.1 reports the extraction counts and matching rules.

Comparison results. The network recovers 443 of the 535 material–product pairs (82.8 percent) through exact or clear product-category matches. Allowing same-

[20] A vocabulary-frequency measure in the spirit of Hassan et al. (2019) is also feasible for these filings. We do not use one because the benchmark asks whether the text establishes an operational firm–material relationship. A material may be central to production yet appear only once, whereas repeated references can come from risk-factor or boilerplate language.

family products and adjacent production stages raises recovery to 92.5 percent. Appendix Table E.2 reports the full matching breakdown.

Table 1 reports the supplementary comparison of distinct material–HS6 pairs. The reports identify 80 such pairs, of which 69 appear in the network (86.3 percent). In the reverse direction, these same 69 pairs account for only 1.1 percent of the network's 6,365 material–HS6 pairs across the 39 overlapping materials. The network thus captures most of the relationships identified by the reports, while the reports cover only a small part of the network at this coding level. Appendix Table E.3 extends this comparison to four-digit product categories and industry assignments.

[Insert Table 1 about here]

4.2 Comparison with Regulatory Filings

Annual SEC filings provide a firm-level benchmark based on companies' own disclosures of operational relationships with individual materials. Appendix E.2 describes the filing sample, text extraction and review, scope rules, and benchmark statistics.

Comparison results. We perform two different comparison exercises. The first asks how well our quantified measures recover the self-disclosed relationships, i.e., if the disclosed firms receive positive cost dependence or positive producer-side exposure at the individual material level. The second is based on the firm's network presence, requiring that we identify a product of the firm that lies within our constructed network of the pertinent material. They are nested because every relationship with positive exposure naturally has network presence as well.

Table 1 reports these two nested recovery rates. Positive exposure captures 1,513 of the 2,815 disclosed relationships (53.7 percent). At the broader network-presence level, the measure captures 2,109 (74.9 percent). Appendix E.2 reports the material-level results and examines the disclosures that neither criterion captures.[21]

Reverse recovery is 8.7 percent among the 17,472 firm–year–material relationships with positive exposure and 2.9 percent among the 72,288 relationships with network presence. Such sparse coverage may partly reflect the selective nature of corporate disclosure (Verrecchia, 1983). We examine this possibility in Section 6.4.

5 Data and Sample

For the empirical analysis, we first describe the full-information sample, which uses annual reports to refine firms' product identification after validation. We then define

[21] Internet Appendix IA4 reports the filing comparison for the industry-based extension.

the firm-level variables and combine the sample with material prices and supply conditions. Appendix Table A.1 defines the empirical variables and lists their sources.

5.1 Full-Information Dependence Measure

The baseline measure identifies products from firms' reported business lines without using their annual reports. An annual report can, however, resolve what a short line name leaves open, such as the products behind a broad label or whether the firm makes or resells them. For the empirical analysis, we therefore reconstruct the full-information sample, exploiting the additional information from these reports for product identification refinement.

Construction. Our matched filing sample contains 9,364 SEC annual reports from 1,743 firms. Of these, 3,266 reports from 738 firms name at least one of the 55 materials. For these firms, we repeat the product-identification procedure of Section 3.3 after adding every annual-report passage that names a material. The LLM also links each reported material relationship to the corresponding business lines. We then compute full-information dependence from Equation (12) using the same aggregation and valuation rules. Internet Appendix IA5 provides the details.

Sample. We reapply the filters of Section 3.4 to the refined lines, retaining fiscal years 2017–2025 and the coverage condition of at least 75 percent of revenue being identified. This leaves 75,510 firm-years for 12,481 firms, of which 29,039 firm-years for 4,823 firms have positive dependence. Filing answers govern at least one line in 1,423 firm-years for 248 firms and change total dependence in 529 firm-years for 98 firms, 43 of them from zero to a positive value. Yet the overall ranking barely moves: the rank correlation of total dependence across firm-years measured both ways is 0.999. Compared with the baseline sample, 40 firm-years enter and 6 leave. The sample contains 1,263 producer firm-years for 253 firms, 160 merchant firm-years for 42 firms, and 74,087 user firm-years for 12,283 firms.[22]

5.2 Firm-Level Variables

Firm characteristics. LSEG Refinitiv Workspace provides firms' product, business, and geographic segment revenues, reported primary industry, and headquarters country. Annual data on firm characteristics come from Compustat North America and Global, linked to firms by their International Securities Identification Numbers (ISINs). When a firm-year is missing from Compustat, we use Refinitiv annual accounts where available.

[22]Our firm universe is the set of primary ordinary-share listings with a live quote in LSEG Refinitiv Workspace on January 31, 2026, so firms delisted before that date were not retrieved.

Standard financial characteristics are log assets, profitability measured by operating income over assets, leverage measured by debt over assets, and annual asset growth. Production characteristics cover capital expenditure, net property, plant, and equipment (PP&E), inventory, R&D, and employment, scaled by assets, together with the cost-of-sales ratio, the share of sales absorbed by cost of goods sold. Unreported R&D is set to zero and accompanied by a non-reporting indicator. The number of business segments and the share of geographic-segment revenue earned outside the firm's home country describe operating scope. After conversion at period-end exchange rates, log assets and log employees per unit of assets enter as deviations from their country–year means.

Quarterly accounts. For the quarterly analysis, we use Compustat North America and Global accounts for firms reporting at least 16 quarters with positive total assets from 2019Q3 through 2025Q4, retaining earlier eligible quarters for these firms as well. Compustat supplies sales and production costs, working-capital balances, cash, debt and interest expense, and cash-flow items. Refinitiv supplies quarterly operating income. Compustat reports operating cash flow, capital expenditure, and cash flow from financing activities year to date, so we difference them within each fiscal year to obtain quarterly flows. Changes and cumulative flows require a common reporting currency across the relevant dates.

Market outcomes. Compustat security files provide daily closing, high, and low share prices, shares outstanding, and exchange rates. They also provide adjustment factors for stock splits and total-return factors that account for dividends. Together with the annual accounts, these data give return volatility, market beta, idiosyncratic volatility, historical expected shortfall, log market-to-book, and the high–low bid–ask spread. Idiosyncratic volatility follows the country-market specification of Ferreira and Laux (2007), and downside risk uses historical expected shortfall as in Atilgan et al. (2020). For valuation and illiquidity, we use log market-to-book (Pástor and Veronesi, 2003) and the spread estimator of Corwin and Schultz (2012).

For the quarterly regressions, we use log total stock returns in the trading currency between quarter ends, computed from the adjustment and total-return factors, with the price return where the total-return factor is missing. Annual accounting variables refer to the fiscal year of each observation, while annual market outcomes are measured at, or over the twelve months ending at, the fiscal-period end. We winsorize annual financial variables within year and quarterly outcomes within quarter at the 1st and 99th percentiles.

Substitution variable. The disclosure analysis in Section 6.4 uses each product's minimum end-use substitution score from Section 3.2. A product is hard to substitute when this score is 2 or below, meaning that at least one end use has no practical

substitute or only an alternative carrying a major penalty. At the firm–material level, the hard-to-substitute share is the fraction of scored dependence arising from these products.

5.3 Material Supply Conditions

Material supply information comes from three sources: trade-policy records from Global Trade Alert (GTA), mining news on mine disruptions, and monthly benchmark prices for the materials. The trade-policy records give us an annual count of narrowly targeted export restrictions for each material, which we relate to firms' disclosure in Section 6.4. Combined with the mining news, they also yield a dated list of major, unanticipated supply tightenings. We use this list in Section 7.2 to separate supply-driven price changes from other price movements. Material prices, weighted by each firm's dependence, measure the cost pressure whose effects Section 7.2 traces. Finally, we summarize each material's price volatility and its exposure to Chinese export restrictions in two fixed measures for the market-outcome tests in Section 7.3. We describe each component in more detail below.

Trade restrictions. We construct the annual material-level *export-restriction count* from GTA, which records policy instruments affecting foreign commercial interests, their covered products, and their announcement and implementation dates. The count includes export bans, licensing requirements, quotas, taxes, and other export restrictions that GTA classifies as discriminatory. Each act counts once for every affected material, in the year of its first implementation. Because GTA records acts by the customs products they cover, we link the acts to materials through those products.

The count is intended to capture restriction activity specific to a material. We therefore exclude broad acts, such as sanctions and catch-all export controls, that list hundreds of products across the tariff schedule. Keeping acts that cover at most three of our 55 materials leaves 135 of the 355 acts. For each filing, we use the material's count in the preceding calendar year.

Supply events. We construct a dated list of supply-tightening events for 2017–2025 from two sources. We screen the 549 GTA acts linked to our 55 materials for export-policy events and a fixed corpus of mining news for mine disruptions. The 549 acts are a wider set than the 355 behind the restriction count, which keeps only implemented acts of the listed discriminatory instruments. We date each candidate to its announcement and review it against the underlying sources. An event enters only when preceding coverage did not anticipate it and the affected operations or trade flows represent a meaningful share of the material's world supply. We also require source evidence of the supply impact and no confounding material news in the event

window. Each decision is recorded with the event. The resulting list contains 28 events covering 30 materials: 12 export-policy acts and 16 mine disruptions. Appendix Table A.10 reports them. The same screening identifies seven easing events. These enter only by marking quarters excluded when fitting the counterfactual price models in Section 7.2.

Cost pressure. We combine material-price changes with predetermined dependence to measure each firm's predicted cost pressure. Because D_{mft} gives the response of variable costs to the material's price as a share of identified revenue, we define

$$Pressure_{ft} = 100 \sum_{m \notin \{\text{gold, silver}\}} D_{mft} \left(\ln p_{mt} - \ln p_{m,t-4} \right), \tag{13}$$

where $\ln p_{mt}$ is the quarterly mean of the log of the material's monthly benchmark price in U.S. dollars, and D_{mft} is the most recent annual measure of the full-information sample (Section 5.1) whose fiscal-period end precedes the quarter by at least 90 days. A value of one represents a predicted cost increase of one percent of identified revenue over the year, with bauxite counted net of aluminum.

Prices are available for 44 of the 55 materials, covering 95 percent of measured dependence. Gold and silver trade as financial assets, however, so we leave them out of pressure and enter them separately as $G_{ft} = 100 \sum_{m \in \{\text{gold, silver}\}} D_{mft} \left(\ln p_{mt} - \ln p_{m,t-4} \right)$. The other 42 priced materials cover 71 percent of measured dependence. Unpriced materials contribute zero. The monthly series come from Refinitiv for 25 materials, Investing.com for 14, Xinhua for holmium, lutetium, thulium, and ytterbium, and the World Bank for phosphate rock. Months without a quote inside the silicon series are filled from a second silicon series, using the ratio of the two where both are quoted. The Xinhua quotes end in December 2019, and from January 2020 the four series follow the equal-weight mean monthly log change of the priced heavy rare-earth oxides (dysprosium, terbium, erbium, gadolinium, and yttrium), so every priced material has a four-quarter price change throughout the sample window.

Material risk. Two fixed measures describe the risk attached to each of the 42 priced materials other than gold and silver. We define Price Risk σ_m to be the standard deviation of the material's monthly log price changes from January 2009 to December 2025, using consecutive months only. China-restriction Risk ρ_m is the log of one plus the number of export acts implemented by China over 2017–2025 that GTA links to the material. Here, an act can cover the material as part of a broader package. Because this measure captures exposure to China's restriction regime, we retain those broad controls. It uses the same GTA instrument types as the export-restriction count above, and each Chinese act counts once per material.

A firm's *basket* for each measure is the dependence-weighted average of that mea-

sure across the firm's materials,

$$PriceRisk_{ft} = \sum_m w_{mft}\,\sigma_m,\ ChinaRestriction_{ft} = \sum_m w_{mft}\,\rho_m, \quad \text{with } w_{mft} = \frac{D_{mft}}{\sum_k D_{kft}}, \tag{14}$$

where the sums run over the 42 materials. The baskets capture the composition of exposure, and their correlation across firm-years is 0.01.

5.4 Empirical Samples and Summary Statistics

Material groups. We use a fixed grouping of materials by market role throughout the analysis. Metals with liquid exchange-traded futures before our return sample form the investment-metal group, namely gold, silver, platinum, and palladium, traded at CME Group, and aluminum, copper, lead, nickel, tin, and zinc, traded at the London Metal Exchange. In the pre-sample data, their price changes comove more strongly with world equity returns than those of the other materials with available price series.[23] We refer to the remaining 45 as other materials and focus the public-coverage comparison in Section 6.4 on this group.

Empirical samples. For the firm-level analysis in Section 7, we use downstream users in the full-information sample and exclude upstream firms, whose revenues can rise with material prices.[24] The sample includes measured zero dependence. The main regressions using the dependence rank or the rank of the largest single material retain these observations. Regressions using log dependence or the number of materials retain only positive dependence, as do the appendix checks that re-estimate the ranks in this subsample. The market-outcome regressions require at least five firms per six-digit NAICS industry–year cell, and sample sizes vary with outcome availability.

The quarterly analysis in Section 7.2 uses the subset with quarterly accounts from 2018Q2 through 2025Q4, in four-digit NAICS–quarter cells of at least five firms. This leaves 158,068 firm-quarters for 7,642 firms in 82 countries, including 3,053 firms with nonzero cost pressure. At the firm–material level, each observation is a firm, a priced material outside gold and silver with positive dependence, and a quarter. This gives 374,543 observations for 3,054 firms. Positive dependence does not require nonzero pressure, because a material's price can be unchanged over a firm's quarters, so the firm counts differ by one.

[23] Over 2009–2017, the median coefficient of monthly price changes on the world equity return is 0.95 for these ten metals, significant at the 5 percent level for eight of them, compared with 0.12 for the 28 other materials with monthly prices.

[24] Product-line revenues are available for most firms from 2018. Annual market outcomes extend through 2025, with coverage varying by outcome (Appendix Table A.6).

Table 2 summarizes the full-information sample. Panel A covers all 75,510 firm-years, of which 38.5 percent have positive dependence. Panels B and C report market outcomes and firm controls. Table 3 summarizes the quarterly sample: material price changes, cost pressure, a broad-window supply-driven component, and the firm outcomes through quarter 4. The broad-window component assigns the full four-quarter price change to supply-driven pressure when the event list contains a tightening for the material in the current or previous quarter. Section 7.2 instead uses the event-driven component estimated relative to a no-event price path. Cost pressure is zero for firm-quarters with no dependence on a priced material, and it takes both signs where prices move, so the absolute rows report its size. Among firm-years with positive dependence, mean total dependence is 0.115 and the median is 0.022. To put the median in perspective, a hypothetical one percent increase in all material prices would raise variable costs by about 0.02 percent of identified revenue, holding output fixed. Dependence on the largest single material has a median of 0.013, and the median number of materials with positive dependence is 4. The regressions in Table 15 use the users among these firm-years, with 48,476 observations from 9,172 firms for market-to-book, 20,282 of them with positive dependence. The regressions restricted to positive dependence use a separate sample of 18,330 observations, formed with the same minimum of firms per industry-year among firms with positive dependence.

[Insert Table 2 about here]

[Insert Table 3 about here]

6 Material Dependence across Industries and Firms

We examine how material dependence varies across firms and industries and how selectively it appears in public disclosures. Unless otherwise noted, dependence refers to a firm's total dependence summed over the 55 materials.

6.1 Extent and Distribution of Material Dependence

Material dependence varies widely, even among firms in the same narrow industry. Industry-level assessments can therefore obscure heavily exposed firms whose industry peers have little or no material dependence.

Prevalence and magnitude. In the 2019 descriptive cross section, 3,158 (39.1%) of 8,070 firms have positive measured dependence on at least one of the 55 materials.[25]

[25]We use 2019 because it falls within the 2017–2019 period used to construct trade weights and precedes the pandemic and the disruptions of 2022. Other years show similar industry profiles. The descriptive statistics weight firms equally.

Among exposed firms, total dependence is skewed toward high values. It averages 0.114, compared with a median of 0.022 and a 90th percentile of 0.400. These firms depend on 6.9 materials on average.

For scale, we compare dependence on individual materials with a natural-gas benchmark calibrated from Ganapati et al. (2020). Expressing their marginal-cost response relative to revenue under constant returns gives a reference value of approximately 0.05.[26] At least one material-specific dependence exceeds this rounded benchmark for 953 exposed firms (30.1%). Under equal proportional price increases, the implied cost effect for nearly one-third of exposed firms is therefore comparable to or larger than this natural-gas benchmark. Copper, gold, and aluminum are most often these firms' largest exposures.

Across industries and production chains. To assess what industry classifications reveal about material dependence, we report industry profiles and the variation explained by industry membership in Table 4. Panel A reports industry means for downstream users in 2019, including firms with zero measured dependence, with semiconductor and electronic component manufacturers as the benchmark industry. Other electrical equipment manufacturers have significantly higher total and largest single-material dependence than this benchmark.

Industries also differ in the breadth and length of their material supply chains. We measure breadth by the number of material–product nodes reached through a firm's product lines, and length by the production steps from material entry points to those products.[27] Communications equipment reaches 12.7 measured nodes on average, compared with 2.9 in basic chemicals, yet its dependence is lower. Its paths are also longer, at 6.1 steps compared with 5.1. Motor vehicle parts have the longest paths among the manufacturing industries in the table, at 6.2 steps.

Nickel and copper are used across a broad industrial base. Potash, phosphate rock, scandium, and ruthenium are concentrated in far fewer user industries.[28] This breadth can limit adjustment through final demand. Shifting expenditure away from one nickel- or copper-intensive product may redirect it toward another product that uses the same material. When substitution within production is also difficult, losses

[26] Ganapati et al. (2020) estimate a coefficient of 0.291 on the interaction of log natural-gas prices with the gas share of electricity generation (Table 3, Panel A, column 6), with an average natural-gas generation share of 0.20. This implies a marginal-cost elasticity of 0.0582. Dividing by their Table 2 mean markup of 1.15 gives a comparable dependence measure of $0.0582/1.15 \simeq 0.0506$, rounded to 0.05. The comparison holds output fixed.

[27] Related measures describe the number of production stages embodied in a product and its distance from final use (Antràs et al., 2012).

[28] We summarize this concentration using the inverse Herfindahl index of each material's dependence shares across three-digit industries. The index equals the number of industries if dependence is distributed evenly among them. Its median across materials is 2.9, compared with 1.6 for potash, 1.4 for phosphate rock, 1.4 for scandium, and 1.9 for ruthenium, and with 5.1 for nickel and 4.6 for copper. Appendix Figure A.1 shows the industry shares and the index.

can exceed first-order predictions, as shown by Baqaee and Farhi (2019).

For upstream firms, a material-price increase can raise producer revenue as well as input costs. Panel B therefore reports cost dependence, the producer revenue share, and their difference for the firms defined in Section 3.4. All three measures are multiplied by 100. Metal ore miners have greater revenue than cost exposure, with a net of -27.91. For nonferrous metal and aluminum firms, the two largely offset, leaving positive net cost exposures of 5.54 and 6.10.

[Insert Table 4 about here]

Within-industry differences. Narrow industries explain only part of the variation. Six-digit NAICS industry–year averages account for 52 percent of the variation in users' total dependence ranks, leaving about half among firms in the same narrow industry and year (Table 4, Panel C). At the individual material level, industry affiliation explains even less. Six-digit industry–material–year averages account for 28 percent of the variation, leaving nearly three quarters within industries.[29]

Among exposed semiconductor and other electronic component manufacturers (NAICS 3344), the tenth and ninetieth percentiles of total dependence are 0.004 and 0.327. A common one percent rise in material prices would raise variable costs by about 0.004 and 0.327 percent of identified revenue at these two points in the industry distribution.[30]

These differences persist over the sample. Among the 2,753 firms with positive dependence in both 2019 and 2023, the rank correlation of total dependence is 0.95. Appendix Table A.4 separates persistent firm differences from changes within firms. The latter account for only 3.4 percent of the variation in dependence ranks. The measure's construction contributes to this persistence: we hold product coefficients and the goods assigned to each business line fixed, so measured dependence changes only with business-line revenue shares.

Dependence itself is persistent, but cost pressure varies within firms as material prices move (Equation (13)). The series obtained by holding each firm's dependence at its first-year value has a correlation of 0.96 with actual cost pressure among exposed firms, and within-firm changes account for 67.5 percent of the variation in actual cost pressure. This source of time-series variation differs from Hassan et al. (2019): their firm-level political risk measure varies mostly within firms because exposure changes over time, whereas our dependence measure is largely fixed and cost pressure moves with material prices.

[29]The shares are adjusted R^2 from regressions of the rank on the indicated fixed effects alone. Country-by-year averages account for only 4 percent, and broader industries for less than narrow ones, 24 percent at the two-digit and 44 percent at the four-digit level. Firm averages account for 97 percent of the variation in dependence ranks, consistent with persistent differences across firms.

[30]The percentiles use the 2019 cross section.

6.2 Concentration across Firms and Materials

Material dependence is concentrated among a small share of firms. Across materials, the most dependent 10 percent of firms with positive dependence account for a median 93.6 percent of summed dependence (Figure 3, Panel A), while the median share of the top 1 percent is 42.0 percent. Excluding upstream firms lowers the median top-decile share only to 87.3 percent, indicating that dependence is also highly concentrated among users.[31] Even for scandium and bauxite, where dependence is most evenly distributed, the top decile accounts for 46 and 64 percent.

The raw count of materials and the economic concentration of exposure differ sharply within firms (Panel B). The median exposed firm has positive measured dependence on four materials. Its largest exposure accounts for 74.4 percent of total dependence, and its three largest account for 99.9 percent. Expressed as the number of equally weighted materials that would produce the same concentration, the median effective number is 1.6.[32] Appendix Table A.5 shows that this pattern is not confined to firms with few material links. Among the 739 firms exposed to at least ten materials, the three largest exposures still account for a median 87.0 percent of the total. This concentration describes price-mediated cost exposure. A shortage of an essential, hard-to-replace input can interrupt production even when its cost contribution is small.

Diversification across product lines can still moderate a firm's overall cost exposure. Among firms with at least two exposed product lines, total dependence is a median 63 percent of the dependence of their most exposed line.

[Insert Figure 3 about here]

6.3 Material Links and Economic Vulnerability

Figure 4 places measured dependence in the production network. Panel A compares each product group's share of total dependence with its share of mapped goods. Goods at the refining stage or earlier carry only 16.7 percent of dependence. Metals and metal articles carry 40.1 percent while accounting for 18.6 percent of mapped goods. Vehicles and instruments account for 15.5 percent of goods but carry 2.9 percent of dependence.

Distance from material entry points may also shape visibility. Panel B reports this distance from ores, recovered materials, or purchased refined inputs. For each firm, we average the shortest-route distances to its products using their contributions to

[31] The comparison excludes the 147 upstream firms among these firms, as defined in Section 3.4.

[32] The effective number is measured by the inverse Herfindahl index, $N^{\text{eff}}_{ft} = 1/\sum_m \pi^2_{mft}$, where π_{mft} is material m's share of the firm's total dependence. We compute the index for each firm in 2019 and report the median.

dependence as weights. The median exposed firm lies 5.2 production steps from an entry point, with the 10th and 90th percentiles at 3.1 and 7.4 steps. At that distance, firms often buy processed chemicals or subassemblies. The raw material is harder to see, and filings and specialist reports may therefore give the underlying requirement little attention. Public reports, for example, commonly relate mineral supply to batteries, electric vehicles, and other energy technologies (International Energy Agency, 2025), while the firm-level mapping reaches products well beyond these prominent uses.

[Insert Figure 4 about here]

Product categories can also connect otherwise separate material chains. We measure this overlap by the share of each material's total dependence carried by categories that also map to another material (Appendix Figure A.2).[33] Categories shared by nickel and cobalt carry 65.8 percent of cobalt dependence and 37.0 percent of nickel dependence. Where this overlap represents joint use and both materials are hard to replace, a shortage of either can constrain the same production process.

A recent sequence illustrates this joint vulnerability. China announced tungsten export controls on February 4, 2025. On February 24, the DR Congo suspended cobalt exports. Cemented-carbide cutting tools use tungsten carbide with a cobalt binder, and categories shared by the two materials carry 51.6 percent of tungsten dependence and 6.1 percent of cobalt dependence. Producers using both inputs therefore faced restrictions on the same production activity within twenty days.

Some downstream chains for the same material pass through a common intermediate. To measure how much dependence a disruption there could reach, we calculate, for each good other than an entry point, the share of a material's dependence carried by products whose mapped routes all pass through it. The largest share is the material's bottleneck share.[34] Table 5 reports this share for the 37 materials where it reaches at least half. It is at least 90 percent for 25 of the 55 materials. Tantalum, osmium, niobium, antimony, feldspar, and manganese, which the table omits, instead have parallel routes, with shares below 30 percent. A disruption at a high-share bottleneck would therefore reach several downstream chains at once, generating correlated losses among firms that may not transact with one another.

[Insert Table 5 about here]

[33] All overlap shares use the 2019 cross section. We sum contributions to firm-level dependence across firms within each HS6 category.

[34] Total dependence is summed across firms with positive dependence on the material in 2019.

6.4 Disclosure Selectivity

Patterns of firm disclosure. We test whether firms disclose larger or harder-to-replace material requirements more often. We match observed filings to 7,691 firm-years of the full-information sample across all 55 materials, including 65 firm-years at 16 upstream firms. Disclosure denotes an affirmed current operational relationship under the rules of Appendix E.2. The full-information measure draws on these filings at firms whose reports name a material, so part of the association between measured dependence and disclosure is built in. We deliberately use this full-information sample to minimize measurement error and better distinguish non-disclosure from incomplete product identification.

We estimate a linear probability specification at the firm–material–year level

$$Disclose_{mft} = \mu_{mt} + \eta_{ft} + \sum_{q=1}^{4} \lambda_q \mathbf{1}\{Quartile_{mft} = q\} + \delta\, Upstream_{mft} + e_{mft}, \qquad (15)$$

where $Disclose_{mft}$ equals one if firm f discloses a relationship with material m in year t and zero otherwise. The variable $Quartile_{mft}$ groups the positive dependence of downstream firms into quartiles within each material and year, $Upstream_{mft}$ marks an upstream firm's relationships with positive cost dependence or producer share, and all other pairs form the reference group. Firm-by-year fixed effects η_{ft} compare materials within the same filing, while material-by-year fixed effects μ_{mt} absorb reporting differences common to each material.

Table 6 presents raw disclosure rates in Panel A and the regression estimates in Panel B. Downstream firms disclose 8.8 percent of positive measured relationships, with the rate rising from 2.7 percent in the bottom dependence quartile to 20.0 percent in the top quartile, compared with 0.4 percent of firm–material–year observations with zero measured dependence. Upstream firms disclose 58.2 percent of their positive relationships.[35]

The pattern also holds within filings. In the estimates of Equation (15) (Panel B, column 1), the top-to-bottom quartile gap remains 16.3 percentage points, and upstream relationships are disclosed 56.9 percentage points more often than the reference group, both significant at the 1 percent level.

Among downstream firms' positive relationships, we separate replacement difficulty from cost importance by adding the hard-to-substitute share of Section 5.2 while controlling for the material's share of total firm dependence.[36]

[35]The 0.4 percent disclosure rate conditions on measured-zero relationships, whereas the validation match rate conditions on disclosure. Appendix E.2 discusses the scope and product-information gaps behind these disagreements.

[36]We weight customs codes by their contributions to dependence and exclude codes without a substitution score from both numerator and denominator. The share is positive for 17.9 percent of

Cost importance explains part of disclosure, yet replacement difficulty retains an independent association. A ten-percentage-point increase in the material share is associated with a 3.08-percentage-point increase in disclosure probability, significant at the 1 percent level (column 2). At the same material share, moving the hard-to-substitute share from zero to one is associated with a 3.38-percentage-point higher predicted disclosure rate, a difference significant at the 5 percent level (column 3). The material-share coefficient remains essentially unchanged at 3.06 percentage points.[37] Firms thus appear more willing to report inputs that would be costly to replace, even conditional on their cost importance.

To examine whether disclosure follows changes in supply risk, we relate it to the export-restriction count of Section 5.3 in the calendar year before filing, conditional on the material's importance to the firm. Here we retain firm-by-year fixed effects and replace material-by-year effects with material fixed effects, which absorb persistent differences in disclosure across materials. A one-standard-deviation increase in export restrictions is associated with a 0.50-percentage-point higher probability of disclosure, significant at the 5 percent level (column 4), consistent with restrictions drawing attention to supply risks. Adding the hard-to-substitute share leaves the export-restriction coefficient at 0.51 percentage points, also significant at the 5 percent level, while the hard-to-substitute coefficient remains similar at 3.39 percentage points but is no longer statistically significant under two-way clustering by firm and material (column 5).

[Insert Table 6 about here]

Economic coverage of public information. Because firms disclose larger exposures more often, relationship counts may understate how much measured dependence public sources cover. Table 7 therefore reports dependence-weighted coverage for corporate filings and specialist reports. For this comparison, we use the 2025 dependence measure of downstream users, the latest year in our sample, and focus on the 45 materials outside the investment-metal group (Section 5.4). The specialist corpus includes 38 reports published in 2026 among its 117 dated reports, so its coverage rate is an upper bound on the coverage of reports in our corpus available by 2025.[38] The filing comparison covers 1,480 positive relationships across 248 firms and 45 materials, while the specialist-report comparison covers 2,696 firms with positive dependence on the 30 of these materials the reports mention.

We measure each source's *dependence-weighted coverage rate*, the share of measured dependence carried by relationships the source identifies. For the matched filing

positive firm–material–year observations in the matched filing sample.

[37] Column (3) excludes relationships with no code carrying a substitution score.

[38] Four in five dated reports were published between 2023 and 2026.

sample, this rate is

$$\frac{\sum_{m,f} D_{mf} Disclose_{mf}}{\sum_{m,f} D_{mf}}.$$

Unlike firm-level filings, specialist reports identify relationships at the product level. We therefore decompose each firm's measured dependence into product contributions, $D_{mf} = \sum_j \kappa_{mjf}$, where κ_{mjf} is the contribution of product j through the firm's product lines. The coverage rate over the same 30 materials is

$$\frac{\sum_{m,f,j} \kappa_{mjf} Named_{mj}}{\sum_{m,f,j} \kappa_{mjf}},$$

where $Named_{mj}$ equals one if specialist reports name product j as using material m.

Even after weighting by dependence, public coverage remains limited. Filings identify 38.3 percent of measured dependence in their matched sample, and specialist reports identify 41.0 percent in theirs.

Coverage remains low in some groups carrying substantial measured dependence. Mining and primary metals carry the largest share of dependence among filers, 19.0 percent, yet filings disclose only 5.9 percent of it (Panel A). Specialist reports identify 21.7 percent of it in this industry group. Machinery and vehicles show a similar gap, with coverage rates of 6.3 percent in filings and 25.8 percent in reports. Electrical equipment is much better represented. Filings and specialist reports cover 77.2 and 78.9 percent, respectively.

Minerals and ceramics combine a large contribution to dependence with little public coverage (Panel B). They account for 22.8 percent of dependence among filers and 15.1 percent in the report comparison. Yet filings disclose only 7.8 percent of this dependence and reports identify 2.1 percent. Filings also disclose 8.8 percent of the dependence carried by vehicles and instruments and almost none of that carried by machinery and pharmaceuticals. These gaps show that public sources miss substantial downstream dependence.

[Insert Table 7 about here]

7 Material Dependence and Firm Outcomes

7.1 Material Dependence and Firm-Level Heterogeneity

We first assess how much of the documented variation in material dependence conventional firm characteristics capture.

Table 8 relates material dependence to firm characteristics. Downstream firm-years with positive total dependence are ranked across all firms and years on a scale

from zero to one, while those with measured zero dependence receive a rank of zero. Continuous characteristics also enter as percentile ranks. Our set of financial and production characteristics follows Bolton and Kacperczyk (2021) and Florackis et al. (2023). All specifications include year effects, with standard errors clustered by firm.

More dependent firms are larger and faster growing in the specification with standard firm characteristics (column 1). Production characteristics account for more of the variation: the largest coefficients in column 2 are on the share of sales absorbed by the cost of goods sold (the cost-of-sales ratio), R&D intensity, and inventory. Capital expenditure and fixed assets are also positively related to dependence. Adding foreign sales and the number of business segments leaves these relationships broadly similar on the smaller sample in column 3.

These associations weaken substantially after adding six-digit industry and country fixed effects (column 4). The cost-of-sales, inventory, and R&D coefficients remain positive and significant, as do those on capital expenditure and fixed assets. Leverage is also positively related to dependence within industries and countries. Yet the observed characteristics explain only 3.7 percent of the within-group variation. Much of the variation in material dependence therefore remains beyond these conventional measures of firm activity.

The specification with firm effects accounts for nearly all total variation ($R^2 = 0.976$), consistent with the persistence of dependence (column 5). Within firms, the cost-of-sales ratio, fixed assets, inventory, and leverage remain positively related to dependence, while capital expenditure and R&D are no longer significant. The relation with leverage raises the question of whether material-price increases create financing needs. We turn to the operating and financial responses that accompany these price movements.

[Insert Table 8 about here]

7.2 Responses to Material-Price Changes

A rise in material prices increases production costs for exposed firms. What follows depends on how firms adjust sales, working capital, financing, and investment. We examine these responses jointly, complementing evidence on incomplete energy-cost pass-through (Ganapati et al., 2020) and on investment and external borrowing after mineral-price increases (Ho et al., 2026).[39]

Cost pressure, defined in Equation (13), weights each material's four-quarter price change by the firm's dependence on that material. It is a predicted increase in variable costs, in percent of identified revenue. In the benchmark of Section 2, holding

[39]Related work studies potential short-run output losses from critical-input disruptions (Panon et al., 2026) and longer-run material-saving innovation (Alfaro et al., 2025).

output and demand fixed, one percentage point of pressure lowers operating profit by one percent of identified revenue. Because neither the accounting outcomes we study nor the specialist reports used for validation enter the measure's construction, the exercise also tests whether the measured links carry information about firms' operations, including relationships that the reports do not document.

For firm f in quarter t and horizon h, we estimate

$$\begin{aligned}\Delta y_{f,t+h} = \beta_h \, Pressure_{ft} + \delta_h \, Later_{f,t,h} + \mathbf{G}'_{f,t,h}\lambda_h \\ + \gamma_1 D_{ft} + \gamma_2 \ln A_{f,t-1} + \alpha_{f,q(t)} + \alpha_{s,t} + \alpha_{c,t} + \varepsilon_{f,t+h},\end{aligned} \tag{16}$$

where β_h is the response to a one-percentage-point increase in pressure. The full-sample standard deviation of pressure is 1.19 percentage points. For sales, ratios, balance-sheet items, and stock returns, $\Delta y_{f,t+h}$ is the change from $t-1$ to $t+h$. For the remaining flow variables, it is the sum from t through $t+h$. Balance-sheet items and cumulative flows are scaled by assets at $t-1$.[40]

Prices continue to move after quarter t. For $h > 0$, we therefore control for subsequent price changes following Teulings and Zubanov (2014):

$$Later_{f,t,h} = 100 \sum_{m \notin \{\text{gold,silver}\}} D_{mft}(\ln p_{m,t+h} - \ln p_{mt}),$$

which applies dependence at t to the price change after t. The vector $\mathbf{G}_{f,t,h}$ collects the corresponding initial and later-price terms for gold and silver.[41] We also control for total dependence D_{ft}, taken from the annual measure used to construct pressure, and for lagged log total assets.

Firm-by-fiscal-quarter effects $\alpha_{f,q(t)}$ remove firm levels and firm-specific seasonality. Four-digit NAICS industry-by-quarter and country-by-quarter effects absorb common movements. The coefficient β_h therefore relates within-firm changes in pressure to outcomes net of these movements. Standard errors are clustered by firm and quarter. Material prices reflect demand as well as supply, so the estimates also capture demand movements accompanying higher input costs. As a timing check, we estimate the same regression for the change from $t-9$ to $t-5$, before the four-quarter price-change window begins.

Figure 6 plots the estimates from quarter 0 to quarter 8. Higher pressure is followed by rising nominal sales and production costs, while operating cash flow falls.

[40] Amounts are in the firm's reporting currency. We measure a change only when that currency is the same at both endpoints and sum a flow only when it remains the same throughout the window. The last outcome quarter is 2025Q4, when the price data end.

[41] Pressure, the later price change, and the gold and silver terms are each winsorized at their 1st and 99th percentiles. At horizon h, every material with positive priced dependence at t must also have a price at $t+h$.

Table 9 compares the dollar changes in sales, the cost of goods sold, and gross profit. At quarter 4, each percentage point of pressure is associated with an increase of 1.02 percent of initial quarterly sales in the cost of goods sold and 0.98 percent in sales. Gross profit shows no significant response. These estimates are close to the pattern implied by full dollar pass-through with quantities held fixed, and the cost response matches the scale predicted by the measure.

[Insert Table 9 about here]

Per percentage point of pressure, cumulative operating cash flow falls by 0.21 percent of initial assets through quarter 4. Higher inventories and receivables tie up cash in inputs and credit extended to customers, while higher payables provide a partial offset through supplier credit. On the common sample reporting operating income and all three working-capital items, these items account for about two-thirds of the cash-flow decline through quarter 2 but only a third through quarter 4.[42] By quarter 8, working capital no longer accounts for the cumulative shortfall.

Borrowing rises alongside the cash-flow shortfall. By quarter 4, cumulative net financing reaches 0.29 percent of initial assets and total debt rises by 0.23 percent. Short-term debt also increases, while cash holdings show no significant response. Because total debt rises, the short-term-debt response is not solely a shift in maturity.[43]

Nominal capital expenditure rises by 0.08 percent of initial assets through quarter 4 and 0.13 percent through quarter 8. Over the first year, the pattern is consistent with firms financing part of the working-capital burden through supplier credit and borrowing while maintaining investment spending. Ho et al. (2026) study a contraction in capital formation through the cost of new capital. Here, material prices enter through the goods firms produce. Nominal expenditure does not identify quantities: higher capital-goods prices can raise spending even when real investment falls.

Several outcomes move significantly in the opposite direction before the price-change window, so part of the subsequent adjustment may reflect reversal.[44]

[Insert Figure 6 about here]

The baseline estimates use material-price changes from any source, including movements in demand. Table 10 separates the pressure accumulated from $t-4$ to

[42] On these firm-quarters, operating income changes by -0.02 and -0.02 percent of initial assets at quarters 2 and 4, while operating cash flow falls by 0.13 and 0.19 percent. Operating income less the changes in inventories and receivables plus the change in payables gives -0.10 and -0.08. Each term comes from a separate regression of Equation (16) on the same firm-quarters.

[43] Appendix Table A.8 reports the additional estimates. Net financing reaches 0.37 percent through quarter 8, while the accounting interest rate on debt shows no significant response.

[44] The Before column estimates responses over $t-9$ to $t-5$, before the price-change window. The coefficients for the cost-of-sales ratio, operating cash flow, total debt, and short-term debt are significant at the 5 percent level, as are those for inventories, receivables, and payables, and the coefficient for net financing at the 10 percent level, with signs opposite to the later responses.

$t + h$ into a component associated with identified supply tightenings and the remaining price variation. For each material with a tightening event, we estimate a quarterly price model after excluding the announcement quarter and the next two quarters for every identified event, tightening and easing alike. Its seasonal terms and common commodity factor give the price path predicted in the absence of an event. Beginning in the announcement quarter, the event-driven component cumulates actual price changes net of this path, following the event-study counterfactual in MacKinlay (1997). The remainder combines pre-announcement movement with the part of the subsequent price change predicted by the model. Both components enter the regression. Figure 5 places each tightening event's two-quarter material-price change against the contemporaneous cross-material distribution. The changes vary widely and can be negative when an event arrives during a demand downturn or does not bind. The decomposition retains these signs.

[Insert Figure 5 about here]

Both components enter in percentage points of identified revenue, so their coefficients compare price changes with the same predicted cost increase. The responses are larger for the event-driven component. At quarter 4, operating cash flow falls by 0.39 percent of initial assets per percentage point, compared with 0.13 percent for the remaining variation. Net financing rises by 0.69 percent, compared with 0.09 percent. Cost-of-sales and long-term-debt responses show the same ordering.[45] Figure 7 traces the two responses through quarter 8 and includes the pre-window estimates. The event-driven component subtracts the fitted seasonal and common-factor movements, but other material-specific movements within an event window remain in it.[46]

[Insert Table 10 about here]

[Insert Figure 7 about here]

The aggregate estimates do not reveal whether specialist reports tend to cover the relationships with larger operating consequences. We therefore repeat the event decomposition at the firm–material level, with one observation for each firm, material, and quarter with positive priced dependence:

[45] Sales rise with both components, by 1.86 and 0.70 percent at quarter 4, while stock returns rise by 0.58 percent with the remainder and show no significant response to the event-driven component. Before the window, the event-component coefficients are insignificant for sales, the cost-of-sales ratio, cash flow, capital expenditure, short-term debt, and returns. Net financing and long-term debt are the exceptions: their pre-window coefficients have the opposite sign, as in the reversal pattern discussed above.

[46] A variant that drops the common-factor term for the ten investment metals, whose complex-wide movements an event itself can drive, gives similar estimates: the cash-flow and net-financing responses at quarter 4 are -0.37 and 0.58 per percentage point.

$$\begin{aligned}\Delta y_{f,t+h} = {} & \beta_h^E \, Pressure^E_{mf,t,h} + \beta_h^R \, Pressure^R_{mf,t,h} + \phi_h \, Pressure_{-m,f,t,h} \\ & + \mathbf{G}'_{f,t,h}\lambda_h + \gamma_1 D_{ft} + \gamma_2 \ln A_{f,t-1} \\ & + \alpha_{f,q(t)} + \alpha_{s,t} + \alpha_{c,t} + \alpha_{m,t} + \varepsilon_{mf,t+h},\end{aligned} \tag{17}$$

The superscripts E and R split material m's dependence-weighted price change from $t-4$ to $t+h$ into its event-driven part and the remainder, using the same construction as Table 10. The two terms enter in percentage points of identified revenue. One standard deviation of material-specific pressure is 0.47 percentage points. $Pressure_{-m,f,t,h}$ is the firm's full-window pressure from its other priced materials excluding gold and silver, while $\mathbf{G}_{f,t,h}$ collects the corresponding gold and silver terms, as in Equation (16). These full-window terms absorb the later-price controls. Material-by-quarter fixed effects $\alpha_{m,t}$ absorb the price change common to firms exposed to the same material, so β_h^E compares firms with different dependence on that material.

At this level, we consider a relationship documented if specialist reports name goods carrying at least half of the firm's dependence on the material. Some reports postdate the quarters they classify, so documentation measures coverage in the report corpus and need not match the information available to investors at the time. We compare the two groups in one regression, interacting every regressor and fixed effect in Equation (17) with the documentation indicator.

Table 11 reports the responses separately for documented and undocumented relationships. The operating-cash-flow point estimates are negative in both groups and are not statistically different from each other. Financing separates them more clearly. Net financing rises by 1.09 percent of initial assets per percentage point where the use is documented and is near zero otherwise, a difference significant at the 10 percent level. Short-term debt follows the same pattern in the point estimates. Specialist reports thus tend to document the relationships whose cost shocks firms finance.

[Insert Table 11 about here]

Firms' capacity to absorb a working-capital shock may vary with size. Table 12 compares the response to event-driven pressure for firms below and above the country–quarter median of total assets in U.S. dollars. At quarter 4, operating cash flow at small firms falls by 0.73 percent of initial assets, while the large-firm estimate is smaller and insignificant. Short-term debt also rises more at small firms. Its small-minus-large contrast is 0.42 percent of initial assets and significant at the 10 percent level. The other size differences remain suggestive.[47] The point estimates place more of the cash-flow burden and borrowing response with small firms. This firm-side pat-

[47]Net equity issuance, reported chiefly by U.S. filers, shows a positive but imprecise response to the event-driven part (the memo row of Table 13).

tern complements bank-side evidence that disruption-time liquidity flows through committed credit lines held mainly by large firms (Greenwald et al., 2025).

[Insert Table 12 about here]

The cash-flow shortfall can be met with external finance, lower investment, or cash drawdowns. Table 13 organizes these margins in a sources-and-uses budget at quarters 4 and 8, using a common sample at each horizon. At quarter 4, operating cash flow falls by 0.31 percent of initial assets while net external financing rises by 0.62 percent. Cash holdings do not respond significantly, and nominal capital expenditure rises by 0.17 percent. The pattern is consistent with firms using external finance to cover the cash-flow shortfall while maintaining investment. The smaller quarter-8 sample is less precise, and its confidence intervals include the quarter-4 estimates.[48]

[Insert Table 13 about here]

Table 14 converts the quarter-4 responses into dollars by applying them to firms' 2024–2025 exposures. In Panel A we impose a uniform 10 percent price increase, and in Panel B we replay the 28 events at each event's observed price change. Both use the same pooled response coefficients for documented and undocumented relationships. The uniform scenario implies $17.8 billion of operating cash flow at risk and $31.5 billion in external financing, or roughly $1.8 of external funds for each dollar of cash flow at risk. Replaying the events produces a similar $35.4 billion in financing. Notably, relationships absent from specialist reports account for 48 percent of cash flow at risk in the uniform scenario. This share exceeds 80 percent for the 2023 and 2025 controls on gallium, germanium, bismuth, indium, and tungsten, suggesting that market attention may overlook much of firms' exposure to major material restrictions.

[Insert Table 14 about here]

7.3 Material Dependence and Equity Risk

Firms within the same narrow industry can face different equity risk because the goods they produce expose them differently to input shocks. Research on operating flexibility shows that a firm's ability to adjust production changes how operating leverage translates into stock risk (Gu et al., 2018). Material dependence should characterize one such production-side exposure. This is also the prediction from Section 2.5, which suggests that return risk rises with both the amount of dependence and the price volatility of the materials behind it.

[48] A broader definition counts the full four-quarter price change as supply-driven when the event list contains a tightening for the material in the same quarter or the preceding one. It gives the same qualitative pattern. Appendix Table A.9 reports the estimates.

We focus first on total return volatility, idiosyncratic volatility, and downside risk measured by historical expected shortfall. We then hold dependence fixed and examine whether the materials in a firm's portfolio account for further differences in risk. We use market beta to assess how much of this risk is systematic. Bid-ask spreads and valuation address related market implications.

To compare firms under similar industry and country conditions in the same year, we estimate the following regression by ordinary least squares (OLS) for each market outcome,

$$Outcome_{ft} = \alpha_{s,t} + \alpha_{c,t} + \theta\, Rank_{ft} + Controls'_{ft}\gamma + \varepsilon_{ft}, \tag{18}$$

where $Rank_{ft}$ is the firm's dependence rank. The terms $\alpha_{s,t}$ and $\alpha_{c,t}$ are fixed effects for the firm's six-digit NAICS industry s and country c in year t. The coefficient θ, estimated separately for each outcome, compares outcomes at the bottom and top of the dependence ranking, controlling for standard firm characteristics.[49] We omit firm fixed effects to retain persistent differences across firms and cluster standard errors by firm.

Table 15 reports the estimates. The first three columns contain the main risk outcomes. Panels A and B include firms with positive as well as zero dependence, using the dependence rank and the rank of the firm's largest single-material dependence. Measured zeros sit at the bottom of each ranking. Panel C restricts the sample to positive dependence and uses the number of materials. The remaining columns report market beta, valuation, and bid-ask spreads.

[Insert Table 15 about here]

More dependent firms have more volatile returns and larger downside losses. Moving from the median to the 90th percentile of dependence ranks is associated with 1.79 percentage points higher annualized return volatility and 1.08 percentage points higher idiosyncratic volatility (Panel A, columns 1 and 2).[50] Expected shortfall is 0.15 percentage points higher over the same comparison (column 3). These differences are 4.8 percent of the sample median for return volatility and 2.7 percent of the median loss on the worst trading days. Ranking firms by their largest single-material dependence gives a similar pattern (Panel B), and the number of materials is positively associated with all three risk measures among firms with positive dependence (Panel C). The associations remain positive when we compute the ranks only among firms with positive dependence.[51]

[49] The controls are log assets, operating income/assets, capital expenditure/assets, net PP&E/assets, inventory/assets, the cost-of-sales ratio, and R&D/assets, with an indicator for unreported R&D.

[50] The corresponding dependence ranks are 0.00 and about 0.74 in this estimation sample, because more than half of the firm-years have measured zero dependence.

[51] Appendix Table A.7 re-estimates the dependence rank and largest-material rank within the positive-dependence sample and adds log total dependence. It also reports the beta, valuation, and

The composition of a firm's material portfolio provides a sharper test of this mechanism. Firms with the same total dependence can face shocks of different sizes because material prices differ in volatility. Under the model, return risk should rise with this material-price risk, and the association should become stronger as dependence increases. Limited substitution can further magnify the cost response to large shocks, much as limited operating flexibility magnifies stock risk (Gu et al., 2018).

We implement this comparison using the price-risk and China-restriction baskets defined in Section 5.3, Equation (14). They weight the material measures σ_m and ρ_m by the firm's dependence on each of the 42 priced materials other than gold and silver. The China-restriction basket asks whether firms exposed to materials subject to more Chinese restrictions also carry more equity risk. Studies of oil exposure combine industries' stock-return sensitivity to oil with oil-price volatility (Stein and Stone, 2013; Alfaro et al., 2024), while Klepacz (2026) uses spending to measure oil exposure. Our price-risk basket instead uses production-based cost dependence across 42 materials.

For firm-years with positive dependence on these materials, we estimate

$$\begin{aligned} Outcome_{ft} = \alpha_{s,t} + \alpha_{c,t} + \phi\, PriceRisk_{ft} + \psi\, ChinaRestriction_{ft} \\ + \theta\, Rank_{ft} + \delta\, Other_{ft} + Controls'_{ft}\gamma + \varepsilon_{ft}, \end{aligned} \tag{19}$$

where $Other_{ft}$ is the share of dependence on gold, silver, and unpriced materials. The remaining terms follow Equation (18). Both baskets are scaled by their standard deviations, so ϕ and ψ compare firms with the same dependence rank whose materials differ in risk. Table 16 reports the estimates. Panel A enters the two baskets linearly, and Panel B interacts each basket with the dependence rank to test whether the association strengthens with dependence. We use exposure-robust standard errors that treat each material as one draw (Adão et al., 2019), since firms share exposure to a concentrated set of materials.[52]

[Insert Table 16 about here]

Portfolio composition helps account for the risk differences among firms with similar dependence ranks. A price-risk basket one standard deviation higher, equivalent to 1.7 percentage points of monthly price volatility, is associated with 0.74 percentage points higher annualized return volatility, 0.50 points higher idiosyncratic volatility, and 0.06 points higher expected shortfall (Panel A). All three estimates are significant at the 1 percent level. The China-restriction basket is also associated with higher return and idiosyncratic volatility, by 0.61 and 0.39 points, but has no significant as-

bid-ask-spread checks discussed below.

[52]The inverse Herfindahl index of these exposures is about 10, and copper accounts for 23 percent of aggregate dependence. Firm clustering alone does not account for common material exposure across firms.

sociation with expected shortfall. Dependence retains its associations with volatility and expected shortfall when both baskets enter.

The association between material-price risk and return outcomes strengthens with dependence (Panel B). From the bottom to the top of the dependence ranking, it increases by 2.19 percentage points for return volatility, 1.40 points for idiosyncratic volatility, and 0.19 points for expected shortfall. At the top, a price-risk basket one standard deviation higher is associated with 1.86 points higher return volatility.[53] These interactions support the model's prediction that dependence amplifies exposure to material-price risk. The China-restriction interactions are not statistically significant. The basket records restrictions, while the model's physical-rationing mechanism concerns realized input shortages, making this a less direct test.

The evidence is stronger for total and downside risk than for systematic risk. In Table 15, the median-to-90th-percentile move in dependence corresponds to a beta about 0.03 higher, and both rank coefficients are significant at the 5 percent level. The number of materials is unrelated to beta. Within the positive-dependence sample, the rank and log-dependence coefficients remain positive but are not statistically significant. In Table 16, the price-risk basket is associated with a beta about 0.02 higher, while the China-restriction estimate is not significant. Outside the investment-metal group, material prices also comove only weakly with world equity returns in the presample data (Section 5.4). This pattern is consistent with much of material-price risk being distinct from aggregate market movements. Input prices that fall in downturns can offset firms' other risks (Kogan et al., 2023), whereas material shocks contribute to systematic production-network risk when they comove with aggregate conditions (Herskovic, 2018).

The return-risk differences also appear in trading costs. Bid-ask spreads are wider under all three dependence measures, and each association remains significant at the 1 percent level in the positive-dependence sample. The price-risk and China-restriction baskets are associated with spreads that are 0.96 and 0.79 basis points wider. At the top of the dependence ranking, a one-standard-deviation higher price-risk basket is associated with a 2.18-basis-point wider spread. Greater return volatility may increase the risk borne by liquidity providers.

Valuation presents a different pattern. In Table 15, firms at the 90th percentile of dependence have market-to-book ratios 8.7 percent higher than firms at the median, conditional on the controls and fixed effects. The largest-material rank carries a similar association, and valuation rises with the number of materials among exposed firms. When ranks and log total dependence are estimated within the positive-dependence sample, however, their coefficients are positive but not significant. The

[53] At the bottom of the ranking, the estimated associations with expected shortfall and spreads are close to zero and that with return volatility is slightly negative.

baseline valuation difference is therefore clearest between firms with and without measured dependence.

Neither material-risk basket has a significant average association with market-to-book. Within exposed firms, however, the valuation association of material-price risk changes sign with dependence. In Table 16, Panel B, a one-standard-deviation higher price-risk basket is associated with 3.1 log points lower market-to-book at the bottom of the dependence ranking and 6.9 log points higher market-to-book at the top. The 10.0-log-point difference is significant at the 1 percent level. More volatile materials therefore accompany lower valuations among lightly dependent firms but higher valuations among the most dependent firms.

Higher risk alone cannot explain higher valuation: for given expected cash flows, a higher required return would lower it. The positive valuation association could reflect growth opportunities in material-dependent products, consistent with the higher investment and R&D spending of more dependent firms (Table 8). It could also reflect the learning mechanism in Pástor and Veronesi (2003), where uncertainty about average profitability can raise market-to-book ratios, especially for firms that pay no dividends. Our estimates do not distinguish between these interpretations. For the risk outcomes, the mechanism is more specific: material-price volatility is associated with more stock risk precisely where dependence is high.

8 Conclusion

A firm need not purchase a raw material directly to depend on it. The requirement may be embodied in the components and chemicals it buys. We trace these connections across 55 materials and 12,481 firms. Technical evidence on physical production, combined with input cost shares and product-segment revenues, yields a firm-level measure of structural material dependence and the multi-stage production paths behind it.

These connections often lie far downstream. Shared intermediate inputs several stages beyond refining carry much of measured dependence. Firms in the same narrow industry can therefore face very different material requirements, including outside prominent applications in green energy and semiconductors.

Public reporting reveals only part of this dependence. Firms disclose materials more often when they represent a larger share of costs and are harder to substitute. Both filings and specialist reports nevertheless leave most measured dependence undocumented in their respective comparison samples. Minerals and ceramics, among the largest product groups, remain largely undocumented in both sources, illustrating what a focus on headline technologies or self-reported use can miss.

Material-price changes reveal the operating consequences of this dependence. Greater predicted cost pressure is associated with lower operating cash flow and more external financing. Through the first year, these patterns are consistent with firms financing a working-capital burden while maintaining capital expenditure, and operating income shows no significant response. The cash-flow and financing responses are larger when the price rise coincides with a supply tightening on the material, such as an export restriction or a mine disruption.

These within-industry differences extend to financial outcomes. After controlling for standard firm characteristics and industry-by-year effects, more dependent downstream firms exhibit higher stock-return volatility, greater downside tail risk, and higher market-to-book valuations than their industry peers. The return-risk differences are modest, while the valuation differences are larger and clearest between firms with and without measured dependence. The return-risk estimates also vary with firms' material portfolios: their association with price volatility is stronger among firms with greater dependence. Together with the cost response, these patterns support the economic content of the measure.

Our findings speak to how material risk should be assessed. Assessments based only on supplier relationships or corporate disclosures will miss a large part of this exposure. Most material requirements enter through purchased components and chemicals, so changing direct suppliers may leave the upstream bottleneck unchanged. Disclosure-based measures also omit relationships firms do not report. Mapping the physical requirements of firms' products can identify these exposures before a disruption makes them visible in market prices.

Policy assessments face the same problem. Their usual focus on extraction and prominent clean-energy or defense applications overlooks the points at which different material chains converge. Industrial chemicals, catalysts, and specialty alloys, for example, are shared intermediate inputs across the wider economy. Locating these bottlenecks can help resilience policies and supply-chain stress tests target the relevant production links.

References

Adam, C., Arezki, R., and van der Ploeg, F. (2026). The scramble for critical minerals. *Oxford Review of Economic Policy*, 42(1):1–13.

Adão, R., Kolesár, M., and Morales, E. (2019). Shift-share designs: Theory and inference. *Quarterly Journal of Economics*, 134(4):1949–2010.

Alfaro, I., Bloom, N., and Lin, X. (2024). The finance uncertainty multiplier. *Journal of Political Economy*, 132(2):577–615.

Alfaro, L., Fadinger, H., Schymik, J., and Virananda, G. (2025). Trade and industrial policy in supply chains: Directed technological change in rare earths. Working Paper 33877, National Bureau of Economic Research. Revised March 2026.

Alfaro, L., Fadinger, H., Schymik, J., and Virananda, G. (2026). The rare-earth footprint of the global economy: Input-output evidence. *AEA Papers and Proceedings*, 116:103–108.

Alonso, E., Pineault, D. G., Gambogi, J., and Nassar, N. T. (2023a). Mapping first to final uses for rare earth elements, globally and in the united states. *Journal of Industrial Ecology*, 27(1):312–322.

Alonso, E., Pineault, D. G., and Nassar, N. T. (2023b). Streamlined approach for assessing embedded consumption of lithium and cobalt in the united states. *Journal of Industrial Ecology*, 27(1):33–42.

Antràs, P., Chor, D., Fally, T., and Hillberry, R. (2012). Measuring the upstreamness of production and trade flows. *American Economic Review*, 102(3):412–416.

Atalay, E. (2017). How important are sectoral shocks? *American Economic Journal: Macroeconomics*, 9(4):254–280.

Atilgan, Y., Bali, T. G., Demirtas, K. O., and Gunaydin, A. D. (2020). Left-tail momentum: Underreaction to bad news, costly arbitrage and equity returns. *Journal of Financial Economics*, 135(3):725–753.

Atkeson, A. and Burstein, A. (2008). Pricing-to-market, trade costs, and international relative prices. *American Economic Review*, 98(5):1998–2031.

Bachmann, R., Baqaee, D., Bayer, C., Kuhn, M., Löschel, A., Moll, B., Peichl, A., Pittel, K., and Schularick, M. (2024). What if? The macroeconomic and distributional effects for Germany of a stop of energy imports from Russia. *Economica*, 91(364):1157–1200.

Baldwin, R., Freeman, R., and Theodorakopoulos, A. (2023). Hidden exposure: Measuring US supply chain reliance. *Brookings Papers on Economic Activity*, 54(2):79–134.

Baqaee, D. R. and Farhi, E. (2019). The macroeconomic impact of microeconomic shocks: Beyond Hulten's theorem. *Econometrica*, 87(4):1155–1203.

Baqaee, D. R. and Farhi, E. (2020). Productivity and misallocation in general equilibrium. *Quarterly Journal of Economics*, 135(1):105–163.

Barrot, J.-N. and Sauvagnat, J. (2016). Input specificity and the propagation of idiosyncratic shocks in production networks. *Quarterly Journal of Economics*, 131(3):1543–1592.

Bartram, S. M., Brown, G. W., and Conrad, J. (2011). The effects of derivatives on firm risk and value. *Journal of Financial and Quantitative Analysis*, 46(4):967–999.

Bastianin, A., Castelnovo, P., Frattini, F. F., and Vona, F. (2026). Induced innovation in

critical mineral saving technologies. Working Paper 05.2026, Fondazione Eni Enrico Mattei.

Boehm, C. E., Flaaen, A., and Pandalai-Nayar, N. (2019). Input linkages and the transmission of shocks: Firm-level evidence from the 2011 Tōhoku earthquake. *Review of Economics and Statistics*, 101(1):60–75.

Boer, L., Pescatori, A., and Stuermer, M. (2024). Energy transition metals: Bottleneck for net-zero emissions? *Journal of the European Economic Association*, 22(1):200–229.

Bogmans, C., Jerez-Osses, M., Miranda-Pinto, J., and Natal, J.-M. (2026). Macroeconomic effects of rare earths supply chain disruptions. Working Paper 2026/073, International Monetary Fund.

Bolton, P. and Kacperczyk, M. (2021). Do investors care about carbon risk? *Journal of Financial Economics*, 142(2):517–549.

Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the Great East Japan Earthquake. *Quarterly Journal of Economics*, 136(2):1255–1321.

Cohen, L. and Frazzini, A. (2008). Economic links and predictable returns. *Journal of Finance*, 63(4):1977–2011.

Corwin, S. A. and Schultz, P. (2012). A simple way to estimate bid-ask spreads from daily high and low prices. *Journal of Finance*, 67(2):719–760.

Dagdelen, J., Dunn, A., Lee, S., Walker, N., Rosen, A. S., Ceder, G., Persson, K. A., and Jain, A. (2024). Structured information extraction from scientific text with large language models. *Nature Communications*, 15:1418.

De Loecker, J. and Eeckhout, J. (2018). Global market power. Working Paper 24768, National Bureau of Economic Research.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, 135(2):561–644.

Dew-Becker, I. (2023). Tail risk in production networks. *Econometrica*, 91(6):2089–2123.

Ersahin, N., Giannetti, M., and Huang, R. (2024). Supply chain risk: Changes in supplier composition and vertical integration. *Journal of International Economics*, 147:103854.

European Commission (2017). On the 2017 list of critical raw materials for the EU. COM(2017) 490 final.

European Parliament and Council of the European Union (2024). Regulation (EU) 2024/1252 establishing a framework for ensuring a secure and sustainable supply of critical raw materials. Official Journal of the European Union.

Ferreira, M. A. and Laux, P. A. (2007). Corporate governance, idiosyncratic risk, and information flow. *Journal of Finance*, 62(2):951–989.

Fetzer, T., Lambert, P. J., Feld, B., and Garg, P. (2024). AI-generated production net-

works: Measurement and applications to global trade. Warwick Economics Research Paper 1528, University of Warwick.

Florackis, C., Louca, C., Michaely, R., and Weber, M. (2023). Cybersecurity risk. *Review of Financial Studies*, 36(1):351–407.

Ganapati, S., Shapiro, J. S., and Walker, R. (2020). Energy cost pass-through in U.S. manufacturing: Estimates and implications for carbon taxes. *American Economic Journal: Applied Economics*, 12(2):303–342.

Gentzkow, M., Kelly, B., and Taddy, M. (2019). Text as data. *Journal of Economic Literature*, 57(3):535–574.

Graedel, T. E., Harper, E. M., Nassar, N. T., Nuss, P., and Reck, B. K. (2015). Criticality of metals and metalloids. *Proceedings of the National Academy of Sciences*, 112(14):4257–4262.

Greenwald, D. L., Krainer, J., and Paul, P. (2025). The credit line channel. *Journal of Finance*, 80(6):3137–3183.

Gu, L., Hackbarth, D., and Johnson, T. (2018). Inflexibility and stock returns. *Review of Financial Studies*, 31(1):278–321.

Hassan, T. A., Hollander, S., van Lent, L., and Tahoun, A. (2019). Firm-level political risk: Measurement and effects. *Quarterly Journal of Economics*, 134(4):2135–2202.

Herskovic, B. (2018). Networks in production: Asset pricing implications. *Journal of Finance*, 73(4):1785–1818.

Ho, P., Knittel, C., and Concordel, A. (2026). The differential impacts of critical mineral prices and oil prices on the economy. *Oxford Review of Economic Policy*, 42(1):91–110.

Hoberg, G. and Phillips, G. (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy*, 124(5):1423–1465.

International Energy Agency (2024). Global critical minerals outlook 2024. Technical report, International Energy Agency, Paris.

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

Klepacz, M. (2026). Price setting and volatility: Evidence from oil price volatility shocks. *American Economic Journal: Macroeconomics*, 18(4):163–199.

Klimek, P., Baum, S., Gerschberger, M., and Hess, M. (2025). Systemic trade risk suppresses comparative advantage in rare earth dependent industries. arXiv preprint 2508.00556.

Kogan, L., Li, J., and Zhang, H. H. (2023). Operating hedge and gross profitability premium. *Journal of Finance*, 78(6):3387–3422.

Loughran, T. and McDonald, B. (2016). Textual analysis in accounting and finance: A survey. *Journal of Accounting Research*, 54(4):1187–1230.

MacKinlay, A. C. (1997). Event studies in economics and finance. *Journal of Economic Literature*, 35(1):13–39.

Menzly, L. and Ozbas, O. (2010). Market segmentation and cross-predictability of returns. *Journal of Finance*, 65(4):1555–1580.

Miranda-Pinto, J., Pescatori, A., Stuermer, M., and Wang, X. (2024). Beyond energy: Inflationary effects of metals price shocks in production networks. IMF Working Paper 24/215, International Monetary Fund.

Müller, M. A., Schweizer, D., and Seiler, V. (2016). Wealth effects of rare earth prices and china's rare earth elements policy. *Journal of Business Ethics*, 138(4):627–648.

Nassar, N. T., Brainard, J., Gulley, A. L., Manley, R., Matos, G. R., Lederer, G. W., Bird, L., Pineault, D. G., Alonso, E., Gambogi, J., and Fortier, S. M. (2020). Evaluating the mineral commodity supply risk of the U.S. manufacturing sector. *Science Advances*, 6(8):eaay8647.

Nassar, N. T., Pineault, D., Allen, S. M., McCaffrey, D. M., Padilla, A. J., Brainard, J. L., Bayani, M., Shojaeddini, E., Ryter, J. W., Lincoln, S., and Alonso, E. (2025). Methodology and technical input for the 2025 U.S. list of critical minerals—assessing the potential effects of mineral commodity supply chain disruptions on the U.S. economy. Open-File Report 2025–1047, U.S. Geological Survey. Version 2.0, revised April 2026.

Nguyen-Tien, V. (2026). Input uncertainty and firm performance: Evidence from critical minerals. Discussion Paper 2197, Centre for Economic Performance, London School of Economics.

Ostry, J. (2026). Chokepoint: The financial effects of rare earth supply disruptions. Working paper, Geneva Graduate Institute, October 22.

Panon, L., Lebastard, L., Mancini, M., Borin, A., Tofani, P. C., Cariola, G., Essers, D., Gentili, E., Linarello, A., Padellini, T., Requena, F., and Timini, J. (2026). Inputs in distress: Geoeconomic fragmentation and firms' sourcing. *European Economic Review*, 187:105355.

Pástor, L. and Veronesi, P. (2003). Stock valuation and learning about profitability. *The Journal of Finance*, 58(5):1749–1789.

Peter, A. and Ruane, C. (2025). The aggregate importance of intermediate input substitutability. NBER Working Paper 31233, National Bureau of Economic Research. Revised May 2025.

Pierce, J. R. and Schott, P. K. (2012). A concordance between ten-digit U.S. harmonized system codes and SIC/NAICS product classes and industries. *Journal of Economic and Social Measurement*, 37(1–2):61–96.

Sautner, Z., van Lent, L., Vilkov, G., and Zhang, R. (2023). Firm-level climate change exposure. *Journal of Finance*, 78(3):1449–1498.

Stein, L. C. D. and Stone, E. C. (2013). The effect of uncertainty on investment, hiring, and R&D: Causal evidence from equity options. Working Paper 1649108, SSRN.

Sun, X., Hao, H., Zhao, F., and Liu, Z. (2018). Global lithium flow 1994–2015: Implications for improving resource efficiency and security. *Environmental Science & Technology*, 52(5):2827–2834.

Teulings, C. N. and Zubanov, N. (2014). Is economic recovery a myth? Robust estimation of impulse responses. *Journal of Applied Econometrics*, 29(3):497–514.

U.S. Geological Survey (2025a). Final 2025 list of critical minerals. Federal Register, 90 FR 50494.

U.S. Geological Survey (2025b). Mineral commodity summaries 2025. Technical report, U.S. Geological Survey.

Verrecchia, R. E. (1983). Discretionary disclosure. *Journal of Accounting and Economics*, 5:179–194.

Waugh, M. E. (2026). Trade in AI-related products. NBER Working Paper 35053, National Bureau of Economic Research.

Table 1: Production Dependence and Reported Material Relationships

Benchmark	Reported	Measured	Both	Recovery (%)	Reverse recovery (%)
Specialist reports	80	6,365	69	86.3	1.1
Filings, positive exposure	2,815	17,472	1,513	53.7	8.7
Filings, network presence	2,815	72,288	2,109	74.9	2.9

Notes: This table compares relationships in specialist reports and corporate filings with the production networks. Reported is the number stated by each benchmark, Measured the number identified by the network, and Both their overlap. Recovery equals Both/Reported and reverse recovery Both/Measured. Specialist reports are matched at the material–HS6 level for 39 materials. Filings are matched at the firm-year–material level for 55 materials in 7,666 Forms 10-K, 20-F and 40-F over fiscal 2017–2025. Positive exposure means positive cost dependence or producer-side exposure. Network presence means that a measured product line maps to a good in the material's network. Appendix Tables E.2 and E.4 report the underlying comparisons.

Table 2: Summary Statistics

	Observations	Mean	SD	P10	Median	P90
Panel A. Material dependence, full-information sample						
Dependence rank (0–1), downstream firms	74,087	0.187	0.299	0.000	0.000	0.732
Total dependence × 100, positive	29,039	11.50	20.64	0.02	2.22	40.39
Number of materials, positive	29,039	7.0	7.7	1.0	4.0	16.0
Panel B. Market outcomes						
Log market-to-book	56,230	0.75	0.94	−0.44	0.76	1.91
Return volatility (%, annualized)	59,940	41.4	22.5	18.4	36.9	69.1
Market beta	53,190	0.92	0.49	0.30	0.92	1.54
Idiosyncratic volatility (%, annualized)	58,175	39.6	17.8	20.8	36.6	61.4
Expected shortfall (%, daily)	61,372	5.85	2.39	3.17	5.56	8.67
Bid-ask spread (bp)	59,767	92.2	45.4	43.4	86.0	142.5
Panel C. Firm controls						
Log assets (local currency)	73,241	22.35	2.47	19.59	22.07	25.73
Operating income/assets (%)	73,193	3.55	12.45	−7.21	4.51	15.09
Capital expenditure/assets (%)	72,593	4.40	4.80	0.31	2.81	10.67
Net PP&E/assets (%)	73,127	25.07	19.56	2.88	21.09	53.18
Inventory/assets (%)	71,511	11.68	11.41	0.10	9.29	25.68
Cost-of-sales ratio (%)	72,732	63.2	23.5	28.8	67.1	88.5
R&D/assets (%)	75,510	2.80	4.39	0.00	1.46	6.91
R&D not reported (%)	75,510	27.1				
Positive dependence (%)	75,510	38.5				

Notes: This table presents summary statistics for the annual sample. Panel A uses the full-information sample, including upstream firms and measured zero dependence. Ranks are for downstream firms. Panels B and C report annual variables where available. Total dependence is multiplied by 100. Log assets enters relative to its country–year mean. Indicator rows report means. SD denotes standard deviation, P10 and P90 percentiles, pp percentage points, and bp basis points. Table 3 summarizes the quarterly sample. Appendix Table A.1 defines the variables.

Table 3: Summary Statistics: Material Prices and Quarterly Outcomes

	Observations	Mean	SD	P10	Median	P90
Panel A. Material price changes (%, material–quarters)						
All 42 priced materials, four-quarter	1,302	3.1	37.9	−37.1	−1.2	51.8
All 42 priced materials, two-quarter	1,302	1.5	24.1	−24.1	−0.1	30.8
Supply shocks (two-quarter, bracketing an event)	39	3.9	29.7	−24.7	3.7	26.7
Panel B. Cost pressure (% of identified revenue, firm–quarters)						
Cost pressure	158,068	0.062	1.188	−0.129	0.000	0.201
conditional on exposure	61,797	0.158	1.896	−1.043	0.000	1.414
conditional, in absolute value	61,797	0.908	1.672	0.002	0.175	3.001
Broad-window component, conditional on exposure	10,325	0.075	0.424	−0.007	0.001	0.156
conditional, in absolute value	10,325	0.121	0.414	0.000	0.004	0.243
Panel C. Firm outcomes through quarter 4						
Sales (log change, %)	125,272	7.55	48.37	−44.20	8.19	59.01
Cost-of-sales ratio (change, pp)	123,852	0.35	21.27	−18.82	0.09	20.11
Inventories (change)	120,026	1.37	5.68	−3.33	0.28	7.07
Accounts receivable (change)	125,496	1.89	7.52	−5.02	0.79	9.95
Accounts payable (change)	123,235	1.45	5.82	−3.54	0.42	7.49
Operating income (cumulative)	93,979	5.81	13.04	−7.23	5.66	19.45
Net income (cumulative)	123,772	3.30	14.02	−10.63	4.24	16.61
Operating cash flow (cumulative)	122,119	7.73	13.10	−5.39	7.56	22.04
Capital expenditure (cumulative)	108,546	6.62	7.39	0.65	4.18	15.67
Cash (change)	126,393	2.71	16.00	−10.48	0.38	16.05
Net financing (cumulative)	116,542	2.07	19.30	−12.10	−1.64	18.25
Total debt (change)	118,896	2.70	11.37	−7.11	0.38	14.65
Short-term debt (change)	114,892	1.15	7.45	−5.72	0.12	9.00
Long-term debt (change)	107,383	1.55	8.91	−5.59	0.01	10.27
Interest rate on debt (change, pp)	106,635	−0.16	9.42	−3.54	−0.05	3.40
Stock return (log change, %)	123,828	1.50	48.68	−54.04	0.46	60.33

Notes: This table summarizes material price changes, cost pressure, and quarterly firm outcomes. The quarterly sample contains 158,068 firm-quarters over 2018Q2–2025Q4. Panel A reports four-quarter log price changes, multiplied by 100, for the 42 priced materials other than gold and silver. The unit is a material-quarter. A broad-window supply-driven observation has a tightening event in the current or previous quarter. Panel B reports cost pressure and this broad-window component, which is distinct from the event-driven decomposition in Section 7.2. The exposed rows retain 61,797 firm-quarters with positive dependence on at least one priced material, while the other 96,271 observations have zero cost pressure. Absolute pressure is the magnitude of the signed measure. The broad-window rows condition on a nonzero component, leaving 10,325 firm-quarters at 2,303 firms. Cost pressure is winsorized at the 1st and 99th percentiles, and the broad-window component on nonzero values. Panel C reports each outcome on its quarter-4 regression sample. Dollar outcomes are in percent of assets at $t-1$ unless marked. Appendix Table A.1 defines the variables.

Table 4: Industry Profiles of Material Dependence

Panel A. Industry profiles, 2019 users

		Total dependence			Largest single material		
Industry (NAICS)	Firms	Level × 100	Nodes	Path length	Level × 100	Nodes	Path length
Semiconductors and components (3344)†	603	9.41	9.5	5.3	7.56	1.2	5.0
Other electrical equipment (3359)	190	15.52***	8.0	5.7***	12.19***	1.6***	5.5***
Basic chemicals (3251)	198	6.72	2.9	5.1	4.97	0.9	5.2
Electrical equipment (3353)	107	6.26	10.3	5.4	5.05	2.1***	5.2*
Motor vehicle parts (3363)	242	5.60	3.7	6.2***	3.99	0.9	6.1***
Communications equipment (3342)	147	3.47	12.7**	6.1***	1.83	1.0	6.6***
Other general-purpose machinery (3339)	136	1.92	5.3	5.6*	1.20	0.7	5.3**
Computers and peripherals (3341)	93	1.32	22.1***	6.1***	0.62	1.1	5.6***
Measuring and medical instruments (3345)	294	1.01	3.2	6.1***	0.73	0.5	6.1***
Industrial machinery (3332)	182	0.76	3.5	5.7**	0.53	0.5	5.4**
Wholesale and retail trade (42, 44–45)	138	7.21	2.6	5.1	5.54	0.7	5.2
Other services	1,610	0.20	0.8	5.8***	0.14	0.1	5.8***
Software publishers (5132)	356	0.13	2.0	6.2***	0.07	0.1	5.9***
All users	7,761	3.47	3.6	5.4	2.69	0.6	5.3

Panel B. Industry profiles, 2019 upstream firms

		Total dependence × 100			Largest single material × 100		
Industry (NAICS)	Firms	Cost (D)	Revenue (P)	Net ($D-P$)	Cost (D)	Revenue (P)	Net ($D-P$)
Metal ore mining (2122)	58	47.85	75.60	−27.75	36.41	58.79	−22.38
Nonferrous metals (3314)	30	75.53	70.17	5.35	57.56	56.60	0.95
Alumina and aluminum (3313)	11	73.06	66.96	6.10	68.68	65.00	3.69
Chemicals and fertilizers (3251, 3253)	12	58.63	37.64	20.99	42.31	28.01	14.30
Other	43	52.53	52.51	0.02	34.10	38.20	−4.10
All upstream firms	154	57.19	64.52	−7.33	42.65	50.66	−8.01

Panel C. Variation in dependence ranks explained by fixed effects

	Firm-years			Firm–material–years		
Fixed effects	Cells	R^2	Adj. R^2	Cells	R^2	Adj. R^2
Year	9	0.000	0.000	495	0.029	0.029
Country × year	718	0.053	0.044	39,490	0.045	0.035
Two-digit industry × year	214	0.238	0.235	11,770	0.085	0.083
Four-digit industry × year	2,253	0.454	0.437	123,915	0.226	0.201
Six-digit industry × year	6,011	0.561	0.521	330,605	0.342	0.282
Country × year + six-digit industry × year	6,728	0.565	0.520	370,094	0.347	0.279
Firm	11,934	0.972	0.966	656,370	0.959	0.951

Notes: This table reports industry profiles and sources of variation in dependence. Panel A gives 2019 means for users, including measured zeros, in the listed primary four-digit NAICS industries. Users without a primary industry label (155 firms in 2019) are excluded. Level is dependence multiplied by 100. Nodes counts the material–product nodes carrying quantified coefficients, and path length is the dependence-weighted number of stages from materials to products. Stars compare means with the semiconductor benchmark (†) in one-sided tests with heteroskedasticity-robust standard errors. Panel B reports 2019 upstream-firm means, where D is dependence and P the producer share. Panel C reports R^2 from separate fixed-effect regressions on 72,194 user firm-years. Firm–material ranks are computed within material and their fixed effects are interacted with material. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 5: Refined Intermediates Carrying Material Dependence

Material	Firms	Bottleneck share (%)	Listed share (%)	Steps	Listed intermediate
Palladium	721	100	94	5	Refined palladium metal sponge
Yttrium	328	100	100	7	Yttrium oxalate precipitate
Terbium	168	100	97	8	Terbium oxalate precipitate
Iridium	49	100	59	5	Refined iridium metal
Platinum	397	100	95	5	Refined platinum sponge
Scandium	32	100	98	6	Sc/Zr-modified 6082 aluminum alloy
Erbium	21	100	98	7	Erbium oxalate precipitate
Gallium	420	100	99	2	4N gallium
Indium	498	100	68	7	Indium tin oxide powder and sputtering targets
Samarium	141	100	100	6	Samarium(III) oxalate precipitate
Ytterbium	83	100	100	7	Ytterbium oxalate precipitate
Lanthanum	157	100	100	6	Lanthanum oxalate precipitate
Holmium	67	100	100	8	Holmium oxalate precipitate
Hafnium	457	100	80	4	Hafnium tetrachloride
Ruthenium	193	100	69	5	RuO2 thick-film resistor paste
Germanium	255	100	84	7	Doped germanium single-crystal boules
Thulium	2	100	100	9	Thulium oxalate precipitate
Praseodymium	323	100	100	6	Praseodymium oxalate precipitate
Europium	19	100	100	6	Europium(III) oxalate precipitate
Gadolinium	185	100	100	7	Gadolinium oxalate precipitate
Lutetium	93	100	100	7	Lutetium oxalate precipitate
Dysprosium	256	100	100	6	Purified dysprosium aqueous solution
Neodymium	312	99	59	7	Neodymium oxide
Cerium	297	98	98	6	Pure cerium solution
Aluminum	1,483	96	86	5	Primary aluminum metal
Lithium	417	89	78	3	Battery-grade lithium carbonate
Silver	919	87	87	1	Electrolytic fine silver
Fluorspar	385	81	57	4	Dried acid-grade fluorspar concentrate
Beryllium	105	80	80	5	High-frequency hybrid circuits on beryllia substrates
Bismuth	176	78	78	2	Refined bismuth metal
Copper	1,556	78	78	2	Cathode copper
Rhodium	87	77	53	6	Soluble rhodium chloride RhCl3·2.5H2O
Gold	759	76	76	2	Wohlwill fine gold
Lead	513	75	75	3	Refined lead ingot
Tin	743	70	70	2	Refined tin
Bauxite	376	55	55	3	Precipitated Bayer hydrate
Phosphate rock	168	52	52	3	Merchant-grade wet-process phosphoric acid

Notes: This table identifies refined intermediates that carry at least half of 2019 firm-level material dependence. A good carries dependence when every recorded route from a material entry point to a downstream product passes through it. Bottleneck share is the largest firm-level dependence share carried by one non-entry good. The table lists the 37 materials meeting the threshold and reports the furthest-downstream qualifying intermediate. Steps is its shortest path from an entry point. The names abbreviate the network descriptions.

Table 6: Dependence and Corporate Disclosure

Panel A. Raw disclosure rates

Dependence group	Firm–material years	Disclosed	Percent
Upstream firms			
Zero	3,318	37	1.12
Positive	251	146	58.17
Downstream firms			
Zero	401,740	1,575	0.39
Quartile 1 (lowest)	4,235	113	2.67
Quartile 2	4,452	139	3.12
Quartile 3	4,331	381	8.80
Quartile 4 (highest)	4,610	921	19.98

Panel B. Disclosure regressions

	Dependent variable: Disclosure, $Disclose_{mft}$ (0/1)				
	All pairs	Downstream firms, positive dependence			
Coefficients × 100	(1)	(2)	(3)	(4)	(5)
Dependence quartile 1 (lowest)	1.96***				
	(0.55)				
Dependence quartile 2	2.83***				
	(0.65)				
Dependence quartile 3	8.17***				
	(1.01)				
Dependence quartile 4 (highest)	18.24***				
	(1.55)				
Upstream firms, positive	56.94***				
	(6.15)				
Material share (10 pp)		3.08***	3.06***	3.09***	3.07***
		(0.45)	(0.45)	(0.58)	(0.59)
Hard-to-substitute share			3.38**		3.39
			(1.68)		(2.42)
Export restrictions (1 SD)				0.50**	0.51**
				(0.24)	(0.24)
Firm × year FE	Yes	Yes	Yes	Yes	Yes
Material × year FE	Yes	Yes	Yes	–	–
Material FE	–	–	–	Yes	Yes
Clustering	Firm	Firm	Firm	Two-way	Two-way
Observations	422,937	17,160	17,151	17,115	17,078
Firms	1,505	311	311	310	310
Materials	55	53	53	55	53
Within R^2	0.076	0.047	0.048	0.048	0.049

Notes: This table relates disclosure to product-based dependence for 7,691 firm-years and 55 materials. $Disclose_{mft}$ equals one when the filing affirms a current operational relationship with the material. Panel A reports disclosure rates by dependence group. Panel B reports linear probability estimates in percentage points. Column (1) estimates Equation (15). Columns (2)–(5) restrict the sample to downstream firms with positive dependence. Material share is the material's share of the firm's total dependence. Columns (4) and (5) use the prior year's export-restriction count, scaled by its standard deviation of 2.1. Because the measure uses these filings when they name a material, part of its association with disclosure is built in. Standard errors are clustered by firm in columns (1)–(3) and by firm and material in columns (4)–(5). ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 7: Material Dependence Covered by Public Discussion

	Regulatory filings (45 materials)		Specialist reports (30 materials)	
	Share of measured	Percent disclosed	Share of measured	Percent named
All downstream relationships	100.0	38.3	100.0	41.0
Panel A. Firm industry				
Chemicals	15.4	71.2	17.8	43.0
Mining and primary metals	19.0	5.9	13.1	21.7
Nonmetallic minerals	0.7	24.8	11.9	11.3
Computers and electronics	18.0	61.9	20.6	48.8
Electrical equipment	13.2	77.2	13.1	78.9
Machinery and vehicles	14.0	6.3	12.6	25.8
Other manufacturing	12.2	20.9	8.0	54.3
Trade and services	7.5	18.4	2.9	36.4
Panel B. Product type of the good carrying the dependence (HS chapters)				
Food and feed (01–24)	0.0	0.0	0.0	0.0
Minerals and ceramics (25–27, 68–70)	22.8	7.8	15.1	2.1
Inorganic chemicals (28)	14.6	96.1	26.4	54.6
Pharmaceuticals (30)	0.1	0.0	0.0	0.0
Other chemicals and plastics (29, 31–40)	16.0	62.1	6.8	41.7
Metals and metal articles (71–83)	4.8	29.8	18.6	14.8
Machinery (84)	2.8	0.5	1.6	16.3
Electrical and electronic equipment (85)	19.7	41.2	28.3	68.3
Vehicles and instruments (86–92)	17.2	8.8	2.9	37.2
Other goods	1.9	76.8	0.3	0.0

Notes: This table reports how much measured dependence appears in regulatory filings and specialist reports. Both comparisons use fiscal 2025, users in the full-information sample, and materials outside the investment metals. The filing columns cover 1,480 positive relationships of 248 firms across 45 materials. The report columns cover 2,696 firms and the 30 materials the reports mention. Share of measured is the row's share of dependence in the comparison sample. Percent disclosed (named) is the dependence-weighted coverage rate. The two sources cover different samples: filings cover firm–material relationships, whereas reports name individual goods. The measure also uses filings that name a material, so part of filing coverage is built in. Panel A groups firms by industry and Panel B goods by HS chapter. Firms receive equal weight.

Table 8: Firm Characteristics and Material Dependence

	Dependent variable: Dependence rank (0–1)				
	Standard characteristics (1)	+ Production characteristics (2)	+ Market scope (3)	+ Industry & country FE (4)	Firm FE (5)
Log assets	0.104***	0.069***	0.060***	0.087***	0.058***
	(0.010)	(0.012)	(0.015)	(0.012)	(0.016)
Operating income/assets	−0.036***	0.022**	0.025**	0.001	0.010**
	(0.009)	(0.009)	(0.012)	(0.010)	(0.004)
Debt/assets	−0.003	0.007	0.017	0.028**	0.018***
	(0.010)	(0.010)	(0.013)	(0.011)	(0.006)
Asset growth	0.067***	0.043***	0.044***	0.010*	0.002
	(0.007)	(0.006)	(0.008)	(0.006)	(0.002)
Log employees/assets		−0.032***	−0.056***	−0.022*	−0.010
		(0.011)	(0.014)	(0.013)	(0.012)
Capital expenditure/assets		0.049***	0.044***	0.025***	−0.001
		(0.009)	(0.012)	(0.009)	(0.004)
Net PP&E/assets		0.049***	0.051***	0.026*	0.035***
		(0.011)	(0.014)	(0.014)	(0.010)
Inventory/assets		0.202***	0.179***	0.056***	0.014**
		(0.010)	(0.013)	(0.013)	(0.007)
Cost of goods sold/sales		0.240***	0.269***	0.102***	0.049***
		(0.011)	(0.013)	(0.013)	(0.009)
R&D/assets		0.218***	0.198***	0.052***	0.005
		(0.016)	(0.017)	(0.015)	(0.010)
R&D not reported		0.000	−0.004	0.010	0.008
		(0.010)	(0.012)	(0.010)	(0.006)
Foreign sales share			0.074***	0.006	0.004
			(0.012)	(0.012)	(0.009)
Number of business segments			−0.014	0.011	0.011**
			(0.012)	(0.008)	(0.005)
Year FE	Yes	Yes	Yes	Yes	Yes
NAICS6 FE	–	–	–	Yes	–
Country FE	–	–	–	Yes	–
Firm FE	–	–	–	–	Yes
R^2	0.013	0.175	0.164	0.568	0.976
Within R^2	–	–	–	0.036	0.017
Zero dependence (%)	61.7	62.9	58.6	58.6	58.6
Observations	57,061	41,697	27,162	27,162	27,162
Firms	11,046	8,913	6,967	6,967	6,967

Notes: This table reports regressions of the dependence rank on firm characteristics for user firm-years in the full-information sample, including measured zeros. Continuous characteristics enter as sample percentiles, so coefficients compare the bottom and top of each distribution. Columns (1)–(3) add standard characteristics, production characteristics, and market scope. Column (4) adds six-digit industry and country fixed effects, and column (5) replaces them with firm fixed effects. Log assets and log employees/assets enter relative to their country–year means. Unreported R&D/assets is set to zero, and ratios are winsorized at the 1st and 99th percentiles within year. Firm-clustered standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 9: Dollar Responses per Percentage Point of Predicted Cost Pressure

		Before	Quarters after the price change		
		$t-9$ to $t-5$	2	4	8
(1)	Cost of goods sold	-0.505^{**}	0.783^{***}	1.025^{***}	0.856^{**}
		(0.223)	(0.166)	(0.203)	(0.364)
	Observations	134,960	136,789	124,119	98,764
(2)	Sales	-0.398	0.813^{***}	0.981^{***}	0.660
		(0.277)	(0.173)	(0.256)	(0.485)
	Observations	136,400	138,350	125,554	99,921
(3)	Gross profit	0.075	0.042	-0.031	-0.205
		(0.101)	(0.094)	(0.120)	(0.164)
	Observations	134,951	136,779	124,113	98,761

Notes: This table reports changes in cost of goods sold, sales, and gross profit per one-percentage-point increase in predicted cost pressure. The regressions follow Equation (16), with pressure in percentage points of identified revenue. Each outcome is measured as a percentage of sales at $t-1$. Columns $t+h$ measure the change from $t-1$ to $t+h$. Before uses $t-9$ to $t-5$ and sales at $t-9$. Both endpoints must use the same reporting currency. Outcomes are winsorized at the 1st and 99th percentiles within quarter. The specification and sample follow Appendix Table A.8. Standard errors are clustered by firm and quarter. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Material Cost Pressure by Event-Driven Price Changes

		Event-driven price changes			Remaining price changes		
	Quarters after the price change	0	2	4	0	2	4
(1)	Sales	0.043 (0.393)	1.691** (0.659)	1.860* (0.933)	0.321*** (0.090)	0.609*** (0.118)	0.703*** (0.155)
(2)	Cost-of-sales ratio	0.188 (0.113)	0.159 (0.245)	0.631** (0.248)	−0.015 (0.051)	0.080 (0.056)	0.134** (0.060)
(3)	Operating cash flow	−0.144*** (0.052)	−0.317*** (0.093)	−0.391** (0.147)	−0.046*** (0.013)	−0.102*** (0.025)	−0.135*** (0.040)
(4)	Capital expenditure	0.000 (0.025)	0.052 (0.061)	0.151 (0.094)	0.010* (0.006)	0.031* (0.018)	0.057* (0.030)
(5)	Net financing	0.120 (0.102)	0.382 (0.226)	0.692** (0.295)	0.038* (0.022)	0.087 (0.051)	0.094 (0.078)
(6)	Short-term debt	0.020 (0.045)	0.040 (0.082)	0.243** (0.104)	0.020** (0.008)	0.056** (0.023)	0.057 (0.036)
(7)	Long-term debt	0.095 (0.057)	0.403*** (0.123)	0.536*** (0.143)	0.010 (0.011)	0.043 (0.027)	0.051 (0.043)
(8)	Stock return	0.233 (0.407)	−0.335 (0.727)	−0.581 (0.838)	0.157* (0.088)	0.378** (0.153)	0.582** (0.225)

Notes: This table decomposes firms' material price changes into an event-driven part and the remainder. For each firm-quarter and horizon h, the dependence-weighted log price change from $t-4$ to $t+h$ pools the pressure of Equation (16) with its later-price terms. For each material with a tightening event in the list announced inside this window, the event-driven part cumulates, from the announcement, the difference between the actual quarterly log price changes and the changes predicted by a per-material model with quarter-of-year terms and a common factor, the median change across non-gold materials without a listed event that quarter, estimated on quarters untouched by any listed event. The tightening events in the list and the seven easing events identified by the same screening each exclude the announcement quarter and the two following. The remainder is the total change minus the event-driven part: the variation predictable without the events plus pre-event movement. The event-driven part is winsorized at the 1st and 99th percentiles of its nonzero values, the remainder like pressure, and both enter in percentage points of identified revenue, so coefficients are per percentage point. The full-sample standard deviation of cost pressure is 1.19 percentage points. The event-driven part is nonzero for 30,045 of 158,068 firm-quarters at $h = 4$. Gold and silver enter as one price change over the same window, which also replaces the later-price terms. The sample and other regression details follow Appendix Table A.8. Standard errors are clustered by firm and quarter. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Event-Driven Cost Pressure at the Firm–Material Level

		Documented			Not documented			Difference
	Quarters after the price change	0	2	4	0	2	4	4
(1)	Sales	−0.877	0.638	0.456	0.749*	1.499**	1.221	−0.765
		(0.697)	(0.912)	(1.272)	(0.396)	(0.683)	(0.858)	(1.424)
	Observations	69,859	63,277	57,146	290,809	265,189	240,880	
(2)	Cost-of-sales ratio	0.225	0.270	0.575	0.219	0.528	0.348	0.226
		(0.236)	(0.315)	(0.391)	(0.196)	(0.325)	(0.333)	(0.466)
	Observations	69,653	63,069	56,953	290,275	264,696	240,412	
(3)	Operating cash flow	−0.252**	−0.429**	−0.694**	−0.074	−0.353**	−0.390	−0.304
		(0.099)	(0.169)	(0.266)	(0.076)	(0.161)	(0.230)	(0.373)
	Observations	69,652	62,326	55,397	290,013	261,860	234,637	
(4)	Net financing	0.047	0.555	1.092**	−0.109	−0.023	−0.036	1.128*
		(0.130)	(0.360)	(0.474)	(0.089)	(0.247)	(0.379)	(0.625)
	Observations	68,498	60,553	53,498	284,466	253,087	225,080	
(5)	Short-term debt	−0.110	0.088	0.490**	0.119*	0.128	0.247	0.243
		(0.121)	(0.165)	(0.217)	(0.065)	(0.156)	(0.188)	(0.292)
	Observations	65,999	59,390	53,323	274,283	248,632	224,878	

Notes: This table compares responses to event-driven material price changes for documented and undocumented relationships. The sample contains firm–material–quarters with positive dependence on one of 42 priced materials. The specification is Equation (17). The material's event-driven and remaining pressure terms cover the window from $t-4$ to $t+h$ and are constructed as in Table 10. The regressions control for the full-window price changes of the firm's other materials and of gold and silver, which absorb the later-price terms. The two groups are estimated jointly with group-specific fixed effects. Specialist reports document a relationship when they name goods carrying at least half of the firm's dependence on the material. The event-driven part is nonzero in 65,486 observations, 14,596 of them documented. Coefficients are the event-driven part's, per percentage point of material-specific pressure, whose standard deviation is 0.47 percentage points of identified revenue. The regressions include material-by-quarter fixed effects. The difference column is documented minus undocumented at quarter 4. Standard errors are clustered by firm and quarter. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 12: Event-Driven Cost Pressure by Firm Size

		Small firms			Large firms			Difference
	Quarters after the price change	0	2	4	0	2	4	4
(1)	Sales	0.551	2.501**	2.009	−0.033	1.465*	1.754*	0.846
		(0.839)	(1.141)	(1.606)	(0.414)	(0.728)	(0.965)	(1.570)
	Observations	76,269	69,306	62,824	75,278	68,765	62,448	
(2)	Cost-of-sales ratio	0.509*	0.621	1.258***	−0.047	−0.019	0.337	0.337
		(0.276)	(0.393)	(0.371)	(0.209)	(0.295)	(0.296)	(0.411)
	Observations	75,317	68,428	62,026	74,554	68,089	61,826	
(3)	Operating cash flow	−0.120	−0.367**	−0.727**	−0.106**	−0.228**	−0.178	−0.439
		(0.074)	(0.174)	(0.295)	(0.052)	(0.105)	(0.174)	(0.308)
	Observations	76,469	68,640	61,178	75,018	67,957	60,941	
(4)	Short-term debt	0.054	0.174	0.668***	−0.006	−0.008	0.072	0.420*
		(0.122)	(0.160)	(0.240)	(0.046)	(0.102)	(0.110)	(0.218)
	Observations	68,663	61,750	55,471	72,056	65,582	59,421	
(5)	Net financing	0.228	1.019**	1.748**	0.044	0.179	0.430	0.567
		(0.167)	(0.427)	(0.673)	(0.104)	(0.213)	(0.273)	(0.604)
	Observations	73,927	64,605	56,796	74,317	66,858	59,746	

Notes: This table compares responses to event-driven cost pressure for small and large firms. Firms are split at the country–quarter median of total assets in U.S. dollars at $t-1$. The specification of Table 10 is estimated once in each group, and the table reports the event-driven part's coefficients. The difference column is the small-minus-large contrast at quarter 4 from a pooled specification in which both parts are interacted with the size group and the controls and fixed effects are shared across groups, so it need not equal the gap between the displayed group coefficients. The event-driven part is nonzero for 12,480 small and 16,479 large firm-quarters at $h = 4$. Coefficients are per percentage point of cost pressure. Outcomes and other regression details follow Appendix Table A.8. Standard errors are clustered by firm and quarter. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 13: Sources and Uses of Funds under Event-Driven Cost Pressure

	Quarter 4	Quarter 8
Panel A. Sources of funds		
Operating cash flow	−0.306** (0.147)	−0.090 (0.305)
inside it, working capital: inventories	0.116 (0.135)	0.341* (0.186)
receivables	0.134 (0.168)	0.165 (0.264)
payables (+)	0.131 (0.156)	0.445** (0.200)
Net external financing	0.617** (0.267)	0.593 (0.406)
balance-sheet change in total debt	0.711*** (0.156)	0.452 (0.279)
memo: net equity issuance (own sample)	1.259 (0.799)	0.998 (1.594)
Panel B. Uses of funds		
Capital expenditure	0.169* (0.097)	0.262 (0.212)
Change in cash holdings	0.144 (0.179)	0.318 (0.229)
Other investing and residual flows	−0.001 (0.098)	−0.077 (0.188)

Notes: This table organizes the quarter-4 and quarter-8 responses to event-driven cost pressure into sources and uses of funds. The coefficients are the event-driven part's from the specification of Table 10, in percent of assets at $t-1$ per percentage point of cost pressure. Each column uses its own common sample of firm-quarters reporting every listed item except net equity issuance: 94,635 at quarter 4 and 67,927 at quarter 8, with the event-driven part nonzero in 23,311 and 19,505 of them. Operating cash flow includes working-capital changes, reported in the three indented rows. Payables enter cash flow positively. Net equity issuance is stock sales less purchases and uses its own sample, 22,174 firm-quarters at quarter 4 and 15,643 at quarter 8, chiefly for U.S. filers. The change in total debt is a balance-sheet change rather than a financing flow, so it need not lie below net financing. Dividends, acquisitions, and other flows enter only through the residual row, which closes each column's identity by construction. Standard errors are clustered by firm and quarter. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14: Dollars at Risk from Event-Driven Price Changes

	Exposed firms	Their assets ($bn)	Cash flow at risk ($bn)	Financing raised ($bn)	Undocumented share (%)
Panel A. A 10 percent event-driven price increase, one material at a time					
Copper	1,429	9,377	4.0	7.0	25
Aluminum	1,360	9,403	2.3	4.1	76
Nickel	1,004	8,839	1.7	3.0	13
Silicon	1,019	6,766	1.5	2.7	72
Cobalt	828	8,955	1.5	2.7	17
Platinum	397	5,072	1.1	2.0	21
Palladium	715	6,425	0.9	1.5	98
Lithium	431	6,147	0.8	1.5	7
Germanium	263	4,146	0.6	1.1	88
Zinc	661	7,515	0.6	1.0	91
All priced materials	2,662	13,686	17.8	31.5	48
Panel B. Events in the list replayed at their observed event-driven price changes					
Li, 2022Q2 export act	431	6,147	7.6	13.5	9
Co, 2025Q1 export act	828	8,955	6.9	12.2	13
Li, 2019Q1 mine disruption	431	6,147	3.6	6.3	8
Al, Ni, 2021Q2 export act	1,582	11,189	3.4	6.1	36
Bi, In, W, 2025Q1 export act	913	8,520	3.4	5.9	96
Ga, Ge, 2023Q3 export act	525	5,722	3.2	5.7	82
All events announced in 2025Q1, jointly	1,380	10,720	10.5	18.5	41
All 28 events, summed	2,256	13,120	20.0	35.4	29

Notes: This table converts quarter-4 responses to event-driven material-price changes into dollar amounts. The cross section uses each firm's last 2024–2025 observation with assets in U.S. dollars. Dollar amounts use firm assets and the quarter-4 event-driven coefficients from Table 10, -0.39 for operating cash flow and 0.69 for net financing per percentage point of pressure, with each firm–material pressure capped at the 99th percentile of the estimation sample in absolute value. In Panel A, a 10 percent price increase, pressure of $100\,D_{mf}\ln(1.1)$, is applied to one material at a time. The panel lists the ten materials with the largest cash flow at risk, and its total covers 41 materials because no sample firm is exposed to thulium. In Panel B, each tightening event in the list is replayed at its observed event-driven price changes, cumulated from the announcement through four quarters as in Table 10 and truncated at the end of the price data for the 2025 events. An event moves all of the materials recorded for it jointly, and the changes keep their signs, so moves in opposite directions offset within the joint rows. The panel lists the six events with the largest cash flow at risk, the events announced in 2025Q1 jointly, and the sum over all 28 events. Undocumented share is the share of cash-flow dollars, in absolute value in Panel B, for which specialist reports do not name goods carrying at least half of the firm's dependence. Both panels value documented and undocumented relationships with the same pooled coefficients. Table 11 estimates the two groups' responses separately.

Table 15: Material Dependence and Market Outcomes

	Return volatility (1)	Idiosyncratic volatility (2)	Expected shortfall (3)	Market beta (4)	Log market-to-book (5)	Bid-ask spread (6)
Panel A. Dependence rank, all firm-years						
Dependence rank	2.414***	1.440***	0.202***	4.176**	11.523***	2.637**
	(0.550)	(0.422)	(0.054)	(1.812)	(3.228)	(1.124)
Observations	50,472	49,021	51,778	43,922	48,450	50,445
Firms	9,763	9,585	9,894	8,978	9,169	9,750
Within R^2	0.081	0.265	0.221	0.018	0.078	0.239
Panel B. Rank of the largest single-material dependence, all firm-years						
Largest material rank	2.239***	1.249***	0.183***	4.279**	10.548***	2.261**
	(0.547)	(0.419)	(0.053)	(1.812)	(3.209)	(1.113)
Observations	50,472	49,021	51,778	43,922	48,450	50,445
Firms	9,763	9,585	9,894	8,978	9,169	9,750
Within R^2	0.081	0.265	0.221	0.018	0.078	0.239
Panel C. Number of materials, firm-years with positive dependence						
Number of materials	0.083***	0.084***	0.008***	0.050	0.443**	0.142**
	(0.030)	(0.025)	(0.003)	(0.101)	(0.177)	(0.063)
Observations	18,150	18,130	18,622	15,813	18,312	18,357
Firms	3,497	3,498	3,550	3,211	3,459	3,533
Within R^2	0.057	0.218	0.193	0.023	0.099	0.215
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
NAICS6 × year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country × year FE	Yes	Yes	Yes	Yes	Yes	Yes

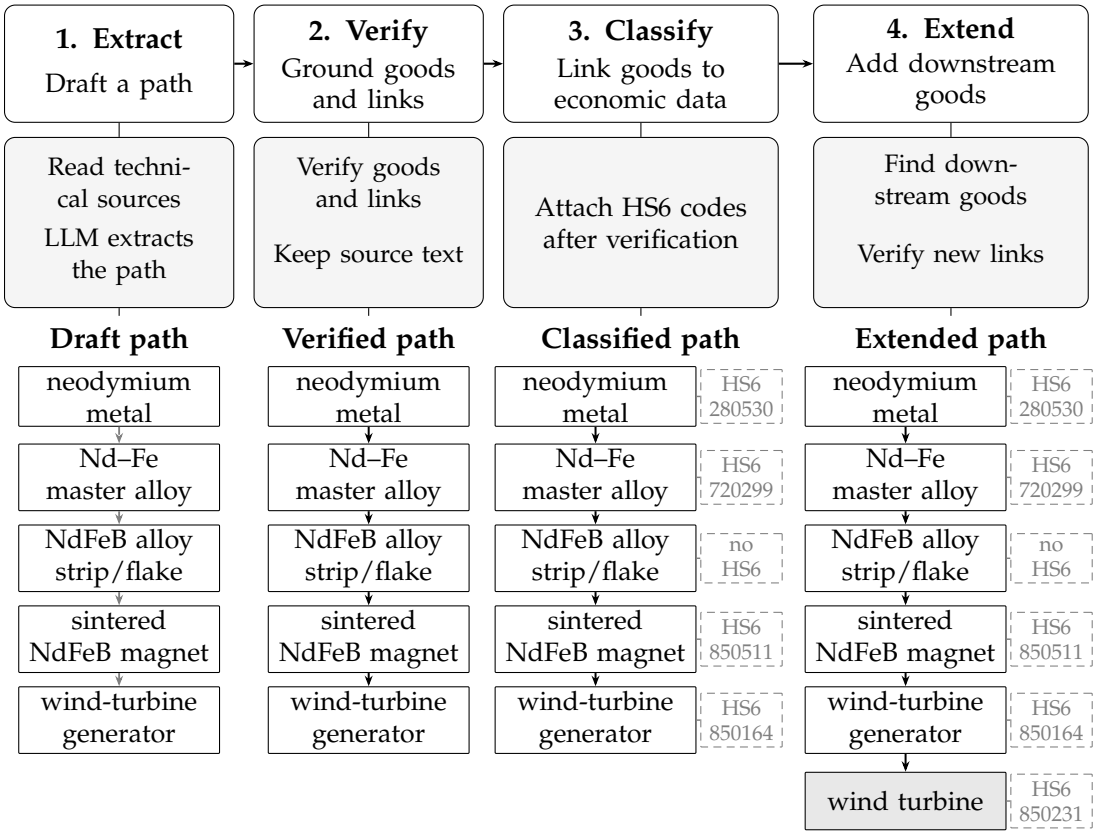
Notes: This table reports regressions of market outcomes on material dependence in the full-information sample, excluding upstream firms. Panels A and B use the dependence rank and the largest-material rank, with measured zeros assigned zero. Panel C restricts the sample to positive dependence and uses the number of materials. Each regression includes the firm controls of Section 5.2, six-digit industry-by-year fixed effects in cells with at least five firms, and country-by-year fixed effects. Outcomes are winsorized within year. Market beta uses 60 monthly returns with at least 36 available. Idiosyncratic volatility and expected shortfall require 120 daily returns. Log market-to-book and market beta are multiplied by 100, volatility and expected shortfall are in percent, and spreads in basis points. Firm-clustered standard errors are in parentheses. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table 16: Material Price Risk and Market Outcomes

	Return volatility (1)	Idiosyncratic volatility (2)	Expected shortfall (3)	Market beta (4)	Log market-to-book (5)	Bid-ask spread (6)
Panel A. Price-risk and China-restriction baskets						
Dependence rank	2.058**	1.291**	0.200**	4.073	5.468	4.803***
	(0.884)	(0.654)	(0.083)	(3.023)	(5.091)	(1.744)
Price-risk basket	0.740***	0.503***	0.063***	1.864***	1.877	0.955***
	(0.200)	(0.131)	(0.014)	(0.702)	(1.250)	(0.342)
China-restriction basket	0.614***	0.389***	0.014	0.682	1.529	0.789***
	(0.113)	(0.087)	(0.015)	(0.536)	(1.208)	(0.171)
Within R^2	0.060	0.226	0.197	0.024	0.101	0.222
Panel B. Baskets and dependence rank						
Dependence rank	1.989**	1.242*	0.195**	4.042	5.166	4.767***
	(0.883)	(0.657)	(0.083)	(3.028)	(5.066)	(1.750)
Price-risk basket	−0.338*	−0.191	−0.030	1.614**	−3.091***	−0.218
	(0.194)	(0.172)	(0.019)	(0.658)	(1.059)	(0.389)
Price-risk basket × dependence rank	2.193***	1.403***	0.189***	0.509	10.034***	2.395***
	(0.449)	(0.247)	(0.028)	(1.989)	(2.512)	(0.706)
China-restriction basket	0.513***	0.230***	0.012	0.605	0.703	0.853***
	(0.134)	(0.087)	(0.015)	(0.763)	(1.198)	(0.172)
China-restriction basket × dependence rank	0.182	0.341	−0.001	0.169	1.703	−0.217
	(0.199)	(0.236)	(0.031)	(1.339)	(3.281)	(0.553)
Within R^2	0.060	0.226	0.198	0.024	0.103	0.222
Observations	17,409	17,390	17,852	15,167	17,528	17,603
Firms	3,340	3,342	3,387	3,059	3,298	3,372
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
NAICS6 × year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country × year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table relates market outcomes to the price-risk and China-restriction baskets of users with positive dependence on priced materials other than gold and silver. Each basket is a dependence-weighted average across the firm's materials. Baskets are centered, and coefficients are per standard deviation: 1.7 percentage points for monthly price volatility and 0.73 for China-restriction risk. All regressions include the dependence rank, the share on gold, silver, and unpriced materials, and the controls and fixed effects of Table 15. Panel A enters both baskets. Panel B also interacts each with the dependence rank. Standard errors for baskets and interactions are exposure-robust using control-adjusted material shifters (Adão et al., 2019). Those for dependence rank are clustered by firm. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

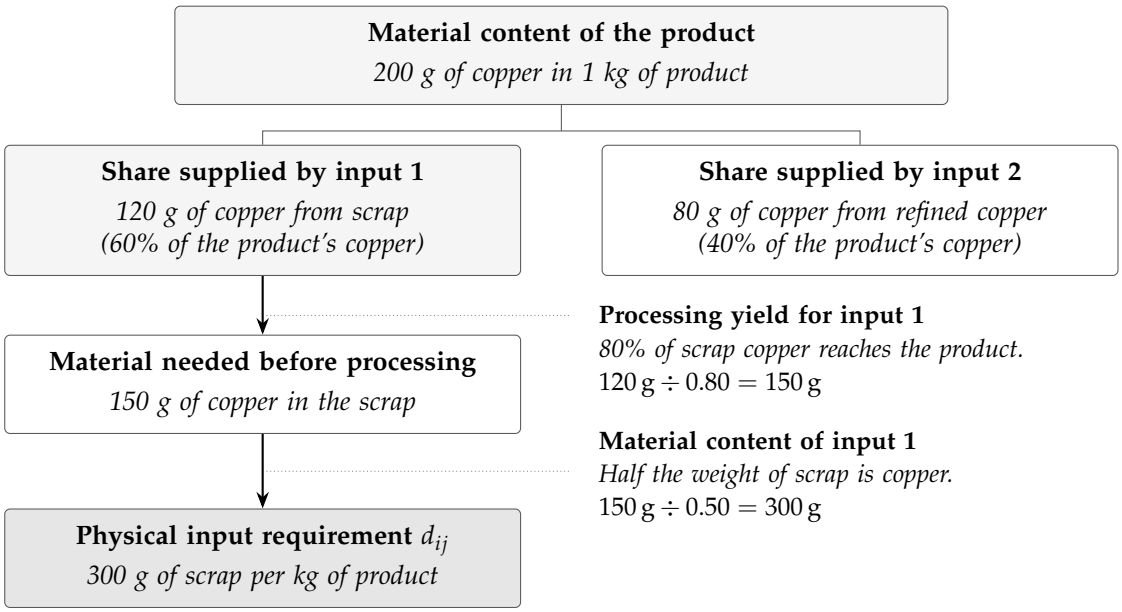
Figure 1: Construction of the Material-Specific Production Networks



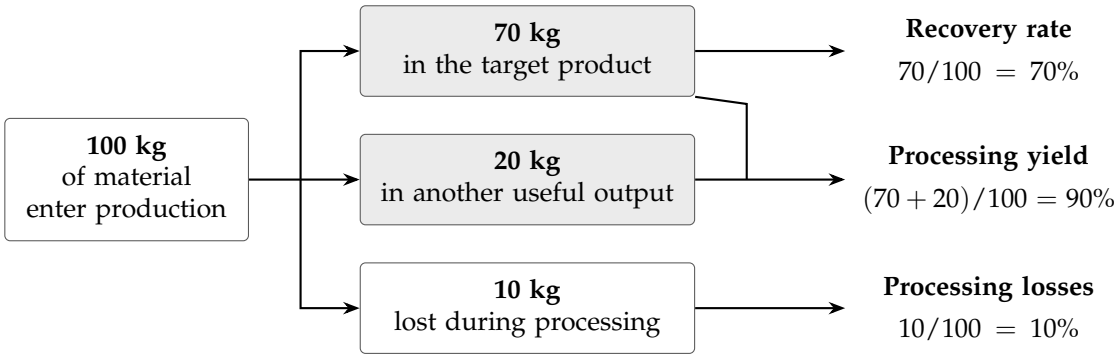
Notes: The figure illustrates the four construction steps using a neodymium production path. Dashed arrows mark candidate links and solid arrows verified links. HS6 labels sit beside the classified goods. Here, "no HS6" marks an intermediate without a standalone code. The shaded node is a downstream addition.

Figure 2: Physical Inputs and Processing Losses

A. How much of an input is needed for 1 kg of product?

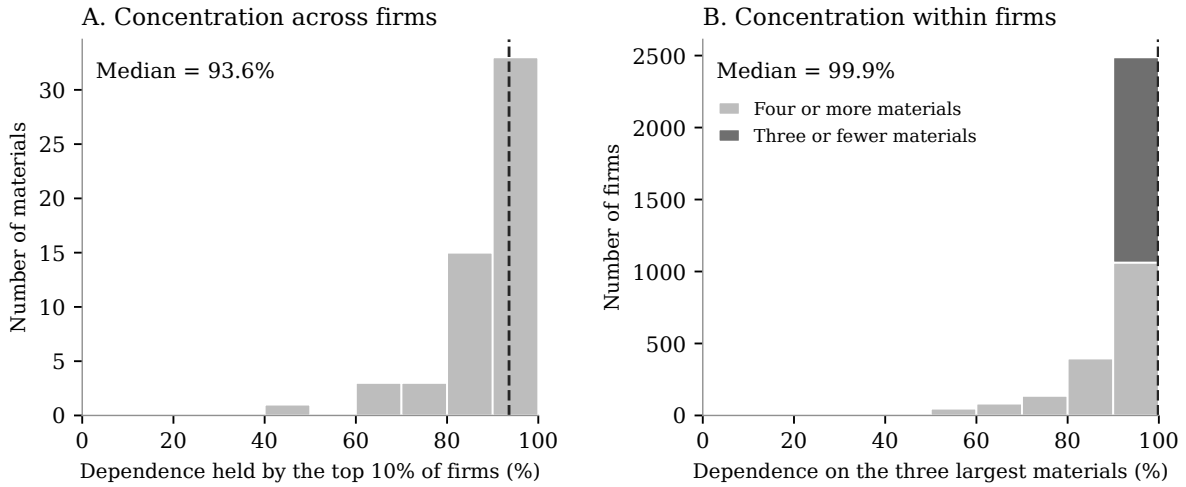


B. Recovery rate and processing yield



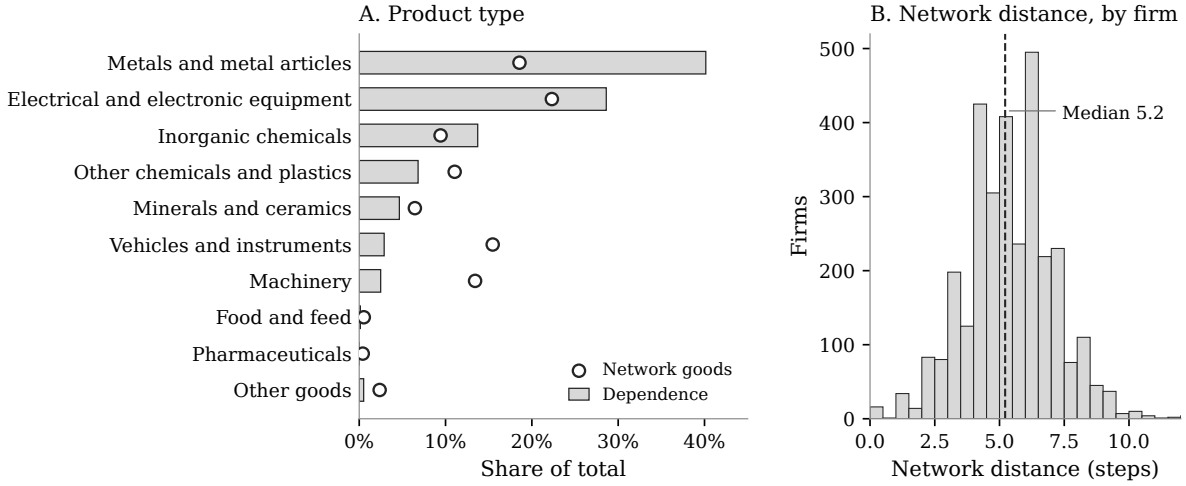
Notes: Both panels use hypothetical quantities. In Panel A, 20% of the copper in the scrap is lost during processing. Panel B traces the material through a separate production step. For an input used up in production, such as a reagent, the input requirement is the net quantity consumed per kilogram of output. Durable equipment is excluded from input requirements.

Figure 3: Concentration of Material Dependence



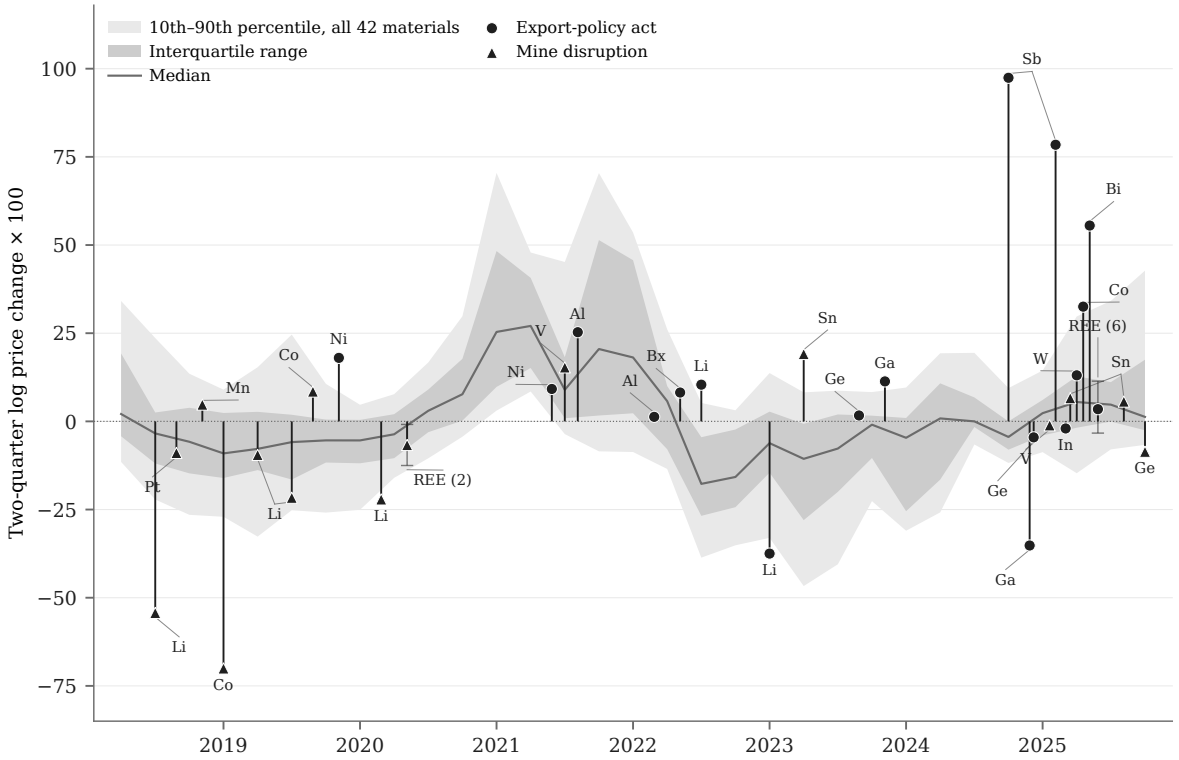
Notes: The sample is the 2019 cross section of 3,165 firms with positive total dependence and 55 materials. Panel A plots the distribution across materials of the share of dependence held by the most dependent 10 percent of firms. Firms are ranked separately for each material among those with positive dependence, and shares use dependence summed across those firms. The top-decile count is rounded to the nearest integer, with a minimum of one firm. Panel B plots the distribution across firms of the sum of their three largest material dependencies divided by total dependence. Dark shading marks the 1,426 firms with positive dependence on three or fewer materials, whose share is 100 percent by construction. The median among the other firms is 94.0 percent. Bins are ten percentage points wide, and dashed lines mark the medians. Panel A counts materials and Panel B counts firms.

Figure 4: Material Dependence in the Production Network



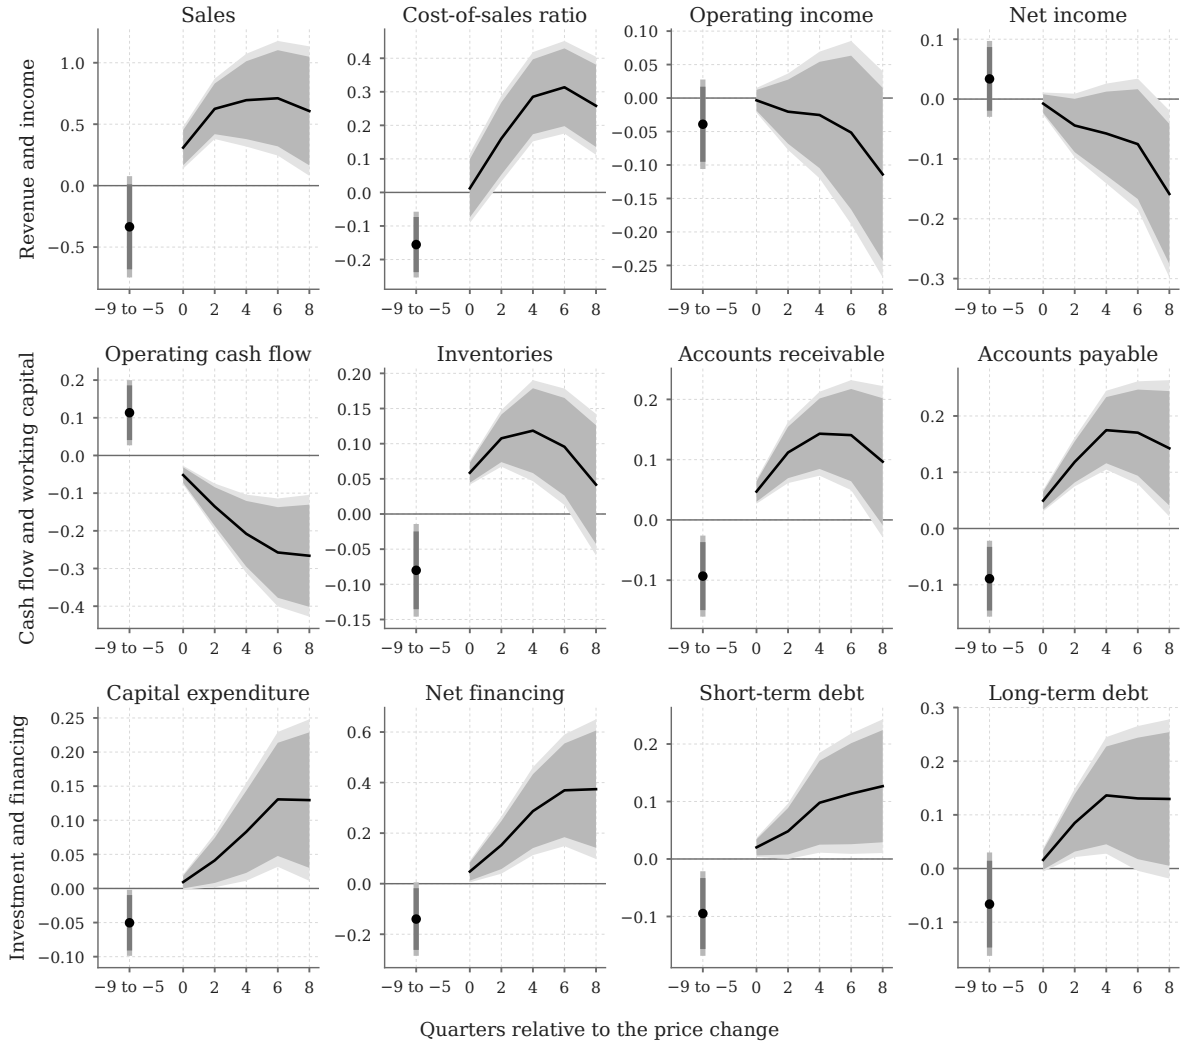
Notes: The figure assigns each firm's material dependence (D_{mft}) to the network goods that carry it, for the 3,165 firms with positive dependence in 2019. Panel A compares each product type's share of total dependence, summed over firms and materials (bars), with its share of the 2,451 network goods carrying that dependence (circles). Product types follow the HS chapter of the firm's product line. A good whose codes span several types counts fractionally. The types are: food and feed (01–24), minerals and ceramics, including glass (25–27, 68–70), inorganic chemicals (28), pharmaceuticals (30), other chemicals and plastics (29, 31–40), metals and metal articles (71–83), machinery (84), electrical and electronic equipment (85), and vehicles and instruments (86–92). Panel B reports the distribution across firms of network distance, the dependence-weighted number of production steps from an entry point of the material's network, such as an ore, a recovered stream or a purchased refined input. The dashed line marks the median.

Figure 5: Material Price Changes around Supply-Tightening Events



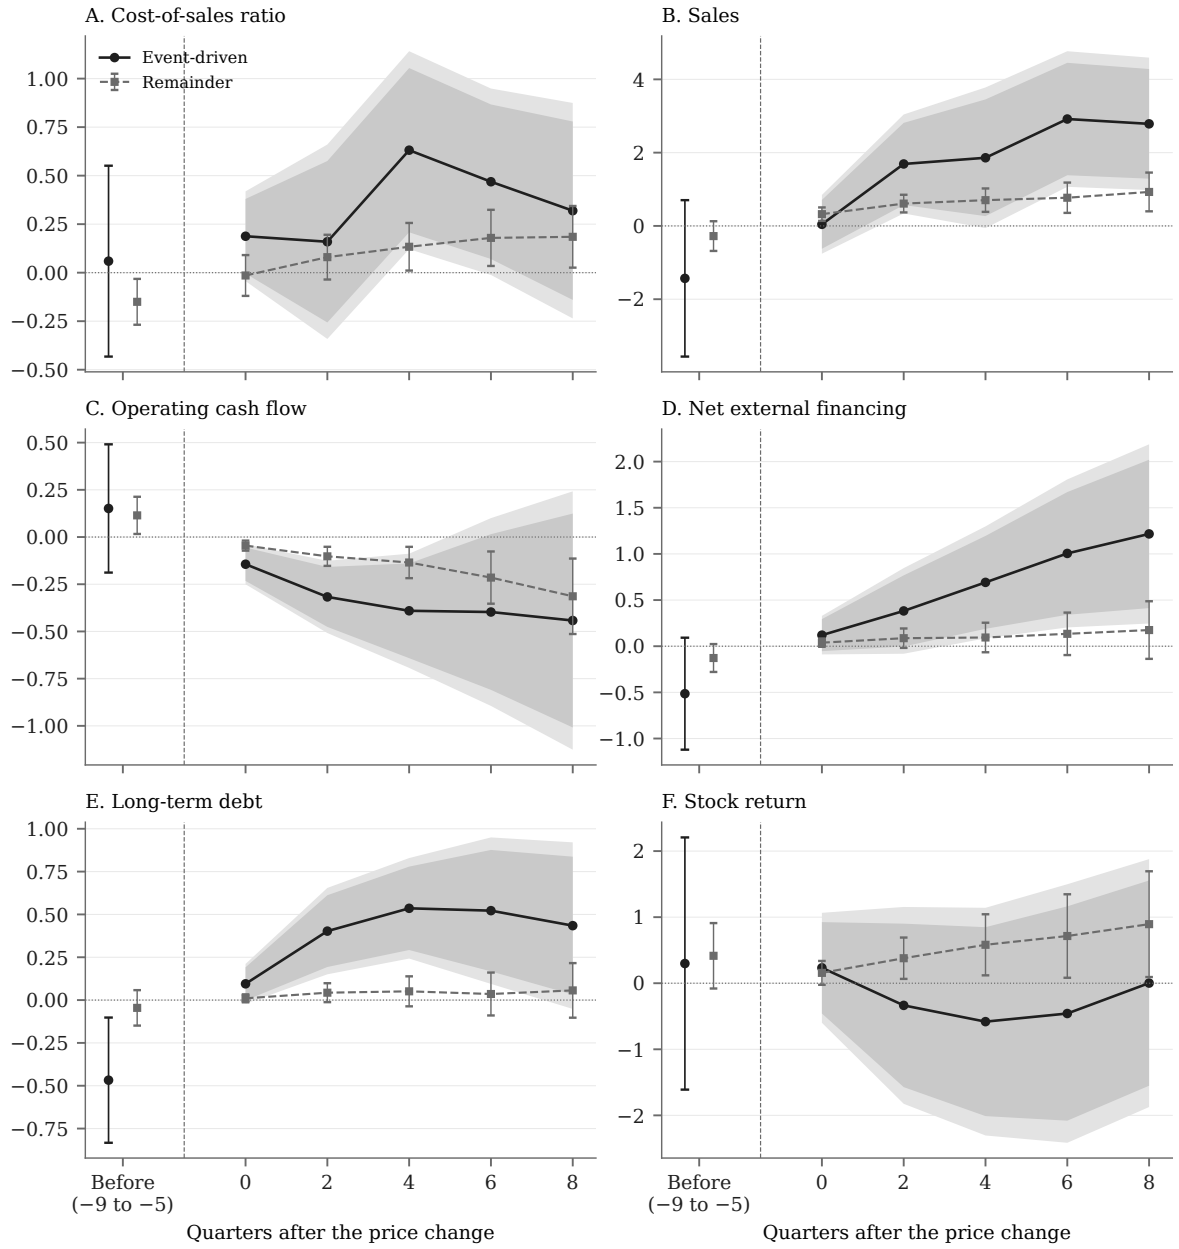
Notes: The figure plots, for each quarter of the regression sample, the distribution across the 42 priced materials other than gold and silver of the two-quarter log price change, multiplied by 100. The line is the median, the dark band the interquartile range, and the light band the 10th–90th percentile range. Circles and triangles mark export-policy acts and mine disruptions in the event list. Each marker gives the affected material's price change from the quarter before the announcement to the quarter after, $100[\ln p(e+1) - \ln p(e-1)]$. Labels identify the materials, with REE denoting rare-earth elements and the number in parentheses their count. Prices are quarterly means of the monthly U.S.-dollar benchmarks used in Equation (13).

Figure 6: Material Cost Pressure, Operations and Financing



Notes: Each panel plots the coefficient on cost pressure from Equation (16) at quarters 0 to 8 as a line, with 90 and 95 percent confidence bands in dark and light gray. The point marked -9 to -5 is the change over the window from $t-9$ to $t-5$, before the price change begins, with the same two intervals. Coefficients are per percentage point of cost pressure, whose standard deviation is 1.19 percentage points of identified revenue. Sales is in percent, the cost-of-sales ratio in percentage points, and the other outcomes in percent of assets at $t-1$. Appendix Table A.8 reports these estimates and adds cash, total debt, and the interest rate on debt. Standard errors are clustered by firm and quarter.

Figure 7: Responses to Event-Driven and Remaining Cost Pressure



Notes: Each panel plots the coefficients on the event-driven part of cost pressure, as a solid line with 90 and 95 percent confidence bands in dark and light gray, and on the remainder, as a dashed line with 95 percent error bars, from the specification of Table 10 at quarters 0 to 8 after the price change. The points marked -9 to -5 are the changes over the window before the price change begins, with 95 percent intervals. Both parts enter in percentage points of identified revenue, so coefficients are per percentage point. Sales and the stock return are in percent, the cost-of-sales ratio in percentage points, and the other outcomes in percent of assets at $t-1$. Table 10 defines the decomposition and reports the estimates at quarters 0, 2, and 4. Standard errors are clustered by firm and quarter.

A Additional Tables and Figures

Table A.1: Variable Definitions and Sources

Variable	Definition	Source
Total dependence	Sum across the 55 materials of each good's dependence per dollar of sales, aggregated over product lines by revenue share. Bauxite with aluminum and magnesite with magnesium are each counted once, and a good's total is capped at one	Section 3.3
Material dependence	The firm's dependence on one material	Section 3.3
Dependence rank	Percentile of positive dependence among downstream firm-years (0–1). Measured zeros receive zero	Section 7.3
Largest-material rank	Percentile of the firm's largest single-material dependence among downstream firm-years (0–1). Measured zeros receive zero	Section 7.3
Number of materials	Materials with positive dependence	Section 3.3
Producer	Firm with at least 25 percent of identified revenue in goods it makes at or above a material's benchmark form or in its near-pure forms	Section 3.4
Merchant	Firm other than a producer with at least 25 percent of identified revenue in such goods it resells	Section 3.4
User	Downstream firm, one that is neither a producer nor a merchant	Section 3.4
Material share	Material's share of the firm's total dependence	Section 6.4
Hard-to-substitute share	Fraction of scored material dependence from products whose minimum substitution score is 2 or below (major penalty)	Section 6.4
Broad-window supply-driven pressure	The part of cost pressure from materials with a tightening event in the event list in the quarter of the price change or the quarter before	Section 5.4
Price-risk basket	Dependence-weighted average of the standard deviation of monthly log price changes over 2009–2025 for the 42 priced materials other than gold and silver	Section 5.3
China-restriction basket	Dependence-weighted average of the log of one plus Chinese export acts covering each material over 2017–2025	Section 5.3
Other dependence share	Share of the firm's dependence on gold, silver and unpriced materials	Section 7.3
Cost pressure	100 times the sum over the 42 priced materials other than gold and silver of dependence times the four-quarter log price change	Equation (13)
Material-specific pressure	One material's term in cost pressure	Section 7.2
Later price change	100 times dependence at t times the log price change from t to $t+h$, summed over the priced materials	Section 7.2
Documented relationship	Specialist reports name the goods carrying at least half of the firm's dependence on the material	Section 7.2
Sales change	Change in log sales from $t-1$ to $t+h$, in percent	Compustat
Cost-of-sales ratio change	Change in cost of goods sold divided by sales from $t-1$ to $t+h$, in percentage points	Compustat
Inventories, accounts receivable, accounts payable, cash	Change from $t-1$ to $t+h$ over assets at $t-1$	Compustat
Total, short-term and long-term debt	Change from $t-1$ to $t+h$ over assets at $t-1$. Total debt is short-term plus long-term debt	Compustat
Operating income	Quarterly operating income, summed from t to $t+h$ over assets at $t-1$	LSEG Refinitiv Workspace
Net income	Quarterly income before extraordinary items, summed from t to $t+h$ over assets at $t-1$	Compustat

Table A.1, continued

Variable	Definition	Source
Operating cash flow	Year-to-date operating cash flow differenced within the fiscal year, summed from t to $t+h$ over assets at $t-1$	Compustat
Capital expenditure	Year-to-date capital expenditure differenced within the fiscal year, summed from t to $t+h$ over assets at $t-1$	Compustat
Net financing	Cash flow from financing activities, summed from t to $t+h$ over assets at $t-1$	Compustat
Interest rate on debt	Annualized interest expense over the previous quarter's debt, in percent	Compustat
Stock return	100 times the log change in the split- and dividend-adjusted total-return index from quarter-end $t-1$ to $t+h$, in the trading currency. Price returns are used when the total-return factor is missing	Compustat security files (daily prices)
Investment metals	Metals with liquid exchange-traded futures before 2018: gold, silver, platinum, palladium, aluminum, copper, lead, nickel, tin and zinc.	Section 5.4
Log market-to-book	Log of market equity over positive book equity	Compustat security files (daily prices) and Compustat book equity
Return volatility	Annualized standard deviation of monthly returns over the twelve months to fiscal-period end	Compustat security files (daily prices)
Market beta	Slope of monthly stock returns on the country market return over the prior 60 months	Compustat security files (daily prices)
Idiosyncratic volatility	Annualized standard deviation of residuals from daily regressions on the country value-weighted market return over the fiscal year	Compustat security files (daily prices)
Expected shortfall	Negative of the mean of the lowest 5 percent of daily returns in the fiscal year	Compustat security files (daily prices)
Bid-ask spread	High-low estimator of Corwin and Schultz (2012), averaged over the fiscal year, in basis points	Compustat security files (daily prices)
Foreign-sales share	One minus the home-country share of geographic-segment revenue, with at least two segments	LSEG Refinitiv Workspace
Export-restriction count	Implemented export acts per material and calendar year, among acts covering at most three of the 55 materials	Global Trade Alert

Notes: This table defines the variables and reports their sources. Section references give the construction of our measures. Annual financial variables are winsorized at the 1st and 99th percentiles within year, and quarterly outcomes within quarter.

Table A.2: Coverage of the Quantified Network and Firm Panel

	Count	Of which
Panel A. Materials and products		
Material–product pairs in the pricing scope	6,332	192 without an estimate
Pairs with a cost- or revenue-side estimate	6,140	759 upstream of the benchmark
Cost-side estimates: benchmark and downstream	5,381	723 bounds
walked estimates and benchmark forms	4,658	163 benchmark forms
Panel B. Firms, product-line measure		
Firm-years with accepted product lines, from fiscal 2017	87,042	14,102 firms
less exposure unavailable	5,180	
Firm-years with the measure	81,862	13,347 firms
positive dependence on at least one material	30,832	5,045 firms
zero dependence	51,030	
less under 75 percent of revenue identified	6,196	
less fiscal 2026	190	
Baseline sample	75,476	12,478 firms
positive dependence on at least one material	28,980	4,815 firms
zero dependence	46,496	

Notes: This table reports coverage of the material–product pricing scope and firm panel. Panel A separates cost-side estimates at or downstream of the benchmark from upstream revenue-side elasticities and reports pairs without an estimate. The treatment of bounds follows Appendix D. Panel B traces accepted product-line firm-years to the baseline sample. Positive and zero refer to the headline cost-side measure, and unavailable firm-years remain blank.

Table A.3: Product-Level Estimates

	Pairs	Mean	SD	P10	Median	P90
Panel A. Cost-side estimates by estimate type						
All cost-side estimates	5,381	19.82	31.21	0.00	1.83	80.72
Benchmark forms	163	93.71	6.35	86.88	99.20	99.21
Walked estimates	4,495	10.77	20.41	0.00	0.93	39.83
Bounds	723	59.47	36.61	0.50	75.34	97.80
Panel B. Product type (HS chapters)						
Food and feed (01–24)	37	6.61	12.11	0.03	0.82	22.62
Minerals and ceramics (25–27, 68–70)	409	28.42	36.10	0.12	6.85	89.48
Inorganic chemicals (28)	591	40.69	33.87	1.38	29.20	86.88
Pharmaceuticals (30)	51	3.95	8.68	0.01	0.55	9.99
Other chemicals and plastics (29, 31–40)	767	19.31	27.65	0.08	4.22	73.23
Metals and metal articles (71–83)	949	41.51	39.04	0.26	27.52	99.21
Machinery (84)	559	5.65	16.21	0.00	0.14	14.62
Electrical and electronic equipment (85)	951	6.73	17.78	0.00	0.24	17.45
Vehicles and instruments (86–92)	750	3.27	11.76	0.00	0.07	4.84
Other goods	317	13.64	24.97	0.01	0.87	57.83
Panel C. Material group						
Investment metals	1,754	26.19	34.24	0.03	5.12	86.88
Other materials	3,627	16.74	29.15	0.00	1.14	76.46
Panel D. Revenue-side elasticity for pairs upstream of the benchmark form						
All upstream pairs	759	57.69	50.98	0.43	50.29	110.89
Ores and intermediates	734	56.27	51.12	0.40	45.12	110.93
Near-pure forms	25	99.52	19.22	86.88	102.04	109.18

Notes: This table summarizes product-level estimates. Panels A–C report cost-side dependence at or downstream of the benchmark. Panel D reports upstream revenue-side elasticities, which can exceed one. Cost-side dependence is per dollar of product sales, and all entries are multiplied by 100. A good enters once for each material it uses. Bounds follow Appendix D. Product types use the good's customs code, and investment metals follow Section 5.4.

Table A.4: Sources of Variation in Material Dependence

	Dependence rank			Log dependence	Material rank
Industry level (NAICS digits)	2	4	6	6	6
Year	0.0	0.0	0.0	0.0	2.9
Industry	23.6	43.6	52.1	43.2	25.4
Permanent firm differences	73.0	52.9	44.5	53.7	66.7
Changes within firms	3.4	3.4	3.4	3.1	5.0
Observations	72,194	72,194	72,194	26,947	3,970,670
Firms	11,936	11,936	11,936	4,483	11,936

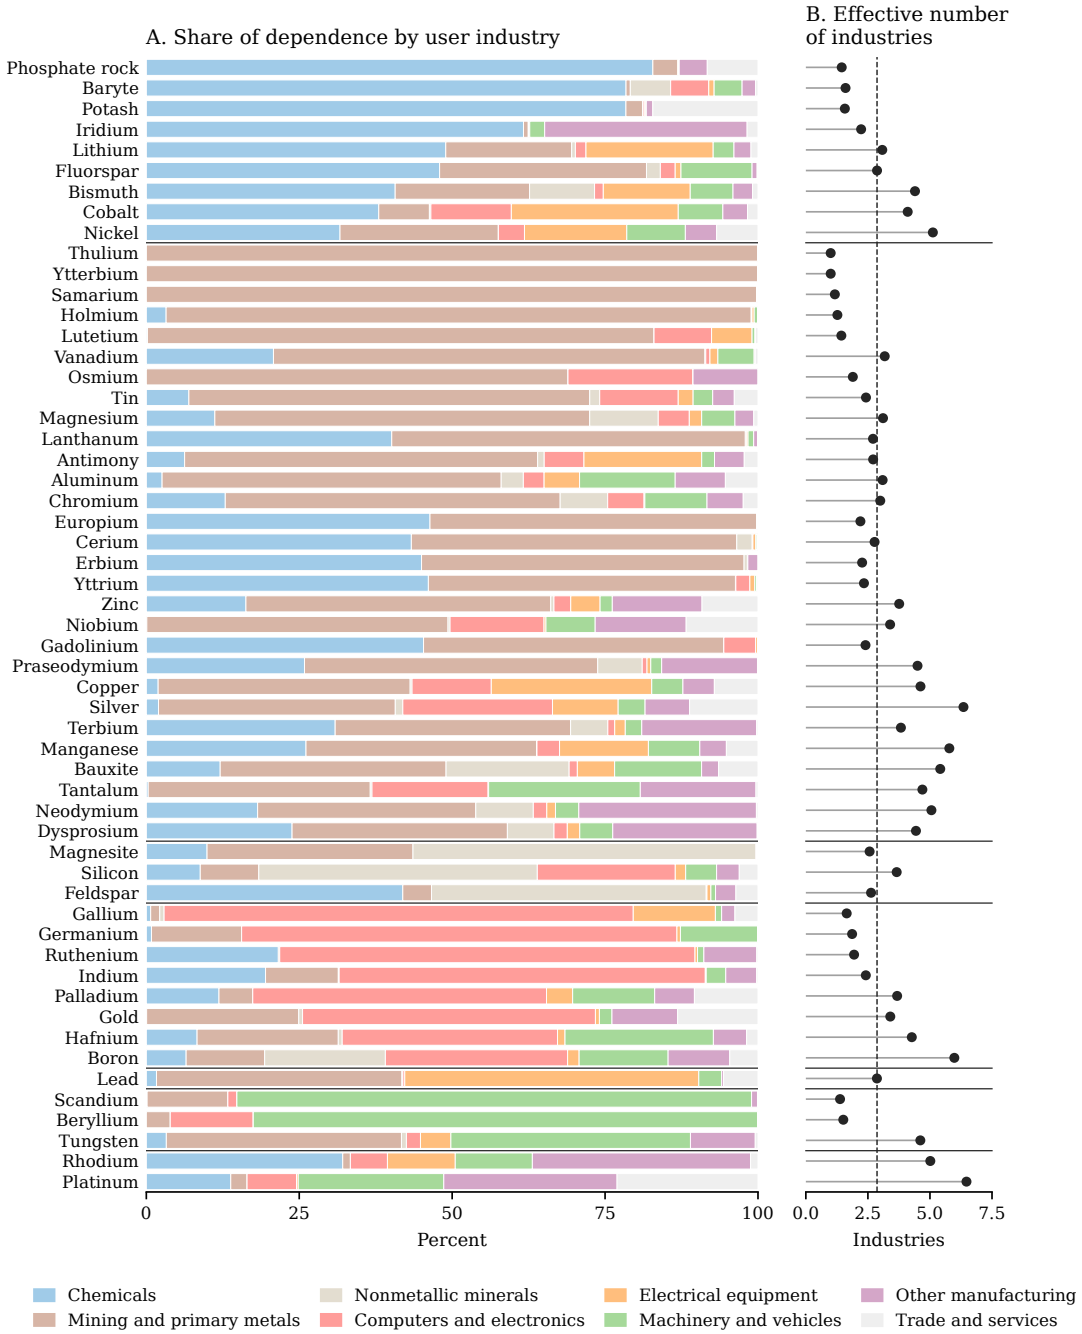
Notes: The entries are increments in adjusted R^2 from nested regressions with material-interacted year, industry-by-year, and firm fixed effects. We assign the remainder to within-firm variation, so the four shares sum to 100 percent. The sample contains user firm-years with a reported industry and country. Ranks are formed within material, and log dependence retains positive values. Because product coefficients and line classifications are fixed, within-firm changes come from revenue shares.

Table A.5: Concentration across Years and Measurement Subsamples

	Firms	Materials	Top share	Top-three share	Effective number
All firms with positive dependence, 2019	3,165	4	74.4%	99.9%	1.6
Listed critical materials only, 2019	3,130	4	79.9%	100.0%	1.5
At least ten materials, 2019	739	15	50.6%	87.0%	2.9
At least 80% of revenue identified, 2019	3,126	4	74.4%	99.8%	1.6
No bound reliance, 2019	2,744	4	73.9%	100.0%	1.7
All firms with positive dependence, 2024	3,948	4	74.3%	99.7%	1.6

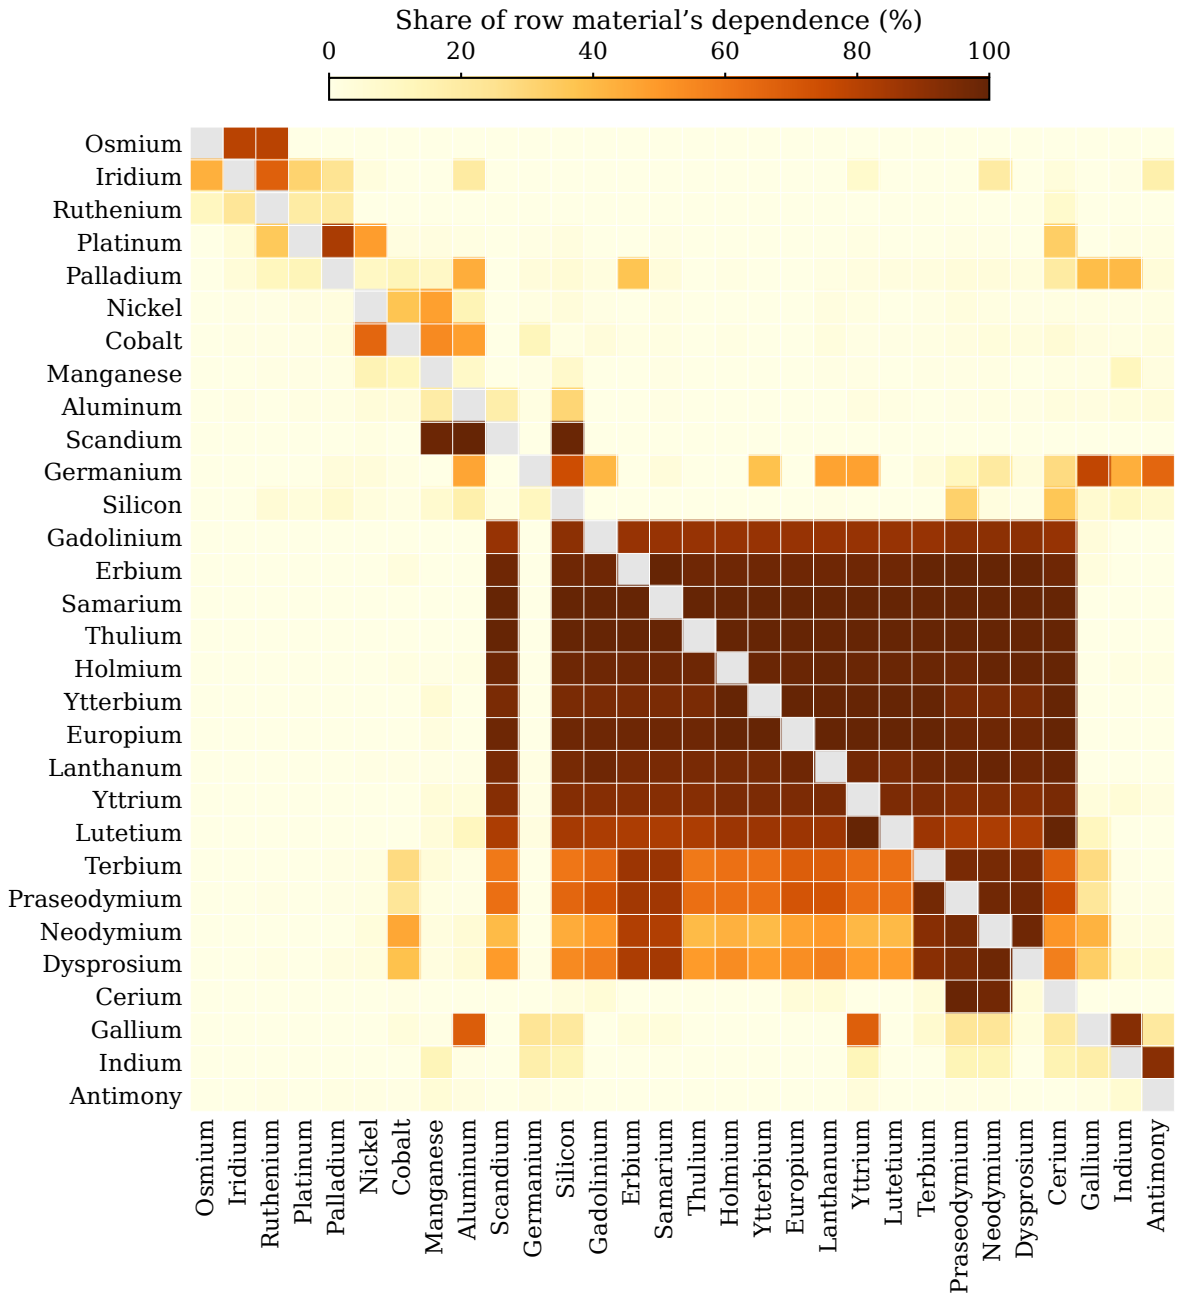
Notes: Entries are cross-sectional medians among firms with positive total dependence. Top shares use each material's share of the total. The effective number is the inverse sum of squared shares. Listed critical materials excludes gold and magnesite. No bound reliance excludes firms whose measured revenue uses a bounded coefficient.

Figure A.1: Material Dependence by User Industry



Notes: Panel A allocates each material's summed dependence across eight groups of firms' reported three-digit NAICS industries, and each bar sums to 100 percent. Mining and primary metals combines NAICS 212 and 331, machinery and vehicles combines 333 and 336, and other manufacturing is the remainder of NAICS 31–33. Panel B reports the inverse Herfindahl index across all 65 industries. The dashed line marks the median. Firms receive equal weight in the 2019 sample of 3,079 users with positive dependence and a reported industry. The large rare-earth shares in other manufacturing mainly reflect animal-feed producers. Rare-earth premixes and compound feed share HS 230990 in the product mapping.

Figure A.2: Material Dependence through Shared Product Categories



Notes: Each cell is the share of the row material's 2019 dependence carried by HS6 categories also mapped to the column material, summed across 3,165 firms with positive dependence. Because each row has its own denominator, opposite cells can differ and row totals can exceed 100 percent. The figure shows the 30 materials with the most links reaching 60 percent in either direction. We use average-linkage clustering based on mean two-way overlap. Shared HS6 categories can contain different goods. Thus, the figure cannot establish that two materials enter the same product.

Table A.6: Market Outcomes by Fiscal Year

	2017	2018	2019	2020	2021	2022	2023	2024	2025
Panel A. Dependent variable: Log market-to-book									
Dependence rank	18.994	6.468	4.649	15.424***	11.869***	10.016**	9.650**	15.462***	14.652***
	(25.021)	(4.569)	(4.530)	(4.874)	(4.311)	(4.120)	(4.114)	(4.166)	(4.440)
Observations	252	3,812	4,678	5,163	6,061	6,631	7,148	7,389	7,316
Within R^2	0.238	0.128	0.107	0.093	0.092	0.079	0.070	0.066	0.060
Panel B. Dependent variable: Realized volatility									
Dependence rank	3.556	1.553*	1.551**	3.391***	3.087***	0.880	1.270*	1.972***	1.266*
	(3.804)	(0.887)	(0.752)	(0.882)	(0.825)	(0.731)	(0.711)	(0.712)	(0.737)
Observations	295	4,093	5,026	5,397	6,419	7,204	7,783	8,084	7,477
Within R^2	0.304	0.193	0.260	0.206	0.176	0.169	0.245	0.279	0.207
Panel C. Dependent variable: Implied volatility									
Dependence rank	–	8.036	14.776	8.601	24.003***	28.790	6.077	−5.935	–
	–	(9.279)	(13.602)	(10.280)	(8.877)	(24.055)	(10.478)	(7.446)	–
Observations	–	144	234	301	369	396	412	380	–
Within R^2	–	0.522	0.521	0.557	0.598	0.527	0.488	0.472	–
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
NAICS6 FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table estimates Panel A of Table 15 separately by fiscal year. The sample, controls, fixed effects, winsorization, and firm-clustered standard errors follow that table. We retain years with at least 100 observations. Realized volatility uses daily returns over the fiscal year and requires 120 trading days. Implied volatility is the 91-day at-the-money OptionMetrics measure and is available through August 2025 for matched U.S. firms. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A.7: Market Outcomes among Firms with Positive Dependence

	Return volatility (1)	Idiosyncratic volatility (2)	Expected shortfall (3)	Market beta (4)	Log market-to-book (5)	Bid-ask spread (6)
Panel A. Dependence rank among firm-years with positive dependence						
Dependence rank	2.666***	1.518**	0.206***	4.529	5.708	5.354***
	(0.834)	(0.617)	(0.077)	(2.810)	(4.884)	(1.657)
Observations	18,150	18,130	18,622	15,813	18,312	18,357
Firms	3,497	3,498	3,550	3,211	3,459	3,533
Within R^2	0.057	0.217	0.193	0.023	0.098	0.215
Panel B. Rank of the largest single-material dependence among firm-years with positive dependence						
Largest material rank	2.431***	1.271**	0.188**	5.281*	3.530	4.958***
	(0.831)	(0.616)	(0.077)	(2.820)	(4.878)	(1.655)
Observations	18,150	18,130	18,622	15,813	18,312	18,357
Firms	3,497	3,498	3,550	3,211	3,459	3,533
Within R^2	0.057	0.217	0.193	0.024	0.097	0.215
Panel C. Log total dependence, firm-years with positive dependence						
Log total dependence	0.232***	0.132**	0.016**	0.370	0.704	0.456***
	(0.082)	(0.062)	(0.008)	(0.265)	(0.478)	(0.156)
Observations	18,150	18,130	18,622	15,813	18,312	18,357
Firms	3,497	3,498	3,550	3,211	3,459	3,533
Within R^2	0.057	0.217	0.193	0.023	0.098	0.215
Firm controls	Yes	Yes	Yes	Yes	Yes	Yes
NAICS6 $\times$ year FE	Yes	Yes	Yes	Yes	Yes	Yes
Country $\times$ year FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table re-estimates Table 15 among firm-years with positive dependence. Panels A and B re-form the two ranks within this sample. Panel C uses log total dependence. Outcomes, controls, fixed effects, units, winsorization, and firm-clustered standard errors follow the main table. ***, **, and * denote significance at the 1, 5, and 10 percent levels.

Table A.8: Material Cost Pressure, Operations and Financing: Estimates

		Before	Quarters after the price change				
		$t-9$ to $t-5$	0	2	4	6	8
Panel A. Operations							
(1)	Sales	−0.335	0.309***	0.625***	0.696***	0.711***	0.607**
		(0.210)	(0.087)	(0.125)	(0.193)	(0.238)	(0.269)
	Observations	136,217	151,547	138,071	125,272	112,411	99,678
(2)	COGS/sales	−0.155***	0.012	0.161**	0.285***	0.314***	0.258***
		(0.050)	(0.053)	(0.065)	(0.068)	(0.070)	(0.075)
	Observations	134,778	149,871	136,517	123,852	111,124	98,538
(3)	Operating income	−0.039	−0.003	−0.020	−0.025	−0.052	−0.114
		(0.034)	(0.009)	(0.029)	(0.048)	(0.070)	(0.079)
	Observations	109,635	115,681	104,746	93,979	83,174	72,751
(4)	Net income	0.034	−0.008	−0.044	−0.058	−0.075	−0.159**
		(0.032)	(0.009)	(0.027)	(0.043)	(0.056)	(0.071)
	Observations	133,166	152,859	138,171	123,772	109,470	95,594
(5)	Operating cash flow	0.113**	−0.052***	−0.135***	−0.208***	−0.257***	−0.266***
		(0.044)	(0.013)	(0.031)	(0.053)	(0.073)	(0.082)
	Observations	132,394	151,487	136,597	122,119	107,819	93,995
(6)	Inventories	−0.080**	0.059***	0.108***	0.119***	0.096**	0.042
		(0.034)	(0.009)	(0.021)	(0.037)	(0.042)	(0.051)
	Observations	132,218	145,532	132,437	120,026	107,813	95,640
(7)	Accounts receivable	−0.093**	0.047***	0.112***	0.143***	0.141***	0.097
		(0.034)	(0.010)	(0.026)	(0.036)	(0.047)	(0.064)
	Observations	138,043	151,759	138,383	125,496	112,752	100,020
(8)	Capital expenditure	−0.050*	0.009	0.041*	0.083**	0.131**	0.130**
		(0.025)	(0.006)	(0.020)	(0.037)	(0.050)	(0.060)
	Observations	119,870	145,262	125,635	108,546	93,020	78,870
(9)	Cash	−0.008	0.000	−0.011	−0.010	−0.036	−0.112
		(0.082)	(0.015)	(0.039)	(0.059)	(0.076)	(0.103)
	Observations	138,941	152,814	139,351	126,393	113,564	100,761
Panel B. Financing							
(10)	Net financing	−0.140*	0.048**	0.153**	0.287***	0.369***	0.374**
		(0.074)	(0.022)	(0.058)	(0.089)	(0.113)	(0.141)
	Observations	125,363	148,244	131,463	116,542	102,109	88,511
(11)	Accounts payable	−0.089**	0.050***	0.118***	0.175***	0.170***	0.143**
		(0.034)	(0.010)	(0.022)	(0.036)	(0.047)	(0.062)
	Observations	135,877	149,085	135,900	123,235	110,719	98,219
(12)	Short-term debt	−0.095**	0.020**	0.048*	0.098**	0.114**	0.127**
		(0.038)	(0.008)	(0.025)	(0.044)	(0.053)	(0.059)
	Observations	121,900	140,719	127,332	114,892	102,680	90,575
(13)	Long-term debt	−0.066	0.016	0.085**	0.136**	0.131*	0.130
		(0.049)	(0.011)	(0.032)	(0.055)	(0.069)	(0.076)
	Observations	108,945	133,417	119,717	107,383	95,591	83,917
(14)	Total debt	−0.181***	0.043***	0.125***	0.230***	0.254***	0.272***
		(0.056)	(0.013)	(0.040)	(0.065)	(0.081)	(0.087)
	Observations	126,313	145,053	131,573	118,896	106,412	93,979
(15)	Interest rate on debt	0.042	0.003	0.006	−0.040	−0.049	−0.025
		(0.027)	(0.015)	(0.020)	(0.026)	(0.031)	(0.037)
	Observations	108,012	131,958	118,454	106,635	95,086	83,773

Notes: This table reports Equation (16) for 158,068 firm-quarters at 7,642 firms over 2018Q2–2025Q4. Coefficients are per percentage point of cost pressure, whose standard deviation is 1.19 percentage points of identified revenue. Columns $t+h$ measure changes from $t-1$ or cumulative flows from t through $t+h$. Before uses $t-9$ to $t-5$. Appendix Table A.1 defines the outcomes and units. Price availability determines the sample at each horizon. Standard errors are clustered by firm and quarter. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.9: Supply-Driven Cost Pressure under the Broad-Window Definition

		Supply-driven			Other		
	Quarters after the price change	0	2	4	0	2	4
(1)	Sales	0.603^{***} (0.147)	1.446^{***} (0.433)	1.464^{**} (0.559)	0.231^{***} (0.068)	0.437^{***} (0.094)	0.468^{***} (0.140)
(2)	Cost-of-sales ratio	0.087 (0.114)	0.180 (0.174)	0.277 (0.227)	-0.009 (0.039)	0.094^{*} (0.048)	0.175^{***} (0.046)
(3)	Operating cash flow	-0.083^{***} (0.028)	-0.270^{***} (0.071)	-0.504^{***} (0.154)	-0.034^{***} (0.009)	-0.092^{***} (0.022)	-0.143^{***} (0.038)
(4)	Capital expenditure	0.024 (0.016)	0.115^{**} (0.051)	0.224^{**} (0.094)	0.005 (0.004)	0.022 (0.015)	0.046^{*} (0.026)
(5)	Net financing	0.142^{*} (0.080)	0.390^{**} (0.188)	0.834^{***} (0.277)	0.031^{**} (0.013)	0.097^{**} (0.038)	0.179^{***} (0.061)
(6)	Short-term debt	0.052^{*} (0.026)	0.127^{*} (0.063)	0.335^{***} (0.110)	0.017^{**} (0.007)	0.043^{**} (0.021)	0.055 (0.033)
(7)	Long-term debt	0.077^{*} (0.038)	0.247^{**} (0.107)	0.406^{**} (0.149)	0.006 (0.007)	0.046^{**} (0.021)	0.079^{*} (0.040)
(8)	Stock return	0.357^{*} (0.188)	0.690^{*} (0.366)	0.589 (0.696)	0.108^{*} (0.060)	0.162 (0.115)	0.142 (0.176)

Notes: This table reports the broad-window alternative to Table 10. A material's full four-quarter price change counts as supply-driven when the event list contains a tightening for the material in the quarter of the change or the one before. The remaining four-quarter changes form the other component. Both components enter one regression with the controls, fixed effects, winsorizing and scale of Equation (16), and outcomes are read at quarters 0, 2, and 4. Standard errors are clustered by firm and quarter. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A.10: Supply-Tightening Event List

Date	Type	Materials ($\Delta \ln p$, %)	Event
2017-09	Mine	Potash (–)	Allan and Lanigan / PotashCorp
2018-04	Mine	Li (-54)	Olaroz
2018-08	Mine	Pt (-9)	Rustenburg Impala: Rustenburg South Africa platinum restructuring ann…
2018-08	Mine	Mn (+5)	Tshipi Borwa / Jupiter
2018-11	Mine	Co (-70)	Kamoto: Kamoto DRC cobalt sales/exports suspended November 6 due uran…
2019-01	Mine	Li (-9)	SQM Salar de Atacama: Approved compliance programme restricts excess…
2019-06	Mine	Li (-22)	Pilgangoora: Company announces June-July concentrator downtime/output…
2019-08	Mine	Co (+8)	Mutanda year-end suspension reported
2019-08	Policy	Ni (+18)	Indonesia: Nickel ore export ban moved forward by two years
2020-03	Mine	Nd (-1), Pr (-13)	Lynas Malaysia processing Mt Weld feed
2020-03	Mine	Li (-22)	Olaroz / Orocobre
2021-04	Mine	V (+15)	Vametco
2021-06	Policy	Al (+25), Ni (+9)	Russian Federation: Introduction of export duties on metals (June 2021)
2022-03	Policy	Al (+1), Bauxite (+8)	Australia: Government bans exports on aluminium to Russia
2022-05	Policy	Li (+10)	Argentina: Government establishes export criterion value on lithium c…
2022-12	Policy	Li (-37)	Zimbabwe: Export ban on unprocessed lithium
2023-01	Mine	Sn (+19)	San Rafael suspends tin mining amid regional unrest
2023-07	Policy	Ga (+11), Ge (+2)	China: Government announces export control measures for gallium and g…
2024-08	Policy	Sb (+97)	China: Government announces export control measures for antimony and…
2024-10	Mine	V (-1)	Vametco funding-related production slowdown
2024-12	Policy	Sb (+78), Ga (-35), Ge (-5)	China: Government adopts further export controls on dual-use items to…
2025-02	Policy	Bi (+56), In (-2), W (+13)	China: Government announces export control measures for various items…
2025-02	Policy	Co (+33)	Democratic Republic of Congo: Cobalt exports suspended for four months
2025-03	Mine	Sn (+7)	Bisie / Alphamin
2025-04	Policy	7 rare earths (-3 to +11)	China: Government announces export control measures for various items…
2025-04	Mine	Sn (+6)	Bisie / Alphamin
2025-09	Mine	Ge (-9)	Guangxi Yusheng halts germanium production after feed-mine safety rec…
2025-10	Policy	12 rare earths (–)	China: Government announces export controls for foreign entities on r…

Notes: The list contains supply-tightening events identified for the priced materials over 2017–2025: export-policy acts screened from the GTA record and mine disruptions screened from a fixed corpus of mining news, each dated to its announcement and reviewed against its sources. The materials column reports the two-quarter log price change around each event, $100\,[\ln p(e{+}1) - \ln p(e{-}1)]$, from the quarter before the event to the quarter after. A dash marks a change outside the sample window. Events on several rare earths report the range across the affected elements.

B Proofs and Extensions

This appendix proves the results in Section 2 and extends the framework to quantity equilibrium and alternative allocation rules.

B.1 First-Order Results

Lemma 1: cost exposure. Under constant returns, Shephard's lemma gives

$$\frac{\partial \ln MC_j}{\partial \ln p_i} = \frac{p_i d_{ij}}{MC_j} = s_{ij}.$$

Constant markups imply $d \ln p_i = d \ln MC_i$. For $j \neq o$, the chain rule yields $\beta_{mj} = \sum_i s_{ij}\beta_{mi}$. At the origin, $p_o = p_m$ gives $\beta_{mo} = 1$, while unexposed inputs have zero exposure. Iterating the recursion gives Equation (3). □

Heterogeneous producers. Suppose product j is supplied by producers h with marginal costs MC_h, dependence $\beta_{mh} = \partial \ln MC_h / \partial \ln p_m$, and a common fixed markup μ_j, so that $p_h = \mu_j MC_h$. Buyers combine the varieties through a CES aggregator with elasticity $\eta > 1$ and price index $P_j = (\sum_h a_h p_h^{1-\eta})^{1/(1-\eta)}$, so that $q_h \propto a_h p_h^{-\eta}$. Differentiating the index gives

$$\frac{\partial \ln P_j}{\partial \ln p_m} = \sum_h \omega_h \beta_{mh}, \qquad \omega_h = \frac{p_h q_h}{\sum_{h'} p_{h'} q_{h'}},$$

the revenue-weighted average of producer dependence. For an input i with requirement d_{ih} per unit at producer h, Shephard's lemma gives the cost share $s_{ih} = \mu_j p_i d_{ih} / p_h$, and its revenue-weighted average is

$$\sum_h \omega_h s_{ih} = \mu_j p_i \frac{\sum_h q_h d_{ih}}{\sum_h p_h q_h} = \frac{\mu_j p_i \bar{d}_{ij}}{\bar{p}_j},$$

where $\bar{d}_{ij}$ is the quantity-weighted average requirement and $\bar{p}_j$ the average unit value, revenue divided by quantity. Baseline average prices and requirements therefore recover the index's expenditure shares. Away from the initial prices, the unit value reweights producers as quantities shift toward cheaper varieties, so its change need not track the index. □

Proposition 2: firm dependence and profit exposure. At fixed product-line outputs, $dC_f = \sum_r q_r MC_r \beta_{mr}\, d \ln p_m$. Because $p_r = \mu_r MC_r$ and $Y_r = p_r q_r$,

$$\frac{dC_f}{Y_f} = \sum_r \omega_r \frac{\beta_{mr}}{\mu_r}\, d \ln p_m = D_{mf}\, d \ln p_m,$$

proving Equation (6). For the profit result, let $\mathbf{p}_f$ collect the firm's product-line prices and let $q_r(\mathbf{p}_f)$ be line r's demand, allowing demand across lines to interact. The firm chooses $\mathbf{p}_f$ to maximize

$$\Pi_f = \sum_r (p_r - MC_r) q_r(\mathbf{p}_f).$$

At the initial optimum, the first-order conditions eliminate the effect of induced price adjustments. Holding the demand system fixed, the envelope theorem therefore gives

$$d\Pi_f = -\sum_r q_r \, dMC_r = -Y_f D_{mf} \, d\ln p_m.$$

Dividing by Y_f proves the profit result. A direct shift in product demand would add a separate term. □

Proposition 3: supply disruptions. For part (i), let $X_h(p_m)$ be user h's derived material demand, $X_m = \sum_h X_h$ aggregate demand, $\varepsilon_h = -d\ln X_h / d\ln p_m$, and $\vartheta_h = X_h / X_m$ its share of material demand. Market clearing gives

$$-\delta = d\ln X_m = -\Big(\sum_h \vartheta_h \varepsilon_h\Big) d\ln p_m = -\bar{\varepsilon}_m \, d\ln p_m.$$

Combining $d\ln p_m = \delta / \bar{\varepsilon}_m$ with Lemma 1 proves Equation (7).

For part (ii), benchmark-material supply falls by $d\ln q_o = -\delta$. Let X_{ij} denote total deliveries of input i to producer j. At a downstream stage's initial cost minimum, $p_i = MC_j f_i$, where f_i is input i's marginal product. Constant returns give $C_j = MC_j q_j$, so the output elasticity $X_{ij} f_i / q_j$ equals the cost share s_{ij}. With differentiable downstream production, fixed other inputs, and proportional rationing,

$$d\ln q_j = \sum_i s_{ij} \, d\ln X_{ij} = \sum_i s_{ij} \, d\ln q_i.$$

This is the cost-exposure recursion initialized at $-\delta$, giving $d\ln q_j = -\beta_{mj}\delta$.

In the Leontief limit, write $r_j = q_j / q_j^0$. With other inputs fixed and every required input binding initially, proportional rationing gives $r_j = \min\{1, \min_i r_i\}$. Starting from $r_o = e^{-\delta}$, induction along an acyclic network gives $r_j = e^{-\delta}$ for every product connected to the origin through a required input, equivalently every j with $\beta_{mj} > 0$, which proves the Leontief result. □

B.2 Quantity Equilibrium under Price Allocation

This formulation nests the standard CES production-network environment, in which CES technologies determine the conditional input demands $a_{ij}(\mathbf{p})$ while CES final demand determines $y_j(\mathbf{p}, E)$. We use the more general notation because the separa-

tion of prices from equilibrium quantities requires only constant returns and fixed markups.

Holding other-input prices and aggregate expenditure fixed, let Q_j denote gross output of product j and $y_j(\mathbf{p}, E)$ its final demand at product prices $\mathbf{p}$ and expenditure E. Write $a_{ij}(\mathbf{p}) = d_{ij}(\mathbf{p})$ for the cost-minimizing quantity of input i per unit of product j, setting $a_{ij} = 0$ when $i \notin \mathcal{I}_j$. Constant returns imply that producer j's total demand for input i is

$$X_{ij} = a_{ij}(\mathbf{p})Q_j.$$

Product-market clearing therefore requires

$$Q_i = y_i(\mathbf{p}, E) + \sum_j a_{ij}(\mathbf{p})Q_j. \tag{20}$$

Let $\mathbf{A}(\mathbf{p}) = [a_{ij}(\mathbf{p})]$. In vector form,

$$\mathbf{Q} = [\mathbf{I} - \mathbf{A}(\mathbf{p})]^{-1}\mathbf{y}(\mathbf{p}, E). \tag{21}$$

Because the production network is acyclic, products can be ordered so that $\mathbf{A}$ is triangular and nilpotent. The inverse in Equation (21) is therefore a finite sum, $\mathbf{I} + \mathbf{A} + \cdots + \mathbf{A}^H$, where H is the maximum path length.

Given p_m, the price system remains

$$p_o = p_m, \qquad p_j = \mu_j MC_j(\mathbf{w}_j, \mathbf{p}_{\mathcal{I}_j}), \quad j \neq o.$$

Under constant returns and fixed markups, this system is independent of $\mathbf{Q}$. We can solve prices and unit costs upstream to downstream, then use Equation (21) to determine gross outputs. Equilibrium scale therefore leaves the cost-exposure recursion, physical-requirement decomposition, finite-price cost bounds and firm-level fixed-output identity unchanged. In the demand problem of Section 2.3, the envelope theorem removes the induced adjustment of q_r from the first-order profit response.

Quantities respond through additional channels. Differentiating Equation (20) gives

$$d\mathbf{Q} = [\mathbf{I} - \mathbf{A}]^{-1} [d\mathbf{y} + (d\mathbf{A})\mathbf{Q}]. \tag{22}$$

The first term captures final-demand responses to prices, and the second captures substitution in unit input requirements. For example, if $y_j = A_j p_j^{-\eta_j}$ and A_j is fixed, then $d \ln y_j = -\eta_j \beta_{mj} \, d \ln p_m$. Material dependence therefore governs the unit-cost response, while the equilibrium response of gross output also depends on demand and substitution.

For price allocation, $Q_o(p_m)$ is aggregate equilibrium demand for the benchmark

material implied by Equations (20)–(21). Market clearing sets material supply equal to $Q_o(p_m)$. Defining $\bar{\varepsilon}_m = -d\ln Q_o / d\ln p_m$, a log supply cut δ gives $d\ln p_m = \delta/\bar{\varepsilon}_m$, as in Proposition 3(i). The quantity system therefore enters that result only through $\bar{\varepsilon}_m$.

If the material price does not clear the market, demand can exceed supply. Equations (20)–(22) then require an allocation rule to determine deliveries to intermediate and final users.

Proposition 3(ii) imposes proportional delivery cuts. Its recursion gives feasible output, which equals realized output when the rationed input binds at each affected stage. Other allocation rules can change output responses and their ranking by material dependence (Appendix B.4), while leaving the cost-exposure results intact.

B.3 Second-Order Responses and Finite-Shock Bounds

This subsection derives Proposition 4 and Corollary 1 under the multiple-input CES technology in Section 2.4. All input quantities are initially positive and cost-minimizing, and $p_o = p_m$ at the origin. The production function at stage k is

$$q_k = \left[\sum_{i\in\mathcal{I}_k} \alpha_{ik} x_{ik}^{\rho_k} + \alpha_{uk} u_k^{\rho_k}\right]^{1/\rho_k}, \qquad \rho_k = \frac{\sigma_k - 1}{\sigma_k}, \qquad 0 < \sigma_k < 1,$$

where x_{ik} denotes material-dependent input i and u_k is the unexposed composite of all other inputs, including labor and capital. Cost minimization yields the dual marginal-cost function in Section 2.4. Inputs are substitutable for every $\sigma_k > 0$ and enter in fixed proportions at the Leontief limit $\sigma_k \to 0$.

A network with several paths. Let $t = \ln p_m$, $\beta_{mi}(t) = d\ln p_i/dt$, and $H_i(t) = d^2\ln p_i/dt^2$. At $t = 0$, $\beta_{mi}(0)$ is the initial dependence β_{mi} used in the main text. At node j, include all inputs in the shares $s_{ij}(t)$, assigning zero exposure to the unexposed composite. Constant markups imply $d\ln p_j = d\ln MC_j$. Suppressing the argument t, differentiating the CES shares gives $ds_{ij}/dt = (1-\sigma_j)s_{ij}(\beta_{mi} - \beta_{mj})$. Differentiating $\beta_{mj} = \sum_i s_{ij}\beta_{mi}$ then gives

$$H_j = (1-\sigma_j)\underbrace{\left(\sum_i s_{ij}\beta_{mi}^2 - \beta_{mj}^2\right)}_{\mathrm{Var}_j^s(\beta_{mi})} + \sum_i s_{ij} H_i. \tag{23}$$

Curvature is generated by differences in exposure across a node's inputs and inherited from upstream nodes. Equal input exposures produce no local variance term. Baqaee and Farhi (2019) identify a related mechanism for aggregate-output curva-

ture. Here it governs a product's cost response when markups and unexposed-input prices are fixed.

Let $\Gamma_k^{\text{cost}} = (1-\sigma_k)\,\text{Var}_k^s(\beta_{mi})$, and let β_{kj} sum initial cost-share products along paths from k to j, with $\beta_{jj} = 1$. Starting from $\beta_{mo} = 1$ and $H_o = 0$, the recursion yields

$$\Delta \ln MC_j = \beta_{mj} z + \frac{z^2}{2} \sum_{k \neq o} \beta_{kj} \Gamma_k^{\text{cost}} + o(z^2), \qquad z = \Delta \ln p_m. \tag{24}$$

Thus $A_{mj} = \sum_{k \neq o} \beta_{kj} \Gamma_k^{\text{cost}}$.

Quantity curvature on the network. Under proportional rationing, write $y_j = -\Delta \ln q_j$ for the supply-constrained output response, with $y_o = \delta$. Include the unexposed composite among each downstream node's inputs, with $y_u = 0$ and $\beta_{mu} = 0$. Let s_{ij} denote initial cost shares, which sum to one over all inputs, and let $\rho_j = (\sigma_j - 1)/\sigma_j$. The exact CES quantity response is

$$y_j = -\frac{1}{\rho_j} \ln\left(\sum_i s_{ij} e^{-\rho_j y_i} \right).$$

Expanding at the initial allocation gives $y_j = \sum_i s_{ij} y_i - \frac{1}{2}\rho_j \,\text{Var}_j^s(y_i) + o(\|\mathbf{y}\|^2)$. Substituting $y_i = \beta_{mi}\delta + \frac{1}{2} B_{mi}\delta^2 + o(\delta^2)$ yields

$$B_{mj} = \sum_i s_{ij} B_{mi} + \Gamma_j^{\text{output}}, \qquad \Gamma_j^{\text{output}} = \frac{1-\sigma_j}{\sigma_j} \left(\sum_i s_{ij} \beta_{mi}^2 - \beta_{mj}^2 \right),$$

with $B_{mo} = 0$ and zero curvature for fixed other inputs. Iterating on the acyclic network gives

$$B_{mj} = \sum_{k \neq o} \beta_{kj} \Gamma_k^{\text{output}}.$$

For $0 < \sigma_k < 1$, both local curvature terms are nonnegative, so $A_{mj}, B_{mj} \geq 0$. Both are strictly positive whenever a stage with positive transmission weight β_{kj} combines inputs with unequal exposures and positive shares, which proves Proposition 4. □

Corollary 1: finite-shock cost bound. Let p_m^0 be the initial material price. Suppressing the material subscript m within this proof, define

$$g_j(t) = \ln \frac{MC_j(p_m^0 e^t)}{MC_j(p_m^0)}, \qquad \beta_j(t) = g_j'(t), \qquad H_j(t) = g_j''(t).$$

The cost functions incorporate upstream price adjustment while holding markups and unexposed-input prices fixed. Shares and derivatives are evaluated at t, with

$\beta_j(0) = \beta_{mj}$ and $H_j(0) = A_{mj}$. At the origin,

$$g_o(t) = t, \qquad \beta_o(t) = 1, \qquad H_o(t) = 0.$$

Unexposed inputs have zero exposure and curvature. For each downstream node, the recursion in Equation (23) holds at every t:

$$\beta_j(t) = \sum_i s_{ij}(t)\beta_i(t), \qquad H_j(t) = \sum_i s_{ij}(t)H_i(t) + (1-\sigma_j)\operatorname{Var}_j^{s(t)}(\beta_i(t)).$$

Shares sum to one over all inputs, including the unexposed composite. Induction on the acyclic network therefore gives $0 \le \beta_j(t) \le 1$ and $H_j(t) \ge 0$. If $H_i(t) \le \beta_i(t)[1-\beta_i(t)]$ for upstream inputs, then, suppressing t,

$$\begin{aligned} H_j &\le \sum_i s_{ij}\beta_i(1-\beta_i) + \operatorname{Var}_j^{s}(\beta_i) \\ &= \sum_i s_{ij}\beta_i - \left(\sum_i s_{ij}\beta_i\right)^2 = \beta_j(1-\beta_j). \end{aligned}$$

The inequality uses $0 \le 1-\sigma_j \le 1$. The origin satisfies the bound with equality, so induction proves $0 \le H_j(t) \le \beta_j(t)[1-\beta_j(t)]$ at every price on the adjustment path. At $t = 0$, this gives the stated bound on A_{mj}.

For $\beta = \beta_{mj} \in (0,1)$, the differential inequality $0 \le \beta_j'(t) \le \beta_j(t)[1-\beta_j(t)]$ implies, for $t \ge 0$,

$$\beta \le \beta_j(t) \le \frac{\beta e^t}{1-\beta+\beta e^t}.$$

The upper inequality follows by integrating the derivative of $\ln[\beta_j(t)/(1-\beta_j(t))]$, which is at most one. Comparison with this logistic solution also rules out reaching one in finite time. Integrating the two bounds from zero to z gives

$$\beta z \le g_j(z) \le \ln(1-\beta+\beta e^z),$$

proving Equation (11). If $\beta = 0$, the inequality $\beta_j' \le \beta_j$ and the zero initial value imply $\beta_j(t) = 0$, hence $g_j(t) = 0$. If $\beta = 1$, monotonicity and $\beta_j(t) \le 1$ imply $g_j(t) = t$.

For products satisfying the ranking condition, these bounds give $g_j(z) \ge \beta_{mj} z > \ln[1+\beta_{mj'}(e^z-1)] \ge g_{j'}(z)$, proving the stated ordering. When the bounds overlap, this argument yields no ordering without the stage elasticities.

For $\beta > 0$ and $x = e^z - 1 > 0$, the exact log-cost response is positive. Its relative approximation error obeys

$$0 \le \frac{g_j(z)-\beta z}{g_j(z)} \le 1 - \frac{\beta\ln(1+x)}{\ln(1+\beta x)} \le 1 - \frac{\ln(1+x)}{x},$$

where the last inequality uses $\ln(1+\beta x) \leq \beta x$. This proves the conservative error bound in the main text, while the middle expression is tighter when initial exposure is known. Evaluating the last expression at $x = 0.1$ and $x = 0.2$ gives $0.046898\ldots$ and $0.088392\ldots$, respectively. Exponentiating gives the corresponding bounds on the level of marginal cost.

The bounds are tight. As every $\sigma_k \to 1$, $\beta_j(t) \equiv \beta_{mj}$ and the lower bound is attained. As every $\sigma_k \to 0$, $MC_j(p_m^0 e^z)/MC_j(p_m^0) = 1 + \beta_{mj}(e^z - 1)$ and the upper bound is attained. The interval is therefore sharp over the admissible substitution elasticities at fixed initial shares. □

Firm-level finite-shock bounds. The product bounds can be aggregated using the same portfolio weights as firm dependence. Hold each product line's output fixed, let $g_r(z)$ denote its log marginal-cost response, and define $\Delta C_f(z) = C_f(z) - C_f(0)$. Line r's initial variable cost relative to firm revenue is ω_r/μ_r. Therefore

$$\sum_r \frac{\omega_r}{\mu_r}\left(e^{\beta_{mr}z} - 1\right) \leq \frac{\Delta C_f(z)}{Y_f} \leq \sum_r \frac{\omega_r}{\mu_r}\beta_{mr}(e^z - 1) = D_{mf}(e^z - 1). \tag{25}$$

The left side sums the lower product bounds, and the right side sums the upper bounds using Equation (5). Firm f therefore has a larger fixed-output cost increase than firm f' whenever its lower bound exceeds the upper bound for f'. Otherwise, $D_{mf} > D_{mf'}$ ranks their local responses. Their responses to a finite shock can be ordered differently.

B.4 Allocation Across Buyers

Allocation determines how much of an aggregate shortage reaches each producer. Here q denotes supply-constrained output, as in Proposition 3(ii). For a local quantity response, let $L_{ij,e} = -\Delta \ln X_{ij,e}$ and $y_{i,e} = -\Delta \ln q_{i,e}$. Suppose the allocation rule satisfies $L_{ij,e} = r_{ij,e} y_{i,e}$ to first order, with $r_{ij,e} \geq 0$, while other inputs remain fixed. Proportional rationing has $r_{ij,e} = 1$, while supply protection or priority can make these factors smaller or larger.

For path ℓ from o to j, let $s_{k,\ell}$ and $r_{k,\ell,e}$ denote the cost share and allocation factor at step k. Its cost-exposure contribution is $\beta_{mj,\ell} = \prod_{k=1}^{h_\ell} s_{k,\ell}$, and its allocation factor is $R_{mj,\ell,e} = \prod_{k=1}^{h_\ell} r_{k,\ell,e}$. On an acyclic network, iterating $y_{j,e} = \sum_i s_{ij} r_{ij,e} y_{i,e}$ from $y_{o,e} = \delta_{m,e}$ gives

$$-\Delta \ln q_{j,e} \approx \delta_{m,e} \sum_{\ell: o \rightsquigarrow j} \beta_{mj,\ell} R_{mj,\ell,e}. \tag{26}$$

For $\beta_{mj} > 0$, define

$$R_{m,j,e} = \frac{\sum_\ell \beta_{mj,\ell} R_{mj,\ell,e}}{\beta_{mj}}. \tag{27}$$

Then

$$-\Delta \ln q_{j,e} \approx \beta_{mj} R_{m,j,e} \delta_{m,e}. \tag{28}$$

Contracts, inventories, and sourcing affect allocation. β_{mj} records production requirements, while R summarizes the local allocation rule.

Cross-sectional implications. For a common material and event, Equation (28) orders local output losses by $\beta_{mj} R_{m,j,e}$. Let Z_j denote product controls. If the allocation factor is conditionally mean-independent of dependence, $\mathbb{E}[R_{m,j,e} \mid \beta_{mj}, Z_j] = \bar{R}_{m,e}$, then

$$\mathbb{E}[-\Delta \ln q_{j,e} \mid \beta_{mj}, Z_j] \approx \bar{R}_{m,e} \beta_{mj} \delta_{m,e}. \tag{29}$$

Mean independence therefore preserves the conditional-mean ranking. Individual pairs can still be ordered differently. Systematic protection of more exposed products can attenuate or reverse the ranking, whereas greater cuts to those products can amplify it.

C Production-Network Construction

Figure C.1 expands the four steps in Figure 1. It separates construction of the source-grounded trunk from customs classification and the downstream extension. The figure also shows where review, evidence grades, and recycling returns enter. Internet Appendix IA1 details the trunk and customs classification. IA2 covers the downstream search and final network assembly.

From the trunk to the assembled network. The empirical measure uses the frozen trunk, the version fixed for customs classification and downstream extension before the final source audit was complete. It contains 2,763 source products and 3,835 relations. Customs classification can divide a source product into several physical product groups. Each group becomes a node, and the original relations are restated between the relevant groups. This produces 3,829 trunk nodes. The downstream extension adds 3,019 supplementary nodes, giving 6,848 nodes across the 55 networks.[54]

The relation and edge counts differ for the same reason. Of the 3,835 trunk relations, 94 are intakes from a source outside the mapped chain and 198 return to an earlier production stage. The remaining 3,543 forward relations become 4,821 edges once they are assigned to the relevant product groups. Network assembly adds 143 links between groups split from the same product and 4,174 supplementary connections. The result is 9,138 forward edges.

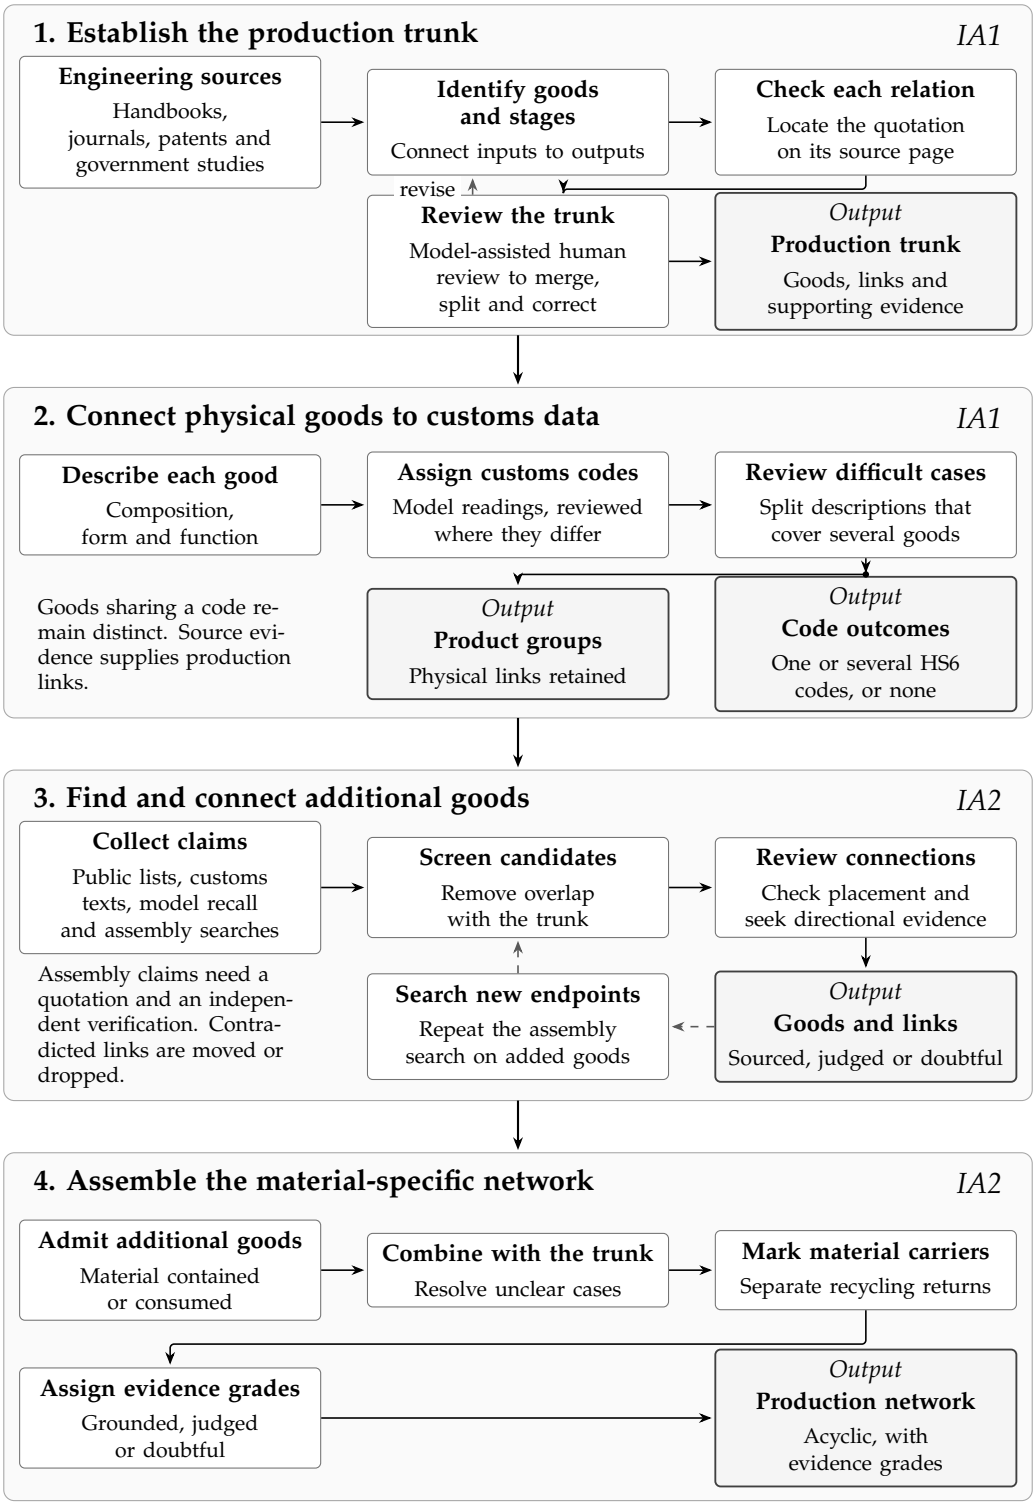
[54] A later review merged duplicate records and restated relationships against their source pages, leaving 2,586 goods and 3,484 relations. Of the 16,747 quotations in this reviewed trunk, 16,613 were confirmed on the cited pages. Because classification and downstream construction use the frozen trunk, IA1 reports the two versions separately.

Evidence grades. Each forward edge carries an evidence grade. Edges inherited from the trunk and links between groups of the same product are grounded. A supplementary connection is also grounded when a source passage states the direction of the physical relationship. Where no such passage is available, we classify the connection as judged if the good's standard recipe or established production practice makes the input its usual route for the material. A doubtful connection reaches the right material through what is likely the wrong form or processing step. The headline measure retains all three grades: 8,036 edges are grounded, 1,055 judged and 47 doubtful. Only 14 nodes require a doubtful edge on their best-supported route, and together they account for less than 0.1 percent of summed product-level dependence. Table C.1 reports the evidence composition separately for each material.

Intakes, returns and material carriers. An intake marks where material enters from a source outside the mapped trunk. Because the mapped trunk ends at that source, we record the intake separately from the edges. A return records a relationship that points back to an earlier stage, as recycling does. Product grouping expands the frozen trunk's 198 return relations into 280 trunk returns, and secondary production routes add another 36. We retain these 316 returns separately from the forward graph. As a result, the graph remains acyclic as required by the model.

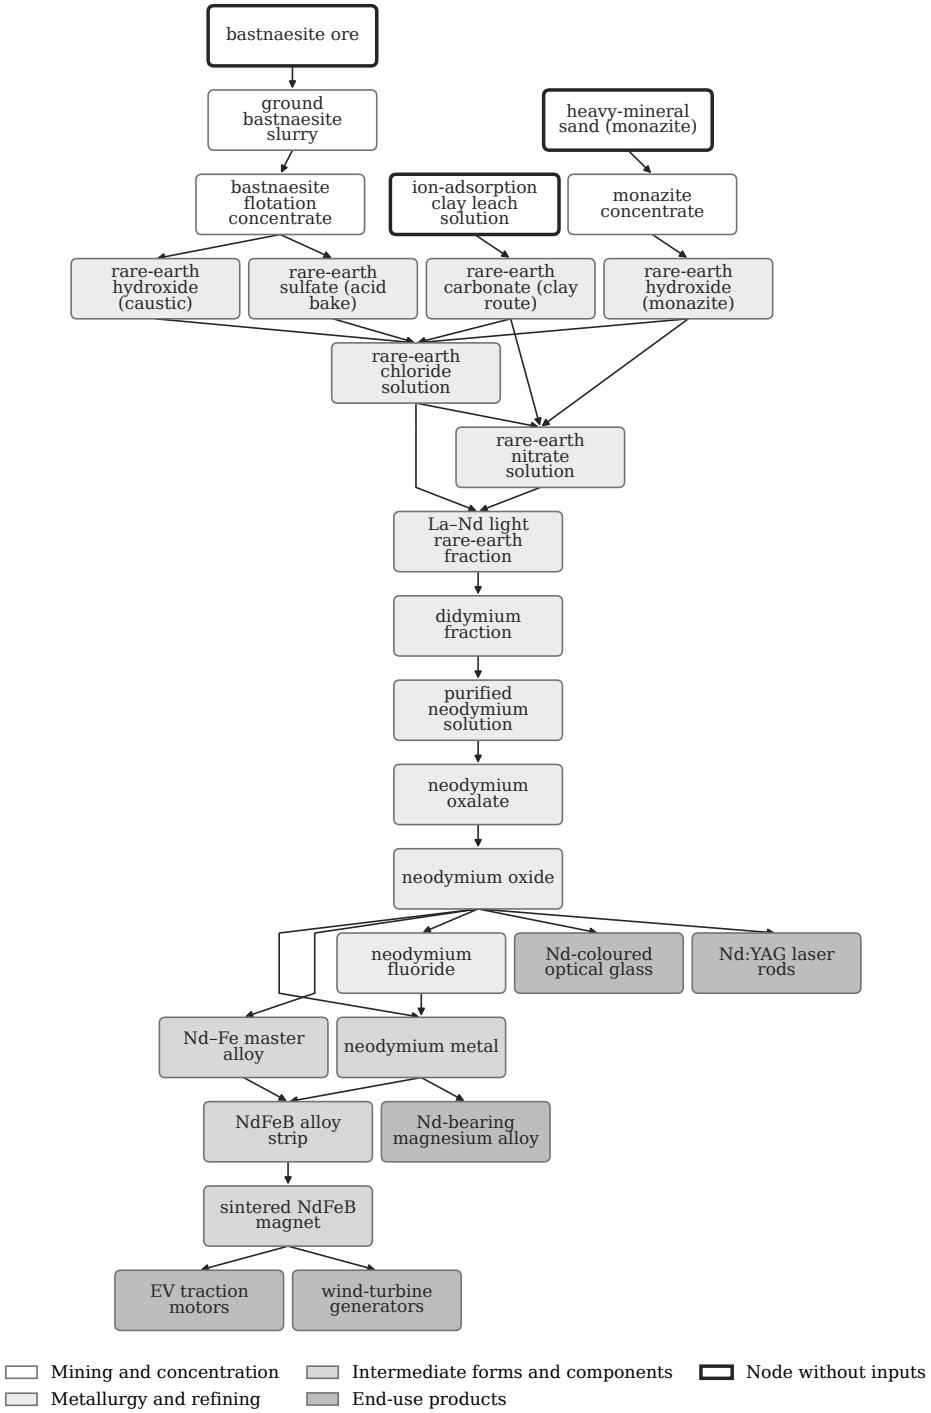
Of the 6,848 nodes, 6,306 contain the material. The others include goods in whose production the material is consumed and co-products created by the same process. The longest forward path contains 20 production steps. Figure C.2 makes the resulting structure concrete. It follows neodymium from ores and concentrates through separated compounds, metal, alloys and magnets into traction motors and wind-turbine generators. Branches into optical glass and laser rods illustrate how the same material reaches different end uses.

Figure C.1: From Source Evidence to Production Networks



Notes: The four stages run from top to bottom. Arrows show the sequence within each stage, boxes labeled Output identify its result, and dashed arrows denote revision or repeated search. Figure 1 follows one neodymium path through the four stages. By comparison, this figure locates the evidence and review decisions. IA1 details stages 1 and 2, IA2 stages 3 and 4, and IA3 the subsequent quantification and firm aggregation.

Figure C.2: Main Routes of the Neodymium Production Network



Notes: The network runs from top to bottom, with each arrow marking a forward production step. The figure displays grounded trunk paths that begin at one of the three nodes without mapped inputs and reach an end-use product. It contains 27 of the network's 99 nodes and 32 of its 113 forward edges. The omitted trunk nodes are by-products, residues, and branches that end before an end-use product. They include flotation tailings, magnet swarf, and the catalyst branch. In addition, the figure omits 50 supplementary goods. Shading marks the production stage, thick outlines mark nodes without mapped inputs, and node names are shortened.

Table C.1: Composition of the 55 Production Networks

	Nodes		Edges					
Material	Trunk	Supplement	All	Grounded	Judged	Doubtful	HS6 codes	Depth
Aluminum	99	237	489	383	102	4	423	11
Antimony	120	55	229	210	18	1	225	12
Baryte	76	56	190	185	5	0	195	13
Bauxite	117	67	255	240	12	3	252	13
Beryllium	74	20	125	121	4	0	105	15
Bismuth	104	38	181	179	1	1	144	14
Boron	97	91	270	248	22	0	233	10
Cerium	46	60	142	108	32	2	119	14
Chromium	61	85	244	173	65	6	167	11
Cobalt	56	104	218	192	25	1	188	14
Copper	127	145	388	314	71	3	343	15
Dysprosium	38	39	81	72	9	0	90	14
Erbium	41	28	73	62	11	0	60	13
Europium	57	25	96	88	8	0	55	17
Feldspar	44	20	86	85	1	0	89	13
Fluorspar	71	71	181	154	27	0	179	11
Gadolinium	53	55	119	102	17	0	102	17
Gallium	44	43	103	92	11	0	102	9
Germanium	31	36	72	62	10	0	87	13
Gold	80	52	184	166	18	0	119	14
Hafnium	65	29	118	108	10	0	113	12
Holmium	43	12	56	54	2	0	53	13
Indium	64	42	138	125	13	0	135	15
Iridium	78	38	139	129	9	1	96	14
Lanthanum	41	54	100	81	19	0	98	15
Lead	113	76	248	231	16	1	222	14
Lithium	56	59	153	142	11	0	120	12
Lutetium	48	11	61	60	1	0	39	13
Magnesite	52	49	155	140	12	3	74	9
Magnesium	46	73	168	144	24	0	141	13
Manganese	73	71	227	196	31	0	192	10
Neodymium	49	50	113	96	17	0	101	16
Nickel	125	118	400	274	120	6	333	18
Niobium	80	33	147	132	15	0	180	11
Osmium	43	12	66	63	3	0	39	12
Palladium	82	45	167	153	11	3	135	13
Phosphate rock	50	56	159	124	35	0	97	10
Platinum	81	52	167	145	22	0	127	12
Potash	57	53	146	130	13	3	98	10
Praseodymium	44	29	82	74	8	0	86	14
Rhodium	62	41	124	107	16	1	100	16
Ruthenium	56	39	105	102	3	0	82	13
Samarium	66	20	125	120	5	0	96	17
Scandium	23	28	63	52	11	0	45	10
Silicon	35	115	175	130	37	8	194	12
Silver	150	68	294	268	26	0	263	13
Tantalum	47	57	133	127	6	0	125	17
Terbium	59	28	98	93	5	0	70	17
Thulium	49	14	64	61	3	0	54	14
Tin	134	66	266	230	36	0	323	17
Tungsten	68	66	209	200	9	0	167	15
Vanadium	77	37	136	134	2	0	181	13
Ytterbium	25	24	50	46	4	0	51	13
Yttrium	64	57	133	125	8	0	158	14
Zinc	188	70	427	404	23	0	311	20
Total	3,829	3,019	9,138	8,036	1,055	47	1,903	

Notes: This table summarizes the 55 production networks. Trunk nodes are product groups identified in technical sources, and supplementary nodes are additional admitted goods. Grounded edges include trunk relations, within-product links and sourced connections. Judged and doubtful edges rely on progressively weaker production evidence, but all enter the headline measure. HS6 counts distinct customs codes within each network, with the total counting unique codes across networks. Edge counts and depth, the longest forward production chain, exclude the 316 separately recorded returns.

D Quantification and Firm Mapping

The main text describes how we construct physical requirements and map network goods to firms. This appendix supplies three additional pieces needed to evaluate the measure: the exact product-level valuation rules, an accounting of the product-line classifications, and the sensitivity of firms' exposure rankings to alternative valuations. Internet Appendix IA3 gives the full implementation. IA4 describes the industry-based extension.

D.1 Product-Level Valuation Rules

Benchmark normalization. We express physical quantities in units of the material's benchmark form, usually a refined metal or separated oxide. Its quoted price p_m^* corresponds to p_m in the model, and dependence at every benchmark node is normalized to one. If c_m^* is the median content across these nodes, the benchmark-form content of good j is $\hat{c}_j = c_j / c_m^*$. The ratio

$$\gamma_{mj} = \frac{\hat{c}_j \, p_m^*}{p_j} \tag{30}$$

is the value of the material held by the good relative to the good's price p_j, before processing losses and upstream markups. To put paths that bypass the benchmark on the same basis, we value the material at $v_m = p_m^* / (\mu_m^* c_m^*)$ per contained metric ton, where μ_m^* is the median markup across benchmark nodes. A non-benchmark root then enters at its material content valued at v_m, together with its own markup. Mining and processing upstream of the benchmark are already reflected in the benchmark price.

Propagation through the network. Let W_j be the requirement per metric ton of good j, expressed in benchmark-form units and adjusted for markups along the production path. Starting from $W = 1$ at a benchmark node, we compute

$$W_j = \mu_j \sum_i f_{ij} \, W_i, \tag{31}$$

where f_{ij} is the physical requirement of input i. For an input retained in the output, it includes processing losses. For an input consumed during production, it is the amount used net of any process return. A path beginning upstream of the benchmark enters the recursion through the material held by that input, valued on the benchmark basis. We apply the recursion to every production link in the network and refer to the resulting propagation as the *walk*. A dependence estimate obtained in this way is

a *walked estimate*. Equation (4) then gives

$$\beta_{mj} = \frac{p_m^* W_j}{p_j}. \tag{32}$$

Intermediate prices cancel, and each producing stage's markup enters once. The product value entering the firm mapping is β_{mj}/μ_j, which expresses cost dependence per dollar of sales.

Price consistency. Equation (32) uses one price p_j for each physical good. We begin with the trade unit value of the tariff line or customs code containing the good and use a recorded price of the good itself where available. Sometimes a code-level price falls below the value of the material contained in the good. We keep that price only when trade under the code carries roughly the same material content. Otherwise, we use the good's own recorded price or construct one from its material value. Where available, the construction also uses the estimated share of the good's value attributable to that material. Internet Appendix IA3 states the full price hierarchy and its consistency checks. As a result, 637 of the 4,122 physical goods receive a constructed price. Another 143 goods remain unpriced.

Bounds. A cost share cannot exceed one. When the walk implies a larger value, the product price, production recipe, and markups cannot be reconciled. We then replace the walked estimate with a bound. The bound equals one when the benchmark value of the material held by the good is at least the good's price, and γ_{mj} otherwise. A good priced by another material's benchmark takes the smaller of one and its content share. We also require the dependence estimates of a physical good to sum to at most one across materials. When the sum exceeds one, we allocate one across the materials according to value shares and classify all of the good's estimates as bounds. Goods without a usable price receive no estimate. Bounds remain in the measure and account for 40 percent of summed cost-side dependence in the material–code table.

Upstream and producer-side measures. A good upstream of the benchmark supplies material used to produce the benchmark form. It therefore receives no downstream cost-side dependence. Its relevant measure is the elasticity of its own selling price to the benchmark price, as in Section 2.2. Because processing charges can make this price a leveraged claim on the benchmark, the elasticity can exceed one.

We also record the producer share P_{mj}. It equals $\min(\gamma_{mj}, 1)$ for the benchmark and near-pure forms of the material, uncapped γ_{mj} for forms upstream of the benchmark, and zero otherwise. This share measures how much of a commodity-form good's price represents the material being sold. Both producer-side quantities are reported separately from cost dependence and are never netted against it.

Table D.1 shows which rule determines each material–product pair. The pricing

scope contains 6,332 pairs across 55 materials. Of the 5,573 pairs at or downstream of the benchmark, 5,381 receive cost-side dependence and 192 remain without an estimate. The remaining 759 pairs carry the producer-side measures just described. The physical-input checks in Section 3.2 led to 234 recorded author decisions. Internet Appendix IA3 reports their objects and outcomes.

Table D.1: Valuation Rules by Type of Material–Product Pair

Type of pair	β_{mj}	P_{mj}	Pairs
Panel A. At or downstream of the benchmark form			
Benchmark node	1	$\min(\gamma_{mj}, 1)$	163
Walked estimate, walk at most the price	$p_m^* W_j / p_j$	by form	4,495
Bound, walk above the price	1 if $\gamma_{mj} \geq 1$, else γ_{mj}	by form	376
Bound, one good's estimates sum above one across its materials	its share of one, by value	by form	339
Bound, priced by another material's benchmark	$\min(1, \gamma_{mj})$	$\min(1, \gamma_{mj})$ or 0	8
No estimate, no usable price	none	0	184
No estimate, content zero and use above the price	none	0	8
Panel B. Upstream of the benchmark form			
Revenue-side elasticity, $D_{mj} = (p_m^* / p_j) \sum_i f_{ij} \hat{c}_i$	$\mu_j D_{mj}$	γ_{mj}, uncapped	755
Priced by another material's benchmark	$\min(1, \gamma_{mj})$	$\min(1, \gamma_{mj})$ or 0	4
All pairs			6,332

Notes: This table assigns the 6,332 material–product pairs to their final valuation rule. β_{mj} is cost-side dependence, P_{mj} is the producer share, and γ_{mj} is the benchmark value of the material held by a good relative to the good's price. By form equals $\min(\gamma_{mj}, 1)$ for benchmark and near-pure forms and zero otherwise. For upstream pairs, $\beta_{mj} = \mu_j D_{mj}$ is stored but excluded from firm dependence. The bound rows sum to 723, the no-estimate rows to 192, and Panel B to the 759 upstream pairs.

D.2 Product-Line Classifications and Coverage

Section 3.3 describes how reported product lines are matched to network goods. Table D.2 accounts for the observations behind that procedure. Panel A shows how the reconciliation of product-line revenue determines the firm-period sample. The remaining panels work with normalized line names. Panel B traces the review that governs each line. Panel C then records whether the line is measured, assigned zero, excluded as resale, or left outside the scope of the measure. A firm-period remains unavailable when no retained product can be valued and the combined excluded and unresolved revenue exceeds revenue on service, software, and financial lines assigned zero.

The revisit of flagged lines in Section 3.3 works as follows. When the only issue is an unidentified product on a residual or generic line accounting for no more than 25 percent of firm revenue in every accepted firm-period, we presume that product lies

outside the catalog. The other flagged lines receive a web-search reading. Sourced answers at high or medium confidence settle the line, while the remaining cases enter a model-based expert review, a reviewer-role model pass, in descending order of revenue or economic stake until 95 percent of each review block is covered. The remaining lines keep their web-search answers. A separate omission check re-asks a line when the firm's text names a material absent from the selected goods but present in other catalog goods under the line's customs codes. The final identification gives priority to this model review, followed by the omission re-ask, the web-search reading, the residual-line presumption, and the original company-context reading.

The original company-context reading remains the governing decision for 42,246 of the 53,727 read lines with positive revenue. Thus, most lines require no override. As a reproducibility check, we re-ask 1,110 unchanged lines. The repeated reading assigns the same role in 97 percent of cases and identifies the same goods in 86 percent.

Table D.2: Classification of Product Lines and Revenue Coverage

Panel A. Firm-periods	
Firm-periods with product-line revenue	94,526
Main sample: lines reconcile to the company total	86,640
Other firm-periods: lines short of the total, no total, or excluded	7,886
Main sample: exposure measured	81,479
Positive on at least one material	30,791
Measured at zero on every material	50,688
Main sample: exposure unavailable	5,161
Mean share of revenue identified, measured firm-periods	0.948
Panel B. Product lines	
Firms with business-segment revenue	34,066
Firms reporting product lines	14,404
Distinct product-line names	85,228
Initial sort: specific products	10,413
product families	28,402
non-goods (assigned zero)	25,777
unresolved	20,636
Lines to be read at firm level	57,039
With revenue, governed by the original company-context reading	42,246
Governed by the residual presumption	6,064
Governed by the omitted-material check	1,466
Governed by a web-search answer	3,221
Governed by the model-based expert review	730
Panel C. Final classification of the lines	
Measured in the main column	44,622
of which merchant lines	393
At zero (non-goods and services)	26,658
Resale lines of made goods, outside the measure	6,896
Geographic lines in which no product was identified	2,262
No revenue in an accepted firm-period	4,790

Notes: A line is a firm and normalized product-segment name pooled across years. Panel A applies the reconciliation and availability rules of Section 3.3. Resolved coverage is revenue from measured and zero-classified lines divided by company revenue. Panel B reports the decision that governs each line, and Panel C its final treatment. Merchant lines trade the material itself or its ore, metal, alloy, semi-fabricated, or compound forms.

D.3 Sensitivity of Firm Exposure Rankings

The key question is whether the valuation rules determine firms' rankings as well as their dependence levels. Table D.3 examines this question. Each comparison revalues the same 4,658 material–product pairs with a walked estimate or benchmark form. Production links, product prices, aggregation weights, and the values assigned to bounds remain fixed.

Alternative valuations. We first set every markup to one and repeat the requirement calculation at the same product prices. A second variant lets retained inputs pass fully into output. Inputs consumed during production keep their baseline values. Finally, the content-value measure uses the benchmark value of the material contained in the finished good as a share of its price. Consequently, it omits material consumed along the production chain, processing losses, and upstream markup adjustments.

Firm rankings change little under these alternatives. Panel A reports rank correlations above 0.95. Depending on the variant, 11 to 24 percent of observations move to another exposure decile. Yet fewer than 4 percent move by more than one decile. The corresponding product-level correlations are at least 0.996 for the first two variants and 0.957 for content value. Content value is zero for a good that uses the material without retaining it. Thus, the valuation choices affect exposure levels more than the ordering of firms.

Panel B asks whether product coverage alone is sufficient. We assign one to every covered good using the material. The resulting measure captures the share of firm activity involving that material, irrespective of the amount required per dollar of product sales. The rank correlation falls to 0.32, and only 16 percent of observations remain in the same exposure decile. Thus, quantified input requirements contribute substantially to the rankings even with product coverage held fixed.

Table D.3: Firm Exposure Rankings under Alternative Valuations

	Rank correlation	Same decile	Within one decile
Panel A. Valuation assumptions			
No markup adjustments	0.999	89.3	99.8
No processing losses	0.988	81.3	97.3
Material content value	0.958	75.6	96.4
Panel B. Dependence intensity			
Material presence only	0.317	16.1	39.4

Notes: Panel A revalues the same 4,658 material–product pairs with a walked estimate or benchmark form. Panel B assigns one to each good with positive dependence. Prices, production links, and mapping weights remain fixed. Rank correlations are Spearman correlations over observations positive under either measure. Deciles are formed within material and year on the common positive sample, retaining cells with at least ten firms (204,398 observations).

E Validation-Benchmark Details

Section 4 reports the comparisons with specialist research and corporate filings. This appendix supplies the construction and matching details needed to interpret those results. Following Gentzkow et al. (2019), we code each claim from its surrounding passage and retain the supporting quotation. The reported comparisons use the classifications after the full review procedure described below.

E.1 Specialist-Report Benchmark

Sources and benchmark units. The registry contains 136 source entries and 124 distinct texts after exact duplicates are combined. Revised editions and separately published summaries remain separate. We exclude five documents available during network construction, leaving 119 texts for the independent comparison.[55] Table E.1 reports the retained sources by publisher, format, and year.

The retained corpus contains 8,404 passages, 2,462 distinct material relationships, and 1,564 product or industry links. The primary benchmark keeps affirmative product expressions, normalizes their descriptions, and combines repetitions within material. This yields 535 material–product units. One reviewed label is excluded because it describes a project or scenario.

Review and matching. Every retained observation includes a quotation located in the cited passage. A second large language model (LLM) reviews a blinded sample of 450 records, divided equally among retained observations and two types of rejected proposal. The reviewer sees neither the original disposition nor the exposure measure. Human review resolves 13 uncertain extractions and 30 product or industry assignments, and it makes 12 document-level corrections. The product matching adds 22 rulings on uncertain cases, 17 of which enter the retained comparison.

Products enter the benchmark when a report affirms an exposure or effect. We combine repeated expressions within a material but preserve distinct product forms. Each description receives a normalized concept and, where possible, a 2022 Harmonized System (HS) or 2017 North American Industry Classification System (NAICS) assignment. Broad descriptions retain their broader category, and ambiguous descriptions retain their candidate set.

We compare each report product with all network products for the same material, including products without customs codes. Exact matches require comparable physical scope. A category match allows a broader or narrower description of the same

[55] The excluded texts are the IEA's *The Role of Critical Minerals in Clean Energy Transitions* (2021), *Critical Minerals Market Review 2023*, and *Global Critical Minerals Outlook* (2024 and 2025), and IRENA's *Geopolitics of the Energy Transition: Critical Materials* (2023).

product concept. Same-family products share a primary function and production stage. An adjacent-stage product is one production, recycling, or assembly step away. General sectors, facilities, and more distant relationships fall outside these categories. Table E.2 reports the resulting matches.

Code-based comparisons. Of the 1,386 participant links attached to a material, 1,107 can be assigned an HS or NAICS code or candidate set. Repeated assignments count once within material. Panel A of Table E.3 asks whether these coded report relationships appear among network products, priced products, or industries. Panel B reverses the comparison after restricting the measure to materials discussed in the reports. A broad assignment is compared at its recorded level. A candidate set matches when one of its eligible categories appears in the measure.

Table E.1: Sources of the Specialist-Report Benchmark

Source category	Documents	Share (%)
A. Publisher type		
Financial research	95	79.8
Policy and public research	24	20.2
B. Document format		
Research report	35	29.4
Research article	54	45.4
Article or transcript	10	8.4
Commodity study	8	6.7
Policy report	8	6.7
Flagship report	4	3.4
C. Publication year		
2016	1	0.8
2018	1	0.8
2020	2	1.7
2021	7	5.9
2022	13	10.9
2023	19	16.0
2024	14	11.8
2025	22	18.5
2026	38	31.9
Undated	2	1.7
Total	119	100.0

Notes: This table summarizes the 119 specialist documents in the benchmark by source, format, and publication year. Exact duplicates are combined, and five sources used in network construction are excluded. Financial research includes banks and their asset- and wealth-management divisions. Policy and public research includes DERA/BGR, IEA, and IRENA. Format and year follow the source registry. Shares may not sum to 100 because of rounding.

Table E.2: Overlap between Report Products and the Production Network

Relationship to production-network products	Units	Share
Exact product	339	63.4%
Clear product-category correspondence	104	19.4%
Same-family product	14	2.6%
Adjacent production stage	38	7.1%
No counterpart	40	7.5%
Direct or category overlap	443/535	82.8%
Including same-family products and adjacent stages	495/535	92.5%

Notes: This table compares distinct material–product expressions in specialist reports with network products for the same material. Direct or category overlap combines the first two rows. Extended overlap adds same-family and adjacent-stage products. Products without an HS code remain in the comparison.

Table E.3: Supplementary Code-Based Comparison with Specialist Research

Panel A. Coded report relationships found in each measure

	Network products		Priced products		Industries	
Resolution	Match	Share	Match	Share	Match	Share
Specific six-digit assignment	69/80	86.3%	62/80	77.5%	7/7	100.0%
Broader category	88/100	88.0%	81/100	81.0%	43/51	84.3%
Multiple candidates	12/15	80.0%	12/15	80.0%	3/3	100.0%

Panel B. Coded measure categories named in the reports

	Six-digit codes		Four-digit headings	
Measure layer	Match	Share	Match	Share
Network products	69/6,365	1.1%	141/3,236	4.4%
Priced products	62/3,348	1.9%	131/2,548	5.1%
Industries	7/1,642	0.4%	52/900	5.8%

Notes: This table compares coded categories in specialist reports with the network, priced-product, and industry layers. The unit in Panel A is a coded relationship in the reports. Broad assignments are evaluated at their recorded level. A candidate set matches when any eligible category appears in the measure. Panel B reverses the comparison and limits the measure to materials covered by the reports, omitting candidate sets. Six- and four-digit matches use separate denominators.

E.2 Corporate-Filing Benchmark

Filing frame and extraction. The reading covers 9,364 Forms 10-K, 20-F, and 40-F matched to an accepted firm-year. Section 4.2 restricts the comparison to exact matches with the baseline sample. A fixed dictionary searches the filings for all 55 materials. Following the contextual treatment of financial language in Loughran and McDonald (2016), gold, silver, tin, lead, and silicon require constrained phrases such as "gold mining," "lead acid," or "silicon wafer." Bare mentions of these five words are omitted.

The search yields 7,411 filing–material cases. An LLM proposes a relationship and timing classification for each case. This initial reading produces 4,593 proposed affirmations, 2,113 mentions, 608 unresolved cases, and 97 cases from rejected extraction batches. We retain complete batches that pass the schema and evidence checks, and every proposed affirmation then enters the review below.

Operational benchmark. Every proposed affirmation is reviewed for timing, operational scope, and evidence that the passage establishes the individual material. Current production, use, storage, and transport qualify. A financial activity qualifies only when the passage also describes an operational relationship. Plans, forecasts, historical references, and mixed timing receive zero. Current exploration and research qualify.

Two classification rules address passages that name a material without establishing such a relationship. Conflict-minerals wording that lists tin, tantalum, tungsten, and gold as a group establishes an individual material only when another passage

does so. Of 88 such cases, 70 remain unresolved. We also review 170 cases that name a material as possible or alternative content. Forty remain unresolved because no passage establishes a current operational relationship with that material. These rules classify 110 proposed affirmations as unresolved. The full review leaves 4,154 current operational relationships, with each material counted once per filing.

A scope reading asks whether the cited passage attests a relationship tracked by the measure. A *sound* relationship says that the company uses the material in what it makes or in its own production, produces or processes it, or trades the material itself. Weaker ties include installed infrastructure, finished-good resale, equipment use, deposits, financial interests, and risk or market discussions. Two independent scope readings agree on 3,993 of the 4,154 relationships. A third resolves 89 disagreements, and human review resolves 72. In the baseline sample, 2,815 of 3,293 current operational relationships are sound and form the benchmark. The remaining 478 enter as no disclosed relationship: 470 have a weaker tie and eight are too fragmentary to decide. Under a broader benchmark that includes all 3,293 relationships passing the operational review, positive exposure recovers 48.0 percent and network presence 68.8 percent.

Unmatched disclosures and sample. Of the 706 disclosures captured by neither comparison criterion, 653 occur when the firm's measured product lines name no good in the material's network. The other 53 occur at firm-years without a measured product line. Network presence alone captures another 596 disclosures, including 580 through a product code and 16 through a named but unvalued good.

We read 48 deliberately selected cells, including 32 misses and 16 matched controls, and subjected ten to a follow-up review. Among the misses, 11 reflect product-information or assignment gaps, 8 activity-scope differences, and 13 product-map coverage or screening differences. Eight of the last group arise from the conservative screening rule. The other five suggest finer-grained coverage, packaging, or connection issues, but none was adjudicated as a network omission. Among the matched controls, nine have direct product support and seven rely on broad code matches that do not establish the specific material-bearing product. The targeted sample identifies these mechanisms but does not estimate their frequency or economic importance.

For each firm and fiscal period, we retain the earliest filing. The final comparison uses 7,666 exact filing–firm-year matches for 1,503 firms over fiscal years 2017–2025. It comprises 7,153 Forms 10-K, 506 Forms 20-F, and 7 Forms 40-F. Of these matches, 7,178 belong to the main sample, 252 to the secondary sample, and 236 to the conditional sample. Table E.4 reports the material-level counts.

Table E.4: Material Relationships in Corporate Filings and Product-Line Dependence

Material	Both positive	of which exposure	Measure only	Filing only	Agreement	Positive agreement
All 55 materials	2,109	1,513	70,179	706	83.2%	74.9%
Aluminum	571	368	1,964	208	71.6%	73.3%
Copper	357	293	1,770	139	75.0%	72.0%
Nickel	172	94	1,737	34	76.9%	83.5%
Zinc	131	91	1,757	46	76.5%	74.0%
Lithium	89	78	1,453	71	80.1%	55.6%
Silicon	102	101	1,645	12	78.4%	89.5%
Platinum	52	36	1,538	27	79.6%	65.8%
Cobalt	60	53	1,610	16	78.8%	78.9%
Magnesium	47	19	1,527	28	79.7%	62.7%
Palladium	44	18	1,007	20	86.6%	68.8%
Gallium	48	38	1,648	1	78.5%	98.0%
Tantalum	38	8	1,198	9	84.3%	80.9%
Manganese	38	37	1,354	8	82.2%	82.6%
Gold	44	44	1,837	0	76.0%	100.0%
Tungsten	32	19	1,414	12	81.4%	72.7%
Potash	25	20	1,156	0	84.9%	100.0%
Silver	24	9	2,232	1	70.9%	96.0%
Boron	23	22	1,974	0	74.2%	100.0%
Chromium	22	16	1,277	0	83.3%	100.0%
Vanadium	20	11	1,133	1	85.2%	95.2%
Iridium	13	4	1,176	8	84.6%	61.9%
Lead	17	17	1,353	2	82.3%	89.5%
Indium	8	8	1,157	11	84.8%	42.1%
Rhodium	11	11	623	7	91.8%	61.1%
Ruthenium	2	1	1,232	14	83.7%	12.5%
Germanium	15	12	1,356	0	82.3%	100.0%
Baryte	10	10	1,539	4	79.9%	71.4%
Fluorspar	12	12	1,290	0	83.2%	100.0%
Antimony	3	3	2,054	8	73.1%	27.3%
Yttrium	8	7	1,451	2	81.0%	80.0%
Other 25 materials	71	53	25,717	17	86.6%	80.7%

Notes: This table compares sound filing relationships with network presence for 7,666 exact filing–firm-year matches and 55 materials over fiscal 2017–2025. The sample contains 421,564 assessed cells after 66 undecidable classifications are excluded. Network presence means that a measured product line maps to a good in the material's network. The "of which exposure" column counts cells with positive cost dependence or producer-side exposure. Agreement is the share on which the filing and network indicators coincide. Positive agreement is the share of filing relationships with network presence. Each material counts once per filing. Rows require at least ten filing relationships, with the remaining 25 materials grouped as Other.