

A Constraint Transmission Framework for Mapping Physical Supply Scarcity to Public Equity Exposure

M.Malukas

Abstract

A mineral can be essential to a growing industrial use and still leave almost no measurable trace in the listed companies associated with it. This paper sets out the Stockatomics Constraint Transmission Framework, a mineral-first specification for connecting downstream demand to the difficulty of expanding or replacing supply, to concentration at a stated production stage, and to the production exposure of listed issuers. The mineral is the stable unit. Sectors and products explain demand. Companies enter only through a compatible measure of current output.

The specification separates four objects that thematic research usually collapses: structural supply responsiveness, a scenario-dependent demand outlook, concentration, and company participation in supply. Supply responsiveness is summarised by the Supply Elasticity Score (SES), an equal-weighted composite of lead time, capital intensity, processing bottlenecks, substitution difficulty and recycling difficulty. Despite the name, SES is not an econometric elasticity. Demand Momentum (DM) is a growth rate under stated sector shares. The Demand Pressure Score (DPS) is the normalised demand term used inside company scores. Mineral Exposure (MES), Supply Leverage (SLS) and the Critical Mineral Catalyst Score (CMCS) aggregate production shares against those mineral conditions. The aggregation identities are stated in full. Component rubrics and the fixed bounds that map raw sums onto 0–100 remain versioned implementation parameters.

A constructed numerical case shows why two minerals with the same SES can imply different adjustment paths, and why a demand revision moves the indicators of a verified producer without creating production exposure for a developer or a cost exposure for a buyer. The paper specifies data requirements, known weaknesses in the current database, and a programme of validity, sensitivity and outcome tests. It does not estimate a production network, does not establish that demand exceeds supply, and does not claim that the scores predict equity returns.

Keywords. critical minerals; supply constraints; physical economy; production exposure; public equities; composite indicators; demand scenarios

1. Introduction

Industrial demand reaches a public equity market only after it has passed through a physical production system. A change in the deployment of batteries, transmission equipment or advanced electronics changes the quantity of particular materials that must be mined, converted, qualified and delivered. Additional supply, if it comes, depends on development time, capital, processing capacity, the practicality of substitutes, and the availability of scrap. Those responses do not share a clock. A smelter bottleneck can bind within a year. A new mine, in many jurisdictions, cannot.

The research problem is therefore not solved by naming a growing end market, or by naming a mineral that market uses. An analyst has to establish where adjustment is difficult, at which stage, and which listed companies have a measurable position in the relevant flow. Company sector labels and thematic index membership do not answer that question. They record an association. Association is not output.

Stockatomics is a financial data platform for the physical economy. It connects critical minerals, products, sectors and public companies, and it carries prices, trade flows, supply descriptions and demand context alongside that map. The framework in this paper is the scoring core of that map. Market prices and trade statistics can inform a research note. They are not inputs to the core score calculations documented here.

The organising principle is mineral-first. Minerals are the stable analytical unit, because a copper cathode, a separated rare-earth oxide or a battery-grade lithium chemical is a thing with units, a production stage and a denominator. Sectors and products explain why demand for that thing might change. Companies provide the public-market connection, and only when a production figure can be placed on a compatible basis with world output. Starting from the material makes the chain inspectable. An analyst can ask why demand is assumed to change, how supply could respond, where production or refining sits, and which issuers actually contribute tonnes.

Constraint transmission is the name used here for that sequence of links. Scarcity conditions may show up as availability, price, volume or an investment incentive. Whether a listed company captures an economic benefit is a further question, turning on costs, contracts, asset ownership, financing and the price already in the equity. The core scores stop at physical positioning. They are built to be argued with, not to stand in for a valuation.

Three things follow for the design. The mineral-level objects — responsiveness, demand and concentration — have to remain readable on their own, because they do not move together. The company connection has to be a production share, not a product mention. And a scenario has to say which input moved, so that a later reader can tell a demand revision from a data correction from a change in the method. The paper reports no backtest, no predictive-accuracy estimate and no realised investment result. Those are later projects, with their own designs, and they are not implied by the existence of a score.

2. What the scores are for

The scores answer a positioning question. Given a stated demand scenario and a stated production basis, which listed companies participate in mineral systems whose supply is difficult to expand or replace, and where that participation coincides with concentration? They do not answer whether the mineral is scarce this quarter. A present shortage is a statement about quantities, inventories, utilisation, trade access and timing. Long development lead times can matter in a temporary surplus. A short disruption can produce scarcity in a market whose long-run supply response is relatively flexible. Every reading needs a horizon, and the horizon has to be the same one used for the demand assumption.

They also do not answer whether the equity should reprice. A constrained mineral can raise costs for one company and support revenue for another. A producer with a large physical share can still face fixed-price contracts, operational disruption, dilution, or a project whose capital cost consumes the rent. No monotonic relationship between the core scores and subsequent equity returns is assumed, and none should be inferred from the construction.

Government critical-mineral lists are a different object again. A designation can reflect strategic use, import reliance, allied concentration, environmental damage, or a political process. Graedel et al. (2015) already separate supply risk, vulnerability to supply restriction and environmental implications, precisely because a single label hides the mechanism. A high Stockatomics score is not a substitute for that designation, and a designation is not evidence that a listed issuer has production exposure.

3. Relation to existing work

Production-network research supplies the reason to care about transmission, not a template for the scores. Acemoglu, Carvalho, Ozdaglar and Tahbaz-Salehi (2012) show that the asymmetry of input–output linkages can keep idiosyncratic disturbances from averaging out in the aggregate. Carvalho, Nirei, Saito and Tahbaz-Salehi (2021) document propagation after the Great East Japan Earthquake through actual supplier and customer ties. Both results motivate attention to where a system is connected and where it is concentrated. Neither result validates this framework. Stockatomics does not estimate a production network, does not recover input–output coefficients, and does not model cascade dynamics. The debt to that literature is architectural: dependencies should be written down, and a disturbance at one point should not be treated as if it arrived everywhere at once.

Mineral criticality research is closer in subject and easier to confuse with the present exercise. Graedel et al. (2015) place 62 metals in a three-axis criticality space. Supply risk in that study is driven in part by geopolitical concentration of primary production, weak substitutes and political instability. The metals of most concern are often by-products, used in small quantities for specialised applications, with no effective substitute. That description is useful here as a warning. A mineral-level score compresses asset-level and host-commodity economics, and a by-product flag is not a model of the host metal. The narrower question in this paper is how difficulty in expanding or replacing mineral supply interacts with a downstream demand scenario and with listed-company production. Environmental implications are outside the core scores.

Official commodity sources fix the units in which an exposure claim is allowed to be made. The International Energy Agency’s *Global Critical Minerals Outlook 2026* describes mineral systems in terms of mining, refining, demand and diversification (IEA, 2026). The U.S. Geological Survey’s *Mineral Commodity Summaries 2026* provides commodity-specific production, reserve and market information, and is explicit that world mine production, reserves and resources are not interchangeable (USGS, 2026). Those publications do not feed the scores automatically. They illustrate the discipline the scores require: the material form, the stage and the year have to be named before a share is computed.

On the construction of composites, the OECD and Joint Research Centre handbook remains the relevant standard (OECD/European Union/EC-JRC, 2008). A composite is an argument about what belongs in the index, how inputs are normalised, and how sensitive the ranking is to those choices. Numerical precision does not retire uncertainty in the evidence. Equal weights, used below for SES, are a declared choice made in the absence of a validated weighting model, not a finding that each mechanism contributes a fifth of economic constraint.

4. Architecture

4.1 Unit of analysis

Each mineral carries a supply-constraint profile, a demand profile and a concentration assessment. A company–mineral relationship joins that profile to a listed issuer. Products and sectors sit on the demand side. Catalysts record developments that may later change a demand or supply assumption. The separation is the point. An observed production figure, a scenario growth rate and a speculative technology announcement do not have the same evidential rank, and the architecture refuses to give them one.

The result is a structured dependency map, not a general-equilibrium model. It does not solve for a market-clearing price, an inventory identity or the speed of adjustment. A mineral–company link does not establish that the mineral is financially material to the issuer. Materiality is a revenue, margin and contract question, asked after the physical relationship has been characterised.




Figure 1. Conceptual structure. Demand, supply responsiveness and concentration remain distinct mineral-level assessments. Company production exposure joins those assessments to the company indicators. Arrows are analytical dependencies, not estimated causal effects. Qualitative catalysts inform research and, when the evidence is sufficient, a dated revision of an input. They do not enter the core calculation.

4.2 Layers that are not allowed to collapse

Four layers do the work. The first describes downstream activity and its claimed relationship to mineral demand. The second describes the capacity of the mineral system to adjust, including substitution and recycling, because both can relieve primary supply. The third describes concentration at a named stage. The fourth connects the mineral to companies through measured participation in production. Company indicators are summaries of selected combinations. They are not a fifth independent fact about the world.

The layers have to stay interpretable when they disagree. A mineral can be difficult to expand while its demand outlook is weak. A concentrated market can have spare capacity. A company can supply an important mineral without that activity dominating earnings. Folding these cases into a single scarcity narrative hides the reason for a score and the condition under which it would change. The rest of the paper is a specification of how the layers are scored, how they are combined, and what is deliberately left out of the combination.

5. Supply responsiveness

The Supply Elasticity Score, SES, expresses how difficult it is for the mineral system to respond to a change in demand. Higher values mean greater constraint and lower responsiveness. The scale is 0–100. The name is retained because it is the label used on the platform. It should not be read as an elasticity. SES is not an estimate of the percentage change in quantity supplied caused by a percentage change in price, and it is not identified from a supply curve.

Five dimensions define the profile: time to add meaningful supply, capital intensity, processing bottlenecks, substitution difficulty and recycling difficulty. Scoring direction is fixed. One hundred means the constraint is extreme; zero means the adjustment route is comparatively open. In the present specification the dimensions are equally weighted:

$$SES_m = 0.20 T_m + 0.20 K_m + 0.20 P_m + 0.20 S_m + 0.20 R_m$$

T time to add supply; K capital intensity; P processing bottleneck; S substitution difficulty; R recycling difficulty. Each component is on 0–100. Equal weights are locked in this version.

Equal weights are a methodological commitment, not an empirical result. There is not yet a body of outcome evidence that would justify putting more weight on lead time than on processing, or on substitution than on recycling. Asymmetric weights would look more precise than the supporting information. They can be introduced in a later version if sensitivity analysis and outcome studies show that a particular dimension dominates realised supply response. Until then, the component profile is part of the output, not a private input to a weighted sum.

Capital intensity is written K here, not C . Concentration, introduced below, is written Γ . The two objects are different, and a shared symbol would make the company formulas unreadable. Nothing in the economics depends on the letter.

Dimension	Mechanism	Reading of a high score
Time to add supply, T	Development, permitting and commissioning delay a production response.	Meaningful additional output takes longer to deliver.
Capital intensity, K	Expansion depends on funding and on the scale of the assets required.	New supply requires a more demanding investment.
Processing bottleneck, P	Conversion or refining can limit usable supply even when mined material exists.	Processing capacity is difficult to replicate or to access.
Substitution difficulty, S	Alternatives may change cost, performance or qualification.	Major uses have fewer practical replacements.
Recycling difficulty, R	Recovery depends on collection, process technology and available scrap.	Secondary supply is a weak adjustment route.

Table 1. Dimensions of the supply-constraint profile. Component scores are judgement-based assessments against a versioned rubric, applied to named evidence. The rubric thresholds are implementation parameters and are not reproduced in this paper.

The profile is informative when the summary is not. Two minerals can share an SES because one is constrained by development lead times and the other by processing, or because one cannot be substituted and the other cannot be recycled. Their response to a price spike, a policy subsidy or a technical change will differ. Publishing only the average would erase the distinction the score is supposed to preserve. Section 10 constructs a pair with identical SES values and opposite profiles for that reason.

Component scores are not measurements in the sense that a production tonne is a measurement. They are structured judgements: a reviewer applies the rubric to a cited body of evidence and records the score. Disagreement between reviewers is evidence about the rubric, not a nuisance to be averaged away in silence. The evaluation agenda in Section 12 treats inter-reviewer agreement as a primary test.

Processing bottlenecks and concentration can describe overlapping features of the same chain. A rare-earth separation circuit can be both hard to replicate and geographically concentrated. Keeping the concepts in separate terms makes the overlap visible. It can also amplify a shared underlying condition when SES and Γ later meet in a company score. Sensitivity analysis has to remove one of the overlapping terms before a small difference between two company scores is treated as independent information.

6. Demand momentum and demand pressure

6.1 Two representations of one assessment

Demand Momentum, DM, is the assumed growth outlook for mineral demand, expressed as a percentage rate. The Demand Pressure Score, DPS, is the normalised form of that demand assessment used inside company scoring, on a 0–100 scale. In the current implementation the two are not independent signals. They are two representations of one sector-share-weighted growth proxy. Agreement between DM and DPS is then a property of the transform, not corroboration.

The operational demand identity assigns each mineral's demand across a defined set of sectors and combines the shares with sector growth assumptions:

$$DM_m = \sum_j \omega_{m,j} g_j$$

$\omega_{m,j}$ is mineral m 's demand share in sector j . g_j is the sector growth assumption, in percent. Shares are required to exhaust the allocation and not to overlap: $\sum_j \omega_{m,j} = 1$.

The documented sector set is artificial intelligence, semiconductors, electricity grids, electric vehicles, batteries, defence, robotics, and a residual category. Several of those labels describe connected activities. A battery installed in an electric vehicle cannot be allowed to generate two claims on the same baseline material. The allocation rule has to say, for each mineral, which sector owns the tonne. Without that rule the shares will not exhaust the allocation in any meaningful sense; they will double-count it.

Weighting by share is what gives scale a role. Rapid growth in a small use need not dominate slower growth in a large use. That is the right instinct, and it is not the same thing as a quantity model. Translating sector activity into mineral demand assumes a stable relationship between activity and material use. Design changes, chemistry, utilisation and recycling rates can break the relationship without any change in the sector growth rate. DM is a scenario-dependent growth proxy under an intensity assumption, not an observed consumption growth rate.

A further arithmetic point is easy to miss. A share-weighted average of sector compound annual growth rates is not, in general, the compound growth rate of aggregate demand. If the sectors have different baselines, the aggregate growth rate depends on those baselines, not only on the shares used as weights. Where baseline quantities exist, the preferred reading of DM is the growth rate of the summed sector quantities. The share-weighted average is the operational proxy when the implementation has shares and growth assumptions but not a full quantity account.

The word momentum is also looser than the measurement. A positive growth assumption does not demonstrate acceleration. Acceleration requires a comparison with an earlier rate over a defined window. The platform label is retained. Users should read DM as a stated growth outlook.

6.2 The preferred pressure identity

The broader demand model is a quantity identity. Incremental pressure on a mineral is the sum, over products, of product growth, mineral intensity and product scale:

$$DPS_m^{raw} = \sum_i (g_i \times \iota_{i,m} \times Q_i)$$

g_i is the growth of product i , $\iota_{i,m}$ is the mineral intensity of that product, and Q_i is product scale in compatible units. A high-growth product with negligible scale then cannot dominate the mineral.

This is the preferred construction, because it keeps scale, intensity and growth as separate evidence. The current implementation does not yet use it as an independent additive term. Product links remain qualitative. They do not add a second demand quantity on top of the sector proxy. A product development can support a revision of a sector assumption inside a named scenario. It should replace the relevant assumption, or be reconciled inside a single demand account. Adding a product estimate to a sector total that already contains the product double-counts the tonne.

A future product-based implementation needs compatible units, intensities, baselines and an explicit reconciliation with sector totals. Until those exist, DPS in company scoring is the 0–100 normalisation of the sector-weighted growth proxy, not the raw product sum. A negative demand scenario should remain visible in economic units even if a bounded pressure score sits on its floor. DPS is not a signed measure of every upside and downside demand effect.

6.3 Company demand momentum

At company level, DM is the production-weighted average of mineral demand outlooks, taken over the minerals for which the company has an eligible exposure weight:

$$DM_c = \sum_m (W_{c,m} \times DM_m) / \sum_m W_{c,m}$$

$W_{c,m}$ is the exposure weight. In the core specification that weight is the production share $E_{c,m}$ defined in Section 8. Minerals without a compatible production share do not enter the denominator.

Company DM therefore describes the covered mineral mix, not the company's revenue mix, and not the minerals it mentions in a filing. Incomplete coverage has to be stated beside the figure. A company whose only measured output is a slow-growing bulk mineral will show a low DM even if a development project, unscored, sits on a fast-growing speciality. That is the correct behaviour of the formula. It is also a reason not to read the figure as a complete demand thesis.

7. Concentration

Supply concentration describes how activity is distributed across the relevant production or refining system. It is not the same object as responsiveness. Responsiveness asks how the system can adjust. Concentration asks where activity sits, or who controls it. A concentrated market can be elastic if capacity is idle and expandable. A fragmented market can be inelastic if every producer faces the same decade-long development path.

Where a reliable set of market shares exists, the preferred concentration factor is a normalised Herfindahl–Hirschman index at a named stage:

$$HHI_m = \sum_k s_{k,m}^2, \quad \Gamma_m = 100 \times (HHI_m - HHI_{min}) / (HHI_{max} - HHI_{min})$$

$s_{k,m}$ is the share of producer or country k at the stated stage. The normalisation maps the index onto 0–100 over a declared range. Country shares and firm shares are not substitutes for one another.

The current implementation does not yet compute that index for every mineral. It uses a structured concentration rubric, scored against named evidence, and that rubric should not be described as a calculated HHI. Assessments have to state whether concentration refers to countries, firms, mining or refining. Geographic dependence does not establish firm-level market power. Refining concentration can be the economically relevant object when mined supply is diverse and conversion is not, which is a familiar pattern in several battery and speciality metals.

Γ enters the company formulas below as a 0–100 factor. Switching a mineral from the rubric to a calculated HHI is a methodology change if it moves scores, and it should be versioned as one. It is not a silent data refresh.

8. From minerals to listed companies

8.1 Production exposure

The company connection is production exposure: the company's output of a mineral relative to world output, on a compatible basis.

$$E_{c,m} = q_{c,m} / Q_m$$

$q_{c,m}$ is company output. Q_m is world output. Numerator and denominator must share a period, a material form, a unit and a stage of production.

Compatibility is a substantive condition, not a formatting one. Contained metal divided by an oxide total is not a share, unless a conversion has been applied and recorded. Refinery output divided by a mine-production denominator mixes stages. Equity-attributable output is not the gross output of an operated asset. Joint ventures, toll processing and successive stages of the same material flow can manufacture an apparent exposure that does not correspond to a comparable production share. Where the basis cannot be reconciled, the quantitative exposure is left unranked. It is not set to zero, and it is not replaced with an unsupported estimate.

A share of global mineral production is not the share of revenue, profit or enterprise value attributable to that mineral. It is not ownership of reserves, and it is not potential future capacity. Shares across different minerals are useful weights inside a company formula. They do not sum to a meaningful share of world mineral production, because the denominators are different commodities. That last point matters for diversified miners: a modest share of a constrained speciality and a modest share of a bulk commodity are not two observations of the same fact.

The database distinguishes producers, processors, developers, explorers and investment vehicles. A qualitative link is the right record when the production measure is absent. Developers and explorers are not assigned current production exposure because a project might supply a mineral later. An investment vehicle would need a separate look-through before holdings could be treated as physical output. The producer-oriented score does not quantify the position of a company that buys the mineral. Input-cost and availability risk for downstream users is a different model, and it is not implied by a low MES.

8.2 Mineral Exposure Score

Mineral Exposure asks how far a company's measured output sits in minerals that combine constrained supply with demand pressure under the scenario.

$$MES_c^{raw} = \sum_m (E_{c,m} \times SES_m \times DP_m)$$

DP_m is the demand-pressure term, the DPS value on 0–100 in the current implementation. The raw sum is then mapped to 0–100 by the fixed normalisation in Section 9.

A high MES means the eligible production book is tilted toward minerals whose supply is difficult to expand or replace and whose scenario demand pressure is high. It does not mean the company is a pure play, and it does not mean the mineral dominates earnings. A broader production portfolio dilutes the aggregate. That

dilution is information. The company score is not a pure measure of the intensity of a single asset, which is why the mineral-level decomposition has to travel with the summary.

8.3 Supply Leverage Score

Supply Leverage asks a different question: how far measured output sits in minerals that are both constrained and concentrated.

$$SLS_c^{raw} = \sum_m (E_{c,m} \times SES_m \times \Gamma_m)$$

Γ_m is the concentration factor on 0–100. Demand does not enter. A demand revision leaves SLS unchanged if exposure, SES and concentration are unchanged.

The phrase supply leverage describes a physical position in this framework. It does not establish control over prices, contractual bargaining power, or the ability to earn excess returns. A three percent share of a concentrated and inelastic mineral is a different research object from a three percent share of a fragmented and flexible one. It is still three percent. Market power is not awarded by the formula.

8.4 Critical Mineral Catalyst Score

CMCS is the summary company signal. It places the greatest weight on mineral relationships in which production exposure, supply constraint, demand pressure and concentration occur together.

$$CMCS_c^{raw} = \sum_m (E_{c,m} \times SES_m \times DP_m \times \Gamma_m)$$

The product inside the sum is intentional. A missing or negligible term pulls the contribution toward zero. A thematic association cannot compensate for the absence of substantiated production exposure.

The name is a platform label, and it is a poor description of the contents. CMCS does not incorporate entries from the qualitative catalysts layer. A substitution announcement, a policy proposal or a laboratory result does not move CMCS. The score moves when an eligible input moves: exposure, a constraint component, the demand scenario, or concentration. Users who read the word catalyst as a claim about news flow will misread the number. The decomposition is the correction.

Indicator	Level	Question it answers
SES	Mineral, 0–100	How difficult is it to expand or replace supply? Higher means less responsive. Not a price elasticity.
DM	Mineral or company, percent	What is the growth outlook under the stated shares and assumptions? Company DM is the covered mix only.
DPS	Mineral, 0–100	What is the normalised demand term used in company scoring? Currently a transform of the same proxy as DM.
Γ	Mineral, 0–100	How concentrated is the named stage? Rubric in the current implementation; HHI where share data are used.
E	Company–mineral, share	What is compatible current output relative to world output at that stage?
MES	Company, 0–100	How far does measured output sit in constrained minerals under the demand scenario?
SLS	Company, 0–100	How far does measured output sit where constrained supply is also concentrated?
CMCS	Company, 0–100	Where do exposure, constraint, demand pressure and concentration coincide?

Table 2. Principal objects. Company scores on 0–100 are normalised raw sums, not probabilities and not percentile ranks.

9. Normalisation, bounds and the meaning of a point

Raw sums are not on a common scale. MES raw multiplies a share by two 0–100 factors; CMCS raw multiplies a share by three. A company with exposure to several minerals accumulates contributions. The published scores apply a fixed affine map and then bound the result:

$$Score_c = \min(100, \max(0, a + b \times Raw_c))$$

The constants a and b are calibration parameters, fixed inside a methodology version. They are not re-estimated when an unrelated issuer is added to the database.

Fixed bounds are a design choice with a specific purpose. A moving cross-sectional normalisation would make every existing score depend on who else had been added to the universe. That is an undesirable property in a research database: an analyst comparing two vintages could not tell a change in the physical evidence from a change in the peer set. The cost of fixed bounds is saturation. Several exposures can share an upper score despite different raw contributions. Component values and raw contributions should be retained, and saturation should be reported rather than interpreted as equality of economic position.

A score is not a percentile rank, not a probability, and not a unit of economic value. A value twice as high does not imply twice the rent or twice the expected return. Comparability across companies still depends on consistent definitions, data vintages and the same calibration. Any change to the constants, the weights or the demand identity is a methodology version, with an explicit rule for historical comparison. The numerical endpoints of the current calibration are implementation parameters. This paper discloses the aggregation identities so that the transmission logic can be inspected; it does not publish the affine constants or the component-rubric thresholds.

Financial interpretation starts where the score stops. The mapping to public equity ends at the issuer's physical positioning. The subsequent questions are whether that positioning is material, whether it is commercially accessible to the listed entity, and whether it is already in the valuation. Those questions are not optional extras. They are the reason a high CMCS is a research prompt rather than a conclusion.

10. A worked transmission scenario

The following case is constructed. Mineral M, mineral Q and mineral N are not descriptions of particular commodities, and Companies A, B and C are not descriptions of particular issuers. The arithmetic uses the identities in Sections 5 to 8. Where a demand-pressure value is required inside MES and CMCS, it is stated as a scenario input, not estimated from a calibrated transform. The point of the case is to show which objects move, and which do not.

10.1 Two minerals, one SES, different constraints

Mineral M is constrained mainly by time and processing. Mineral Q is constrained mainly by substitution and recycling. Both receive SES of 71 under equal weights.

Component	Mineral M	Mineral Q
Time to add supply, T	78	50
Capital intensity, K	64	60
Processing bottleneck, P	86	55
Substitution difficulty, S	55	95

Component	Mineral M	Mineral Q
Recycling difficulty, R	72	95
SES = 0.20 × sum	71.0	71.0

Table 3. Identical summary scores, different mechanisms. A processing subsidy and a substitution breakthrough are not interchangeable events for these two profiles.

For mineral M, $SES_M = 0.20 \times (78 + 64 + 86 + 55 + 72) = 0.20 \times 355 = 71.0$. For mineral Q, $SES_Q = 0.20 \times (50 + 60 + 55 + 95 + 95) = 0.20 \times 355 = 71.0$. A research note that quoted only the summary would treat a reduction in separation lead time as equivalent to the commercial qualification of a substitute. The component profile does not.

10.2 Demand under an explicit allocation

Mineral M's baseline demand is allocated across four sectors. The shares exhaust the allocation. Batteries and electric vehicles are not both credited with the same tonne.

Sector	Demand share ω	Growth assumption g	Contribution $\omega \times g$
Batteries	0.45	18%	8.10
Electricity grids	0.25	9%	2.25
Defence	0.10	6%	0.60
Residual uses	0.20	2%	0.40
DM (share-weighted)	1.00	—	11.35%

Table 4. Baseline demand proxy for mineral M. $DM_M = 11.35$ percent. This is a weighted average of assumed sector rates, not the compound growth rate of a quantity account, and not an observed consumption series.

Raising the battery growth assumption from 18 percent to 24 percent, with shares and other sector rates held fixed, moves DM_M from 11.35 percent to 14.05 percent. SES does not move. Concentration does not move. Company production shares do not move. Any change in a company indicator has to come through the demand term, or it is a bug.

10.3 Three companies, one mineral association

Company A has verified current output: 3.5 percent of compatible world output of M, and 1.0 percent of compatible world output of N. Mineral N is the contrast case, with $SES_N = 40.0$, $DM_N = 4.0$ percent, $DPS_N = 30$ and $\Gamma_N = 41$. For the baseline scenario, $DPS_M = 68$ and $\Gamma_M = 74$. Company B has a prospective project and no current output. Company C buys M for a downstream product. Only Company A has an eligible E.

Company A's demand momentum on the covered book is $DM_A = (0.035 \times 11.35 + 0.010 \times 4.0) / 0.045 = 9.72$ percent. The raw mineral-exposure contribution is dominated by M: $0.035 \times 71 \times 68 = 168.98$, against $0.010 \times 40 \times 30 = 12.00$ from N. Raw MES = 180.98. Raw SLS = $0.035 \times 71 \times 74 + 0.010 \times 40 \times 41 = 200.29$. Raw CMCS = $0.035 \times 71 \times 68 \times 74 + 0.010 \times 40 \times 30 \times 41 = 12,996.5$. These raw sums are not published scores. The published scores are the image of the raw sums under the fixed map in Section 9.

Under an illustrative assumption that DPS scales with the growth proxy until it hits the upper bound, the battery revision lifts DPS_M from 68 to 84.2. Raw MES rises from 180.98 to 221.18. Raw CMCS rises from 12,996.5 to 15,971.2. Raw SLS stays at 200.29, because the revision did not touch exposure, SES or concentration. Company DM rises from 9.72 percent to 11.82 percent. The assumption of a linear DPS map is a device for the example. The production calibration may compress the same revision once a bound is near. Saturation is part of the object an analyst has to check.

Entity	Physical relationship	Effect of a higher battery growth rate
Mineral M	Supply profile held fixed	DM rises from 11.35% to 14.05%. DPS may rise, subject to bounds. SES and Γ unchanged.
Mineral Q	Same SES as M, different profile	Unaffected, unless the scenario also revises Q's demand allocation. Identical SES does not transmit the shock.
Company A	Verified current output of M and N	Demand-sensitive raw terms rise. SLS unchanged. No return forecast follows.
Company B	Prospective development	The project remains on the research map. Current production exposure stays absent, so MES, SLS and CMCS do not acquire a contribution.
Company C	Downstream buyer	Procurement and margin risk may be economically real. The producer-oriented score does not quantify it.

Table 5. Transmission under the constructed scenario. The discrimination is the result: thematic association is not enough to move a company score.

A second scenario would run the other way. Evidence that a substitute reduces M's use in batteries can support a lower demand assumption, if it is commercially credible on the relevant horizon. Evidence that substitution has become practical across major uses can support a reviewed change to the structural substitution component. Those are separate edits, against separate evidence. A catalyst entry triggers neither by itself.

The useful output of the exercise is an explanation, not a ranking. It names the mineral, the revised assumption, the companies with measured participation, and the financial links that are still missing. Another analyst can dispute the battery share, the source of the 3.5 percent, or the decision to treat grids as a separate use, without having to accept or reject a general scarcity thesis.

11. Evidence, coverage and catalysts

11.1 Provenance

Every material input needs a source, an observation period, a unit and a status. Official statistics, company disclosures, analyst judgements, scenario assumptions and derived estimates are not the same kind of number, and the record should not make them look as if they were. USGS commodity tables can support a production denominator. The table, the year and the material basis still have to be checked mineral by mineral. A country production figure is not world output. A reserve estimate is not annual supply. A production quantity displayed beside a demand discussion is not an observation of consumption, and it cannot establish a deficit.

The current database has limitations that bear directly on interpretation. Some company production entries were back-calculated from legacy shares rather than taken from a primary output disclosure. Certain rare-earth denominators were apportioned from a combined oxide total, so the element-level figures are estimates. Some historical price series were repeated or interpolated. Indicative price histories cannot support event studies or return-prediction claims as if they were observed market series. These are properties of the present implementation. They preclude treating every input as an independent observation. Trade-flow scoring is outside the core method.

Missing production, or a missing world-supply denominator, leaves the corresponding quantitative exposure unranked. A company with partial coverage needs a visible statement of which mineral relationships contribute to its score. A count of covered relationships is useful. It does not measure the share of revenue or of economic exposure that has been covered. Company DM likewise describes the eligible subset.

11.2 Identity, units and reconstruction

Quality controls have to test unit consistency, plausible production shares, complete demand allocations and issuer identity. Company output above the compatible world total is an error until it is explained, not a sign of unusual leverage. Alternate listings should map to the same issuer where the economics are the same; unresolved identities should not be promoted into verified investable exposure. Updating a value should keep the previous observation. Erasing it destroys the evidence an earlier user relied on.

Published scores are derived from common inputs and common rules, which removes one source of inconsistency across pages. Central calculation does not make a past result reproducible. Reconstruction needs a dated input snapshot, the source revision, the methodology version and the scenario assumptions in force at the time. Without those, a historical score is a label, not a calculation that can be rerun.

11.3 Catalysts stay outside the sum

Catalysts are technological, corporate and policy developments that might alter supply or demand. A useful record names the mineral, the likely direction, the stage of development and the confidence of the evidence. Presence on that record does not change SES, DM, DPS, Γ , MES, SLS or CMCS. A substitution announcement can be economically empty for years if qualification, cost or production scale remain unresolved. Scoring it on arrival would convert a monitoring list into a false precision.

When the evidence is sufficient to revise a core assumption, the revision should name the affected input, the source and the effective date. That preserves the distinction between watching a possibility and incorporating it. It also lets a user see whether a score moved because the physical evidence changed, because the demand assumption changed, or because the methodology changed. Those three events have different implications, and a single revised number will not say which one occurred.

12. Evaluation agenda

This paper specifies the tests. It does not report completed validation. The order matters. A return study run on economically incompatible inputs would be a precise answer to the wrong question.

The first test is data validity. Mineral units and production stages have to be reconciled. Issuer mappings have to be verified. Derived quantities, including back-calculated shares and apportioned rare-earth denominators, have to be flagged and, where possible, replaced. Coverage has to be measured. A score can be calculated consistently from inputs that do not refer to the same thing. Computational consistency is the initial test, not the interesting one.

The second test is the reliability of the judgement-based assessments. Independent reviewers should apply the SES and concentration rubrics to the same evidence pack, with disagreement reported and then reconciled. If reviewers cannot reproduce a processing-bottleneck score from the cited sources, the score is a private note, whatever its number of decimal places. Sensitivity analysis should vary plausible component scores, drop one overlapping term at a time, swap rubric concentration for an HHI where shares exist, and recompute company scores under incomplete coverage. The reportable objects are rank changes, score changes and the incidence of saturation. Conclusions that survive those perturbations are different in kind from conclusions that depend on a single component.

The third test is coherence with the declared identities, of the sort Section 10 illustrates. Raising a sector growth rate, with physical inputs held fixed, should move DM and the demand-sensitive company terms and should leave SES and SLS alone. Adding an unrelated issuer should leave existing scores alone under fixed calibration. Setting a production figure to missing should remove that contribution and should not silently

insert a zero that still occupies the denominator. These are tests of the implementation. Passing them is not evidence of economic predictive power.

Empirical work on the supply side can then ask whether the constraint profiles relate to observed supply responses, commissioning delays or production changes during documented shocks. That work needs dated inputs, a stated horizon and a comparison group. It has to separate structural responsiveness from outcomes driven by inventories, export restrictions, demand destruction or a failure at a single asset. A mineral with a high T that in fact delivers new tonnes on a short horizon is a finding about the rubric, or about the episode, and either finding is useful.

Any study of equity returns is a separate design. It would need point-in-time company universes, a treatment of delistings and corporate actions, controls for commodity and market exposures, and an out-of-sample period. A trading claim would also have to survive transaction costs. The test should separate information already embedded in prices from subsequent developments. Until that evidence exists, the framework supports exposure research and scenario analysis. It is not an asserted source of excess returns.

13. Discussion and limitations

The framework is most useful when an analyst already has a material demand thesis and needs a public-company supply position that can be checked. The workflow is deliberately narrow: identify the material, read the adjustment profile rather than only the average, examine the demand allocation, and assess the issuer's eligible production share. Decomposition is what makes a company signal available for scrutiny. A hero number without the mineral contributions is a branding choice, not a method.

Sector growth does not uniquely determine mineral demand, particularly when technology changes material intensity. The lithium example is the obvious one, but it is not special: anode chemistry, motor design and grid topology do the same work for graphite, rare earths and copper. Supply responsiveness also varies by asset, region and stage, while a mineral-level SES compresses those differences into five numbers. Global concentration may say nothing about a particular buyer's access to qualified supply. By-product dependence can limit output for reasons that belong to the host commodity. A descriptive flag does not model that relationship, and the capital-intensity component should not be asked to do so in disguise.

The company mapping favours entities with comparable production disclosures. Economically important processors, recyclers and development projects can fall outside the quantitative scores. Comparisons will tilt toward better-disclosed producers. Partial scores cannot be treated as complete company assessments. Aggregation across minerals, and correlation among constraint components, can also move a ranking for reasons that are arithmetic rather than geological. The overlap between processing bottlenecks and concentration, already noted, is the clearest case.

Readers of this preprint can inspect the aggregation identities, the demand allocation rule and the proposed tests. They cannot reproduce a live score from the paper alone. Component rubrics, affine calibration constants and the versioned input database remain implementation objects. That boundary should stay explicit in any claim about transparency. Disclosing the product structure of CMCS is a real disclosure. It is not the same thing as releasing the scoring sheets.

Extensions that would earn a place in a later version are specific. Better material-intensity evidence would allow the product identity in Section 6.2 to replace the sector proxy, provided double-counting is closed. Stage-specific concentration would allow mining and refining to stop sharing a single Γ . Equity-attributable production, recorded separately from operated output, would stop joint ventures from inflating shares. A downstream cost-exposure model would give Company C a score of its own, instead of a disclaimer. Each of

those requires units, an overlap check and a versioned change. A plausible story is not a reason to insert a term into the sum.

14. Conclusion

The Stockatomics Constraint Transmission Framework writes down the path from downstream activity to mineral constraints and on to listed-company production. Growth outlook is not supply responsiveness. Concentration is not firm-level control. A production share is not a revenue share. Physical positioning is not an equity return. The indicators exist so that those distinctions survive contact with a thematic research question.

The current method uses an equal-weighted supply profile, a sector-share demand proxy, a concentration factor, and company production exposure. MES, SLS and CMCS are normalised sums of stated products of those objects. Products and catalysts explain and challenge the assumptions. They do not enter the sums. Detailed financial consequences remain outside the core scores on purpose.

The contribution is a research architecture whose usefulness depends on compatible units, visible coverage and scenario analysis that names the input it changed. Empirical evaluation of supply outcomes, and any later evaluation of investment performance, are still to be done. Until they are done, a high score is a claim about physical positioning under a stated scenario. It is not a claim about what the equity will do.

Appendix A. Notation

Symbol	Meaning
m, c, j, i	Mineral, company, sector, product.
T, K, P, S, R	SES components: time, capital intensity, processing, substitution, recycling. Each on 0–100.
SES_m	Equal-weighted supply-constraint score. Not a price elasticity.
$\omega_{m,j}, g_j$	Demand share of mineral m in sector j , and the sector growth assumption.
DM_m, DM_c	Mineral and company demand-momentum rates, in percent.
$\iota_{i,m}, Q_i$	Mineral intensity and scale in the preferred product-pressure identity.
DP_m, DPS_m	Demand-pressure term used in company scores, on 0–100.
$s_{k,m}, HHI, \Gamma_m$	Stage share, Herfindahl–Hirschman index, and the 0–100 concentration factor.
$q_{c,m}, Q_m, E_{c,m}$	Company output, world output, and the compatible production share.
MES, SLS, CMCS	Normalised company scores defined in Section 8.
a, b	Fixed affine calibration constants inside a methodology version.

Appendix B. Version stance

The identities in Sections 5 to 9 are the scoring specification associated with this preprint. Equal SES weights, the sector-share demand proxy, the separation of catalysts from the core sums, and fixed rather than cross-sectional normalisation are part of that specification. Changes to weights, to the demand identity, or to the calibration constants require a documented revision. The objective is ordinary and strict: two pages in the database should not be implementing two methods.

References

- Acemoglu, D., Carvalho, V. M., Ozdaglar, A., and Tahbaz-Salehi, A. (2012). The network origins of aggregate fluctuations. *Econometrica*, 80(5), 1977–2016. <https://doi.org/10.3982/ECTA9623>
- Carvalho, V. M., Nirei, M., Saito, Y. U., and Tahbaz-Salehi, A. (2021). Supply chain disruptions: Evidence from the Great East Japan Earthquake. *The Quarterly Journal of Economics*, 136(2), 1255–1321. <https://doi.org/10.1093/qje/qjaa044>
- 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. <https://doi.org/10.1073/pnas.1500415112>
- International Energy Agency. (2026). *Global Critical Minerals Outlook 2026*. IEA, Paris. <https://www.iea.org/reports/global-critical-minerals-outlook-2026>
- OECD/European Union/EC-JRC. (2008). *Handbook on Constructing Composite Indicators: Methodology and User Guide*. OECD Publishing, Paris. <https://doi.org/10.1787/9789264043466-en>
- Stockatomics. (2026). *A financial data platform for the physical economy, connecting critical minerals, products, sectors and public companies*. <https://www.stockatomics.com>
- U.S. Geological Survey. (2026). *Mineral commodity summaries 2026* (ver. 1.3, May 2026). U.S. Geological Survey, 222 pp. <https://doi.org/10.3133/mcs2026>

Disclosure

Stockatomics develops the platform and the methodology described in this paper. The author is affiliated with Stockatomics. This preprint is not peer reviewed. The worked scenario uses constructed figures and reports no observed company outcome. No empirical investment-performance validation is reported. Aggregation identities are disclosed so that the transmission logic can be examined. Component-rubric thresholds, affine calibration constants and the underlying database are not released with this paper. The method is intended to support research into physical-economy exposure. It is not a recommendation concerning any security.