The Circularity Gap Report: Critical Raw Materials

Methodology document

Coupled System Dynamics and Dynamic Material Flow Analysis Methodology (SD-dMFA)
Version 1.0
19 August 2026

CIRCLE ECONOMY

Contents

About this document **4**

Goal and scope **5**

Overview 5

Design concepts 9

Details 11

Model architecture **14**

Dynamic MFA formulation 14

System Dynamics formulation 27

Coupling methodology 33

Data **36**

Exogenous input data 36

Parameters and variables 71

Calibration **82**

What is calibrated, and what is not 82

Initialisation: Seeding the pre-study-period stock 83

Historical activation: Testing behaviour against documented episodes 85

Behavioural tuning of the System Dynamics surface 88

Automated parameter fitting: A shipped capability, not the deployed method 90

Sources for the calibration and initialisation evidence base 92

Scenarios **94**

The three axes 94

Axis 1 — Demand 95

Axis 2 — Circularity ambition 96

The wedge ladder: Making results attributable 100

Axis 3 — Shocks and resilience 100

The executed run surface 102

Feasibility bounds and an open finding 103

Sources for the scenario evidence base 104

Assumptions and limitations **107**

How this section is organised 107

Deliberate structural simplifications 107

Evidence and data limitations 108

Temporary states of the current configuration 110

How results should and should not be read 111

Verification: What would catch an error 112

Glossary **114**

Demand and service 114

Material flow and mass balance 115

Behavioural layer and coupling 116

Scenarios and circularity levers 117

Indicators and resilience 118

About this document

This document is the technical reference for the SD-dMFA model. It describes the model as deployed for the *Circularity Gap Report: Critical Raw Materials*, and it is structured according to the Overview, Design Concepts, Details (ODD) protocol for describing simulation models (Müller et al., 2014).[1]

Two distinctions run through the document and are worth stating at the outset. The first is between *general model capability* and *active configuration*. The framework implements more than the study reported here uses: a regional origin–destination trade layer, an automated parameter-fitting workflow, and a set of diagnostic scenario ablations are shipped and runnable but are not part of the executed configuration. Where that is the case, the capability is noted and then explicitly marked inactive, rather than either omitted or presented as though it were in use. The second is between the *reporting window* and the *data extent*. Results are reported for 1950–2050, but the exogenous data and the model's warm-start run wider than that —each input artefact spans 1870–2100, and copper's initialisation begins in 1870. A figure quoted for a year outside 1950–2050 is therefore a property of the underlying series, not a reported result.

Values quoted throughout are those held in the executing configuration. Where a parameter's value in a shared defaults file differs from the value that actually runs, because a later configuration layer overrides it, the document states the value that runs and, where the difference is instructive, both.

Model configuration described

Run entrypoint	cgr-crm (extending cgr-crm-base), the canonical configuration for the reported results
Coupling preset	loose_iterative_sd_full_endogenous_v1 — full SD layer with endogenous circular loops active
Circular-loop surface	circular_loop_response_reman_low_v1
Materials	Copper, nickel, cobalt, lithium, silicon (20 material streams)
Reporting window	1950–2050, annual time step
Data extent	1870–2100, annual (231 years), single global region
Demand scenarios	STEPS, APS, NZE (IEA); active baseline runs STEPS at the L circularity rung
Configuration verified	19 August 2026

[1] Muller, E., Hilty, L. M., Widmer, R., Schluep, M., & Faulstich, M. (2014). Modeling metal stocks and flows: A review of dynamic material flow analysis methods. Environmental science & technology, 48(4), 2102-2113. doi:10.1021/es403506a

Goal and scope

Describes the model's purpose, entities, scale and design concepts. The underlying framework is more general than the configuration described here, and general capability is distinguished from active application throughout.

Circle Economy's SD-dMFA model answers a specific empirical question: under a given trajectory of demand, primary-supply conditions and circularity ambition, how much of the material requirement for copper, nickel, cobalt, lithium and silicon can realistically be met from primary and secondary sources between 1950 and 2050, and what happens to service delivery when it cannot? Answering it requires two things neither a pure stock-flow accounting model nor a pure behavioural model provides alone: a physically consistent account of where material is and how it moves — mass balance, cohort ageing, routing through recycling and disposal — and an explicit representation of how demand, collection effort and reserve policy adapt under stress. SD-dMFA is architected as two coupled submodels for that reason: a dynamic material flow analysis (dMFA) layer enforcing physical conservation and a system dynamics (SD) layer representing adaptive behaviour, connected through a narrow, auditable set of feedback signals rather than merged into a single equation set.

SD-dMFA provides the technical basis for the Circularity Gap Report: Critical Raw Materials (CGR-CRM); that report's own methodological explanation chapter (Chapter 2: Measuring the circularity of critical minerals) explicitly defers full description of data sources, parameterisation and equations to this document. The remainder of this section defines the model's scope formally, following the Overview, Design Concepts, Details (ODD) protocol proposed by Müller et al. (2014) for stock-flow and material-cycle models. Model architecture, data provenance, calibration and scenario design are each developed in their own sections below; this one establishes what the model represents, at what resolution, and why.

Overview

Purpose. SD-dMFA evaluates how primary supply, circular material flows, in-use stock turnover, and adaptive behavioural responses jointly determine whether material-dependent service demand for five critical raw materials can be met over a 1950–2050 horizon. In its current application, it quantifies primary and secondary supply contributions, in-use stock accumulation and retirement, end-of-life routing through recycling, remanufacturing and disposal, delivered versus unmet service, and a set of circularity and resilience indicators, under a baseline configuration and a structured set of scenarios (decarbonisation pathway × circularity ambition × resilience stress test) described in full in the Scenarios section.

Material selection. Copper, nickel, cobalt, lithium and silicon were selected from a longer shortlist of candidate raw materials (including aluminium, graphite, rare earth elements, gallium and tantalum), scored against data availability, short- and long-term relevance, by-product interlinkage, coverage across the value chain, supply risk, and modelling leverage, then confirmed by project-team selection. The five chosen span both host materials and by-product-dependent materials, and together cover applications from bulk conductors and structural alloys to battery chemistries and semiconductors.

Materials and system entities. The model's material dimension (symbol *m*) currently comprises five critical raw materials — copper, nickel, cobalt, lithium and silicon — each tracked in refined-metal-equivalent mass terms (tonnes). Within each material, the model resolves a *material stream* dimension (symbol *s*): a labelled refined-equivalent flow, differentiated by application grade, that activates at the refining stage and is carried through each downstream flow and stock. Twenty material streams are currently defined — four each for copper, silicon and nickel, three for lithium, and five for cobalt — chosen because they differ materially on at least two of three design axes: geopolitical trade concentration, process yield profile, and downstream product eligibility. Table 2 gives the full classification.

Processes. The active physical chain runs mining and beneficiation → smelting → refining → semifinished-product fabrication → finished-product manufacturing → use → collection → sorting/preprocessing → remanufacturing / secondary refining/residue treatment and disposal → system environment. Material stream identity is not carried upstream of refining: primary ore input is allocated to this chain through a source/route allocation layer (deposit type, ore grade, processing route, and route yield) that determines how much primary refined-equivalent material becomes available before the stream-resolved part of the chain begins. Table 3 sets out each stage's role and its primary inputs and outputs.

Spatial scale and extent. The active run scope is global—the model does not currently disaggregate mining geography, refining location, or trade flows by country or bloc. This is a scope choice, not a structural ceiling: the codebase implements a full origin–destination (OD) trade layer, operating over refined metal, concentrates and scrap, that can allocate flows across multiple regions once a regional scope is authored. That capability is disabled in the active `cgr-crm` baseline preset, and statements about primary and secondary supply, service delivery, and resilience in this document describe outcomes for the pooled global system, not for any sub-global geography.

Temporal scale and extent. The model runs at an annual time step over a 1950–2050 horizon. Within that horizon, 1950–2019 serves as the calibration window (against which historical plausibility is assessed) and 2020–2050 as the reporting window (over which scenario results are compared). Table 4 sets this out; per-material burn-in treatment ahead of 1950, used to seed physically plausible initial stocks, is described in the Calibration section.

State variables. Three stock variables carry the model's physical state forward in the dMFA layer — in-use stock, a refinery stockpile, and an optional strategic inventory — and two carry behavioural state forward in the SD layer — a collection-effort integrator and a capacity envelope. Table 5 defines each; the Model Architecture section develops their governing dynamics in full.

Process overview and scheduling. A full run resolves the active configuration, loads the exogenous and calibration-side input tables, and builds an SD desired-demand backbone from `final_demand.csv`. For each material–region slice, the model then iterates the SD and dMFA submodels: SD is run over the full year vector, its outputs are used to construct dMFA inputs, dMFA is run over the full year vector, feedback signals are recomputed from the dMFA result, the signals are smoothed, and convergence is evaluated — repeating until the joint trajectory stabilises or a maximum iteration count is reached. This is a full-horizon iterative reconciliation, not a year-by-year co-simulation: within one iteration an entire submodel is solved before control passes to the other. Once the slice converges, dMFA splits realised service demand across the product/stream mix, routes

end-of-life flows through the collection–recycling chain, updates in-use stock and the refinery stockpile (and, where active, the strategic inventory), balances primary against secondary availability, and computes delivered versus unmet service. Table 1 gives the model dimensions and the resolution each takes in the active configuration.

Table 1. *Core model dimensions and their active resolution.*

Symbol	Dimension	Active values	Status
t	Time	1950–2050, annual	Active
m	Material	copper, nickel, cobalt, lithium, silicon (5)	Active
r	Region	Global (1)	Active — OD-trade layer supports multi-region allocation but is disabled in the active baseline
s	Material stream	20 streams across the 5 materials (Table 2)	Active
z	Product	canonical (1 placeholder)	Placeholder — a richer product taxonomy is designed (see Assumptions and Limitations) but not yet populated
d	Deposit type	canonical (1 placeholder)	Placeholder — source-allocation axis, not a native flodym dimension; a multi-deposit taxonomy is designed but not yet activated
p	Processing route	canonical (1 placeholder)	Placeholder — source-allocation axis, not a native flodym dimension; a multi-route taxonomy is designed but not yet activated

Table 2. *Material-stream classification (20 streams across 5 materials). Streams are separated where they differ on at least two of: geopolitical trade concentration, process yield profile, and downstream product eligibility.*

Material	Material streams (s)
Copper (4)	wire_rod_grade · alloy_grade · tube_sheet_plate_grade · electronic_grade
Silicon (4)	ferroalloy_grade · chemical_alloying_grade · solar_grade_polysilicon · semiconductor_grade_polysilicon
Lithium (3)	battery_grade · industrial_grade · specialty_organolithium_metal
Nickel (4)	battery_grade_sulfate · class_i_refined_metal · class_ii_stainless_feed · chemicals_catalysts
Cobalt (5)	battery_grade_sulfate · superalloy_grade_metal · hardmetal_grade_powder · magnets_grade_metal · chemicals_pigments_catalysts

Table 3. *Active process topology (12 stages).*

Stage	Role	Primary inputs	Primary outputs
mining_beneficiation	Reconstructs primary ore input into beneficiated material via deposit share and ore grade	primary_ore_input, deposit_share, ore_grade	Beneficiated material to smelting
smelting	Converts beneficiated material toward refined-equivalent form	Beneficiation output, route_share	Smelter output to refining

Stage	Role	Primary inputs	Primary outputs
refining	Produces primary refined-equivalent material; material-stream identity activates here via route_stream_share	Smelter output, route_yield, secondary inputs	Primary refined material by stream
semis_fabrication	First manufacturing transformation; generates new (pre-use) scrap	Refined material by stream	Semi-finished product; new scrap
finished_product_manufacturing	Final manufacturing transformation; generates new scrap	Semi-finished product	Finished product to use; new scrap
use	In-use stock accumulation and cohort-based retirement	Finished product inflow, remanufactured inflow	stock_in_use; outflow to collection
collection	Gate 1 of end-of-life routing: collection_rate determines collected vs. uncollected fraction	Outflow from use	Collected material; uncollected → system_losses
sorting_preprocessing	Prepares collected material for downstream routing; generates sorting rejects	Collected material	Sorted material; rejects → system_losses
remanufacturing	Gate 2 (one branch): returns suitable cores to use	Sorted material (remanufacturing_rate share)	Reman flow back to use
recycling_refining_secondary	Gate 2 (one branch): recovers secondary refined material	Sorted material (recycling_rate share)	Secondary supply to refining
residue_treatment_disposal	Gate 2 (one branch): downcycled/non-recovered material and process losses	Sorted material (downcycling_rate share); process losses	→ system_losses
sysenv	Terminal sink for all pre-use and post-use losses	All loss streams	— (system boundary)

Table 4. *Temporal scope.*

Period	Years	Role
Full modelled horizon	1950–2050	Complete simulated period, annual resolution
Calibration window	1950–2019	Historical period against which modelled trajectories are assessed for plausibility
Reporting window	2020–2050	Forward-looking period over which baseline and scenario results are compared
Pre-1950 burn-in	Material-specific (10–80 years)	Synthetic stock spin-up ahead of 1950 to seed physically plausible initial in-use stock; length and method are material-specific and detailed in the Calibration section

Table 5. *Core state variables.*

Variable	Layer	Description
stock_in_use	dMFA	Cohort-resolved in-use material stock; evolves through product inflows and lifetime-distribution-governed retirements
refinery_stockpile_native	dMFA	Buffer stock between variable secondary-material recovery and refining-stage input requirements
strategic_inventory_native	dMFA	Strategic reserve stock, filled/released on SD-generated intents; active in the current baseline for nickel, lithium and cobalt (material-specific thresholds), neutralised via a zero target-coverage setting for copper and silicon
collection_effort	SD	Integrator state driving the collection multiplier applied to the baseline dMFA collection rate, responsive to price pressure and bottleneck conditions
capacity_envelope	SD	Lagged aggregate index of productive-capacity adequacy, driving bottleneck pressure and, through it, price and demand response

Table 6. *Required exogenous inputs (overview level; full column-level detail is given in the Data section).*

File	Feeds	Purpose
final_demand.csv	SD	Desired final demand backbone by material and year
primary_ore_input.csv	dMFA (mining_beneficiation)	Authored primary supply anchor
deposit_share.csv, ore_grade.csv	dMFA (mining_beneficiation)	Allocates ore input across deposit types; converts ore mass toward metal content
route_share.csv, route_yield.csv	dMFA (smelting, refining)	Allocates material across processing routes; applies stage-yield losses
route_stream_share.csv	dMFA (refining)	Maps route output to the 20 material streams at the refining split
product_share.csv	dMFA (use)	Allocates real demand to the (currently single) product category
collection_routing_rates.csv	dMFA (collection onward)	Two-gate routing surface: collection_rate, then recycling/remanufacturing/downcycling split
lifetime_distributions.csv	dMFA (use)	Cohort lifetime distribution parameters governing retirement from use

Design concepts

Basic principles. The model is dynamic, mass-balanced, stock-flow consistent, cohort-based in the use phase, and loosely coupled across its SD and dMFA layers. Physical conservation is enforced entirely within dMFA; SD never overrides a mass balance; it only changes the demand, collection and reserve trajectories that dMFA is subsequently re-solved under.

Dynamic modelling approach. dMFA is inflow-driven: each year's inflow to use is treated as a distinct cohort with its own lifetime probability distribution, and outflow from use in any given year is the convolution of historical inflows with the relevant retirement profile. SD is deterministic and

reduced-form: scarcity-price-demand response, capacity-envelope adjustment, collection-effort response, and strategic-reserve fill/release intent are each governed by a small number of calibrated gains, lags and elasticities rather than by explicit household, firm or regulator decision models. Layered on top of this reduced-form core is a background stochastic noise process on demand and collection, active from each material's historical-dynamics start year, with autocorrelation and variance parameters calibrated per material to reflect real-world demand and collection uncertainty. This noise process is distinct from, but composes with, separately configured scenario shock and surge levers.

Feedback and adaptation. dMFA sends three aggregated signals to SD each iteration: service stress (unmet service relative to demand), circular supply stress (dependence on primary rather than secondary supply), and strategic stock coverage (reserve stock relative to demand, where the reserve is active). SD responds by adjusting scarcity pressure and real demand, the collection multiplier, and — where enabled — strategic fill and release intent. This is a deliberately narrow interface: coupling exchanges a small set of interpretable, auditable variables rather than internal state, which keeps the joint solution inspectable at the cost of not representing each physically plausible SD–dMFA interaction (see Assumptions and Limitations).

Dissipation and loss. Pre-use losses (semi-fabrication and finished-product-manufacturing losses net of new-scrap recovery), uncollected end-of-life material, sorting rejects, and residue-treatment losses all route to a single terminal sink, `sysenv` / `system_losses`. The model does not currently distinguish dissipative from non-dissipative losses, nor does it track in-use or post-use dissipation separately from routed disposal — this is a stated limitation, not an oversight, and is carried through to the indicator design (see Data and Assumptions and Limitations sections).

Spatial dimension. The implementation is territorial in principle — the OD-trade layer operates on named regions — but is not spatially explicit in the active configuration, which resolves to a single global compartment. Within-region heterogeneity (mining geography, refining location, national policy) is therefore not represented; this is a scope choice enabled by the region dimension's current single value, not a structural limitation of the region dimension itself.

Uncertainty and stochasticity. A given run configuration produces a single trajectory, and the operational runtime does not propagate parameter uncertainty via a Monte Carlo ensemble — there is no automated multi-draw sampling of parameters into confidence intervals on reported outputs. This does not mean the active baseline is noise-free: the SD layer's demand/collection noise process runs within each baseline execution from each material's historical-dynamics start year, calibrated per material against real volatility drivers, and its seed can be varied across runs for reproducible ensemble comparisons if needed. Beyond this built-in noise process, uncertainty is otherwise handled through configuration — scenario variants, sensitivity/diagnostic run presets, and (for two materials) explicit calibration against observed series — rather than through formal probabilistic propagation.

Key simplifications. Four simplifications recur across the model, ahead of their detailed treatment in Model Architecture and Assumptions and Limitations: productive capacity is represented as a single aggregate envelope index rather than explicit plant-level capital stock; collection is represented as a single aggregated rate rather than resolved by infrastructure type or geography; the refinery stockpile is a control buffer between recovery and refining demand rather than a market-clearing

inventory mechanism; and the strategic inventory is aggregate — tracked per material rather than resolved by material stream — and, while enabled at the run level in the active baseline, is calibrated to a zero target-coverage level (effectively neutralised) for copper and silicon and left materially active, with material-specific fill/release thresholds, for nickel, lithium and cobalt.

Details

Initial conditions. Because the model's authored data begins at 1950 without a documented pre-1950 flow history for all materials, each material is given a synthetic burn-in: a spin-up period of material-specific length that builds a physically plausible starting in-use stock before the reporting-relevant part of the run begins. Burn-in length and design are set per material against literature or benchmark anchors.[2] The Calibration section provides full burn-in parameters and the corresponding SD warm-start design.

Model input data. Ten exogenous input files are required to execute the active baseline (Table 6); five additional observed-data files are optional and used only for calibration evaluation, never as state drivers. Full schema — dimensions, units, population methodology per file — is given in the Data section.

Main submodel logic

dMFA. Real demand from SD is allocated across the product/stream mix; cohort convolution generates outflow from use; end-of-life flows are collected, sorted, remanufactured, recycled or disposed of according to routing rates and yields; new scrap from fabrication and manufacturing is split between secondary recovery and residue; and primary availability at refining is derived from the ore-input anchor through the deposit/route/yield allocation chain, plus trade terms where the OD layer is active. The full equation set (Eq-MFA-01 through Eq-MFA-37, plus the endogenous circular-loop equations Eq-END-00 through Eq-END-07) is given in Model Architecture.

SD. Scarcity and bottleneck pressure raise price; price reduces demand once the reporting-start gate is reached; utilisation drives a lagged capacity envelope; price and bottleneck pressure drive collection effort, which sets the collection multiplier passed to dMFA; and, where enabled, coverage gaps and stress thresholds trigger strategic reserve fill or release intent. The full mechanism set (Eq-SD-01 through Eq-SD-24) is given in Model Architecture.

Coupling. Each iteration runs SD, then dMFA, then recomputes the three feedback signals; signals are smoothed and, by default, updated only in reporting years; convergence is reached when the largest of the signal, collection-multiplier and strategic-intent deltas falls below a configured tolerance. The full equation set is given in Model Architecture; an origin–destination trade layer exists as a further, currently inactive capability and is not developed in mechanistic detail there.

Model outputs. The model reports in-use stock and use-phase inflows/outflows, primary and secondary supply, end-of-life flows split by destination, new and old scrap, refinery-stockpile and strategic-reserve stocks, delivered and unmet service, coupling and SD diagnostics, circularity

[2] Nickel's burn-in is materially longer than the other four materials' because its documented industrial use, and therefore its pre-1950 in-use stock accumulation, extends further back than the model's 1950 data start; a shorter spin-up calibrated to the same starting conditions as the other materials would understate the stock nickel had already accumulated by 1950. Copper's burn-in is longer still, for the same reason and to a greater degree, reflecting an even longer documented history of industrial use.

indicators (End-of-Life Recycling Rate (EOLRR, Recycling Input Rate (RIR)), resilience indicators, and mass-balance residuals. The full runtime and export variable catalogue — over 100 named series — is given in the Data section.

Software basis. The model is not written from first principles in either half. The material-flow layer is built on `flodym`, an open-source Python library for dynamic material flow analysis,[3] and the behavioural layer on `BPTK-Py`, an open-source Python framework for System Dynamics and agent-based simulation.[4] The choice is consequential for interpretation because each library supplies conventions inherited by the model. .

From `flodym`, the model takes its dimensional machinery and its stock dynamics. Flows, stocks and parameters are declared over named dimension sets and combined by operations that align dimensions automatically, which is why the dimension letters in Table 1 are a property of the implementation and not merely descriptive notation: a flow declared over `(t, m, r, s)` cannot silently be multiplied by a parameter declared over the wrong axes. The cohort convolution in Eq-MFA-08 and the lifetime distribution families described in the Data section are that library's dynamic stock models, parameterised per stream rather than reimplemented here. `flodym` is itself a reimplementation of the concepts of ODYM[5], the framework that established the dimension-set and mass-balance conventions used across much of the dynamic material flow analysis literature, so the accounting conventions in Tables 8 to 11 are the field's standard ones rather than bespoke to this model.

From `BPTK-Py`, the model takes its System Dynamics solver: the stocks, flows and converters of the behavioural layer are declared in that framework's model interface, and the integration of the state equations — the capacity envelope in Eq-SD-12 and the collection-effort state in Eq-SD-16 — is performed by it rather than by hand-rolled difference equations. What the framework does *not* supply is the coupling: the loose iterative reconciliation between the two layers, the feedback signals, and the convergence rule in Table 18 are specific to this model.

Both libraries are pinned in the model's dependency specification and resolved to fixed versions in its lock file, so a given result set corresponds to a specific pair of library versions. The configuration described in this document resolves `flodym` 0.8.1 and `BPTK-Py` 2.3.0. Neither library is vendored into the repository, so reproducing a run requires resolving the lock file rather than the version ranges declared alongside it.

Evaluation. The model is evaluated on historical plausibility rather than on statistical fit. In-use stock is the primary target: each material's pre-1950 initialisation is set against an independent stock estimate at a benchmark year, and the behavioural layer's activation is tested by reproducing documented historical supply and price episodes. Copper, nickel and silicon meet this standard; cobalt and lithium carry provisional stock references pending independently sourced observations (flagged in Assumptions and Limitations). Build-time reconciliation gates additionally abort the build if a stream-share-weighted effective rate drifts outside its tolerance band, and all run checks

[3] Dürrwächter, J., Hosak, M., Weiss, B., & Ueckerdt, F. (2026). flodym: A Python package for dynamic material flow analysis. *Journal of Open Source Software, 11*(119), 10105. doi:10.21105/joss.10105

[4] Transentis labs GmbH. (2026). *BPTK-Py: Business Prototyping Toolkit for Python* (Version 2.3.0) [Computer software]. GitHub:transentis/bptk_py

[5] Pauliuk, S., & Heeren, N. (2020). ODYM—An open software framework for studying dynamic material systems: Principles, implementation, and data structures. *Journal of Industrial Ecology, 24*(3), 446–458. doi:10.1111/jiec.12952

convergence diagnostics and mass-balance residuals. An automated parameter-fitting facility is available and runnable, but is not exercised in the active configuration. The Calibration section provides the full calibration methodology, targets, and governance.

Model architecture

Develops in full the mechanisms and equations introduced at overview level in Goal and Scope. The model distinguishes general capability from active configuration throughout; inactive capabilities are noted briefly rather than developed in mechanistic detail.

The model is architected as three cooperating layers: a dynamic material flow analysis (dMFA) layer that enforces physical mass conservation and cohort-based stock ageing; a system dynamics (SD) layer that represents adaptive behavioural response in demand, capacity, collection effort and strategic reserve policy; and a coupling layer that exchanges a small, auditable set of feedback signals between them. A further origin–destination trade layer exists as an implemented but inactive capability, outlined near the end of this section rather than developed in detail. Each active layer is developed below in the order dMFA → SD → coupling, matching the order in which a solved year's information actually flows: SD proposes demand and policy intents, dMFA resolves them against physical constraints, and coupling reconciles the two under iteration.

Dynamic MFA formulation

Dimensional structure. The dMFA layer resolves flows and stocks over five core dimensions: `t` (time), `m` (material), `r` (region), `z` (product), and `s` (material stream). Two further axes, `d` (deposit type) and `p` (processing route), allocate primary ore to the chain before the native flodym solve begins; they are source/route allocation axes rather than dimensions the physical flow/stock system itself carries. In the active branch scope, `r` resolves to a single value (`Global`), `z` resolves to a single placeholder value (`canonical`), and `m` spans the five active materials; `s` is the only detail axis carrying real, named structure (20 material streams; see Goal and Scope, Table 2), and it is not carried on each flow — it activates at the refining stage and is native on all flows and stocks from refining downstream, while `mining_beneficiation` and `smelting` carry only `[t, r, z]`.

Active contract and topology. The executable MFA behaviour is selected through a named *contract*, which in turn fixes the topology, parameter schema, stream structure, source/route structure, and accounting convention. The active contract's defining properties are a supply-anchored primary path (primary availability is derived from an authored ore-input anchor rather than unconditionally following demand), a disabled product/component detail structure, and a product-only remanufacturing policy. The contract resolves to five named components, each elaborated at the point in this document where its content actually matters rather than here: the *topology* (`cgr_crm_native_stream_v1`) is the twelve-stage process chain set out immediately below; the *stream structure* (`cgr_material_streams_refining_v1`) is the twenty-stream classification given in Goal and Scope, Table 2; the *source/route structure* (`cgr_source_route_v1`) is the deposit/route allocation chain described under Primary supply path, below; the *accounting convention* (`cgr_crm_v1`) is the concrete set of stock-flow identities given in the equation tables through this section; and the *parameter schema* (`cgr_refining_sorting_product_params_v1`) is the executable allowlist of calibratable parameters, listed in full in the Data section.

Process topology. The active executable topology comprises twelve stages: `mining_beneficiation` → `smelting` → `refining` → `semis_fabrication` → `finished_product_manufacturing` → `use`

→ `collection` → `sorting_preprocessing` → { `remanufacturing` / `recycling_refining_secondary` / `residue_treatment_disposal` } → `sysenv` (the model boundary/terminal sink). This is the same topology introduced at the overview level in Goal and Scope (Table 3); it is repeated here because each equation below is stated in terms of these named stages.

Demand-to-use pipeline. For each modelled year, the pipeline runs: (1) SD provides real demand; (2) source/route allocation derives the upstream material-stream mix through `route_stream_share.csv`; (3) aggregate demand is allocated to streams using that upstream mix; (4) stream demand is split by product through `product_share.csv` (currently a single placeholder product, so this step is a formal pass-through rather than a substantive allocation); (5) remanufactured inflow and new-production inflow jointly feed the use phase; (6) cohort ageing generates outflow from use; (7) use-phase outflow becomes old-scrap generation, entering the collection/routing chain described below.

Service-demand boundary convention. A specific modelling convention governs how the exogenous demand input is interpreted, with direct consequences for how several reported indicators should be read. The exogenous `final_demand.csv` series — *service_demand* — is defined as *refined apparent consumption*: refined metal entering fabrication, not an into-use inflow target. This matches the documented provenance of each material's demand series (world refined usage for copper; apparent consumption for nickel, cobalt and lithium; polysilicon/ferrosilicon/silicon-metal demand for silicon), all measured at the refined level that feeds fabrication rather than at the into-use level. The two boundaries differ by the fabrication yield (for example ≈0.80 for copper, roughly a 25% gross-up): refined input required equals service demand directly, while into-use inflow equals service demand multiplied by the fabrication yield, and primary need equals service demand minus secondary and strategic contributions (capped). Service-level and unmet-service indicators are defined consistently at the refined boundary: the supply-side numerator is total fabrication throughput (primary plus secondary plus strategic-reserve releases) against refined service demand in the denominator.

This boundary choice has three practical implications for interpreting model output. First, circularity ratios such as the End-of-Life Recycling Rate and Recycling Input Rate are largely insensitive to it, since numerator and denominator scale together with the fabrication yield. Absolute flow figures are not: an into-use series always reads systematically lower than the corresponding refined-input series for the same material movement, purely because of the fabrication-yield gap, so the two should not be compared without converting boundaries. Second, and most concretely, it shapes how remanufactured material is credited. Remanufactured flow re-enters use directly from collected end-of-life material without passing back through fabrication, so it does not incur a second fabrication-yield loss the way new primary or secondary production does; under the refined-boundary convention, this gives remanufacturing a distinct, yield-free credit against primary-plus-secondary refined-input requirements that recycling (which re-enters upstream at refining and is subject to the full downstream yield chain again) does not receive. Any indicator that compares remanufacturing's contribution against recycling's on a like-for-like basis needs to account for this asymmetry rather than reading the two flows as interchangeable tonnes. Third, resilience indicators (delivered and unmet service) are only meaningful against series measured at the same boundary: combining a refined-boundary service figure with an external into-use statistic will misstate the deficit or surplus unless the fabrication-yield conversion is applied first.

Anchoring demand at the refinery rather than the use stage also shapes how the model relates demand to in-use stock — a choice each dynamic MFA model must make. An *inflow-driven* formulation fixes new-product inflow directly from an exogenous demand series and derives the in-use stock as the accumulated, aged result of that inflow history; a *stock-driven* formulation does the reverse — it fixes a target in-use stock trajectory and solves for whatever new-product inflow is needed each year to reach it, net of retirements. The model's core solve is inflow-driven: `service_demand` is applied directly as this year's new-product inflow requirement, so the in-use stock is an emergent outcome of accumulating that inflow under the lifetime distribution, not an independent target the solve steers toward. This has a counterintuitive consequence for lifetime-extension strategies: because inflow already equals demand by construction, a longer-lived unit does not by itself reduce how many new units are produced — it only delays retirement of those already produced, so in-use stock grows instead of the replacement requirement shrinking, the opposite of the avoided-production effect such a strategy is meant to have. To correct this without abandoning the inflow-driven core, the model carries a second, opt-in stock-driven submodule used for one specific purpose: it derives a reference in-use-stock target from `service_demand` under a reference (un-extended) lifetime, then re-solves new-product inflow so that the actual, potentially lifetime-extended stock still reaches that same target — fewer retirements translate into correctly-signed avoided new production rather than a growing stock. This submodule (`slow_replacement_credit`) is switched on by default in the active contract, but it is a no-op wherever the active lifetime already equals the reference lifetime; in the baseline configuration this document describes, per-stream lifetime multipliers sit at essentially 1.0 throughout, so the submodule is structurally live but currently produces no numerical effect — it only activates once a scenario applies an explicit lifetime-extension lever (see Scenarios).

A related but separate challenge concerns what resilience indicators computed at the refinery boundary can and cannot see. `Service_level` and the resilience scalars built on it measure supply-demand balance strictly at the point material enters fabrication — `Primary_supply` plus `Secondary_supply` plus `Strategic_inventory_outflow`, against refined service demand. By construction, these indicators are blind to any flow that returns to use without passing back through fabrication, for the same reason such flows are excluded from EoL-RR and RIR (see above). Two flows fall into this blind spot. The first is the remanufacturing flow described below under Circular routing—active in the baseline configuration—which returns collected end-of-life material for direct use from the sorting/routing stage. The second is a separate product-recovery mechanism that diverts units to reuse or refurbishment/remanufacturing before they reach end-of-life collection, straight out of first-life retirements; this mechanism exists as an opt-in lever, activating only in scenarios that raise circularity ambition above the current-practice rung. The model provides a second, use-boundary indicator family — `Functional_service_level` and matching resilience scalars — specifically to credit bypass flows that `Service_level` cannot see: it adds a credit for the diverted reuse/refurbishment/remanufacturing flow, discounted by a per-material second-life ratio reflecting that a reused or remanufactured unit typically delivers somewhat less service than a new one, to the same numerator `Service_level` uses, against the same service-demand denominator — so it always reads at least as favourably as `Service_level` and moves consistently with it rather than diverging or saturating independently. One practical consequence follows directly: because the credited pathway is switched off in the baseline configuration, `Functional_service_level` is

numerically identical to Service_level in each reported year here; the two indicators diverge only in scenarios where the product-recovery lever is active.

Circular routing. End-of-life material is routed through collection → sorting_preprocessing → { remanufacturing / recycling_refining_secondary } → system_losses under a two-gate schema. Gate 1 is collection: collection_rate ∈ [0, 1] is the share of end-of-life outflow collected into any recovery system; the uncollected complement, (1 − collection_rate), flows directly to system_losses. Gate 2 applies conditionally on collection and splits collected material three ways — recycling_rate, remanufacturing_rate, and downcycling_rate — constrained to sum to exactly 1. Loss routing through the terminal sink is described below. Remanufacturing remains inside this existing collection-routing loop and returns material to use directly (not through a separate reuse/refurbishment channel); no dedicated reuse or refurbishment routing rates or processes are currently active, and any reuse/refurbishment strategy effect is represented instead as a change to product_lifetime_multiplier (see Scenarios).

Product lifetimes. Product lifetimes are drawn from lifetime_distributions.csv, scaled by a per-stream product_lifetime_multiplier. The active product axis z currently ships one temporary placeholder value, canonical, so lifetime differentiation in the active baseline operates entirely through the material-stream axis rather than through product-level differentiation. Stream-enabled runs keep downstream use-phase quantities native over the material-stream × product combination, so the product axis remains structurally present and will carry real differentiation without an architecture change once richer product data is authored.

Primary supply path. The active authored primary anchor is primary_ore_input, aggregated over (year, material, region) and supply-anchored: final demand does not impose a minimum primary-supply floor on it. Runtime derives primary_concentrate_input, primary_refined_input, domestic_primary_available_to_refining[t,r,s] and primary_available_to_refining[t,r,s] through a source/route allocation chain: ore input is allocated across deposit types and diluted by ore grade (deposit_share, ore_grade), mapped to processing routes (route_share), and carried through mining-beneficiation, smelting and refining yields (route_yield). route_stream_share.csv then maps each processing route's output to the twenty named material streams at the refining split — this is the point at which the stream axis is activated. In the current shipped configuration, deposit_type and processing_route both collapse to a single canonical placeholder value with deposit_share = 1 and ore_grade = 1 — a stated temporary simplification, not a claim that the source/route allocation layer itself is inactive. The stream mapping itself is not a placeholder: route_stream_share carries differentiated, benchmark-anchored shares per material, stream and year, developed in the Data section. The final refining-availability term combines the domestic allocation with net trade where OD-trade is enabled.

Secondary supply path. Secondary material reaching the refinery stockpile comes from three sources: recycled old scrap (collected end-of-life material routed through Gate 2's recycling share), routed new scrap (semis/manufacturing losses net of the new-scrap recovery rate), and — only when OD-trade is enabled — net scrap trade imports. The refinery stockpile (refinery_stockpile_native) is the operational buffer between this variable secondary-material recovery and refining-stage input requirements; its release into refining demand is governed by

refinery_stockpile_release_rate, and its stock evolves as last period's stock plus inflow minus outflow, floored at zero.

Stocks and reserve behaviour. Three stock variables carry the dMFA layer's physical state forward year to year. **In-use stock** (stock_in_use) is driven by product inflows (new production plus remanufactured return) and lifetime-distribution-governed cohort retirements. The **refinery stockpile** (refinery_stockpile_native) buffers secondary-material recovery against refining-stage demand, as above. The **strategic inventory** (strategic_inventory_native) is a reserve stock that fills and releases according to SD-generated intents (see System Dynamics formulation, Mechanism 4); it is aggregate: tracked per material rather than resolved by material stream. The reserve is enabled at the run level and materially active for three of the five modelled materials, with the following per-material calibration: copper and silicon are neutralised via a zero target-coverage setting; nickel carries a 0.1-year coverage target with a 1.25 price-ratio release threshold; lithium carries a 1.15 fill-price threshold, 0.15 fill-service threshold and 0.20 release-service threshold; and cobalt carries a 1.30 fill-price threshold, 1.40 release-price threshold and 0.30 release-service threshold.[6]

Loss routing. Semis-fabrication and finished-product-manufacturing losses generate new scrap, a fraction of which (the new-scrap recovery rate) is recovered to the refinery stockpile; the remainder routes to residue/system losses. Pre-use losses in mining, beneficiation and smelting terminate directly at sysenv. Sorting rejects and recycling/remanufacturing-treatment losses route through system_losses alongside uncollected end-of-life material. The model does not currently distinguish dissipative from non-dissipative losses within this single terminal sink, a stated limitation carried through to indicator design (see Assumptions and Limitations).

Endogenous circular-loop equations. A staged extension governs how far the collection-routing gates and remanufacturing rate are allowed to move endogenously away from their exogenous anchors in response to price, quality and collected-volume signals, rather than following the anchor trajectory unconditionally. The governing principle across all stages is *baseline-preserving response*: each endogenised rate is expressed as a deviation from the exogenous CSV anchor, and each deviation term is constructed to equal zero under baseline conditions, so that a run with all response gains set to zero reproduces the exogenous anchor trajectory exactly — endogenous response only ever shows up as a movement away from that anchor, never as a change to it. This is implemented in 3 stages: recycling/downcycling split, remanufacturing via core-material availability and quality/downcycling cascade. Each stage is independently togglable and collapses to the exogenous anchor when its gains are zero.

Stage 1's recycling/downcycling split is driven by a logit utility with a structural intercept (the anchor's own recycling-versus-downcycling ratio) plus three deviation terms — a price term (elevated recycling responsiveness for copper, more muted for the smaller/thinner markets), a quality term (zero unless Stage 3 is also active), and a collected-volume scale term — each zero at baseline. Stage 2's remanufacturing rate multiplies the exogenous anchor by a response function of core-material suitability (the age-weighted quality of retiring cohorts available as remanufacturing

[6] Copper and silicon are neutralised because both are deep, liquid markets in which a sovereign-scale reserve is judged not to operate meaningfully. Nickel's threshold is calibrated against the roughly 250% nickel price spike of March 2022. Lithium's thresholds reflect that EV-transition demand tightness means a more conservative default fill gate is rarely met. Cobalt's thresholds reflect DRC supply concentration and observed large-scale strategic purchasing.

cores) and price ratio; the function is *bounded* — capped within configured minimum/maximum multipliers so a shock cannot push remanufacturing to implausible extremes — and *mean-preserving*, meaning that when suitability and price sit at their typical run-average levels the multiplier equals one and remanufacturing simply follows its exogenous anchor; it only pulls away from that anchor as conditions move away from typical, calibrated at the "low" responsiveness rung (`gamma_core = 0.25`) consistent with the canonical current-practice circularity scenario. Stage 3 tracks a per-material-stream quality index that mixes primary input, held at a fixed reference quality, with secondary (recycled) input, whose quality is reduced each time it passes through the recycling loop by a fixed *per-cycle quality-loss term* — the same idea as paper or metal losing a little integrity each time it is reprocessed — calibrated for copper, nickel and cobalt; lithium and silicon are held quality-inert because their secondary streams have negligible downcycling room. That quality index can only ever fall relative to the primary reference, never rise (a *monotone-decreasing* deviation), and this decline feeds back into the Stage 1 recycling utility as a one-directional pressure: degraded quality only ever nudges more material toward the lower-value downcycling path, mirroring how real recycled-metal streams accumulate contaminants ("tramp elements") over repeated cycles rather than spontaneously purifying.

Required exogenous inputs and core diagnostics. Ten exogenous input files are required to execute the active baseline configuration — `final_demand`, `route_stream_share`, `product_share`, `primary_ore_input`, `deposit_share`, `ore_grade`, `route_share`, `route_yield`, `collection_routing_rates` and `lifetime_distributions` — full schema, units and population methodology for each are given in the Data section. A further set of observation files exists for potential calibration use but plays no role in the active baseline and is not developed in the Data section (see Data, Calibration data). Core dMFA diagnostics tracked at runtime include `Stock_in_use`, `Primary_ore_input`, `Domestic_primary_available_to_refining`, `Primary_available_to_refining`, `Primary_supply`, `Secondary_supply`, `Refinery_stockpile_stock`, `Strategic_inventory_stock`, `EoL_generated`, `EoL_recycled`, `EoL_remanufactured`, `EoL_downcycled`, `EoL_system_losses` and `Mass_balance_residual_max_abs`, the last of which is checked against a configured tolerance on all runs regardless of which other diagnostics are active. Tables 7 to 11 state the physical layer in full — notation, the primary chain, the use phase and its stock, end-of-life routing, and the supply balance with the indicators built on it — and Table 12 the endogenous circular-loop equations that extend them.

The equations below are stated in a compact notation rather than in the model's internal variable names, so that dimensional structure is visible. Table 7 maps each symbol to the quantity it denotes and to the identifier the model uses for it, and fixes the index conventions. Three conventions apply throughout. First, subscripts are dimension indices in the order `t, m, r, z, s` — time, material, region, product, material stream — and an index is written only where the quantity actually varies along it; a quantity carrying no stream subscript is resolved at material level and broadcast across streams. Second, superscripts in roman type are labels, not exponents: they name a process, a route or a gate. Third, all flows are annual masses in kilotonnes of contained metal, and all stocks are masses at year-end, so a flow and a stock never appear on the same side of an identity without a time difference.

Table 7. *Notation used in the equation tables: symbols, the quantities they denote, the dimensions each carries, and the corresponding model identifier.*

Symbol	Quantity	Dimensions	Model identifier
$t,\ m,\ r,\ z,\ s$	Year, material, region, product, material stream	—	`Dimension letters as given in Table 1`
$d,\ p$	Deposit type and processing route: source-allocation axes collapsed before the flow solve	—	`deposit_type, processing_route`
$a \in \{0,\ \ldots,\ A\}$	Cohort age in years, where A is the maximum tracked age	—	`lifetime distribution age axis`
$Q_{t,m,r}$	Primary ore-feed-equivalent input	$t,\ m,\ r$	`primary_ore_input`
$\delta_{t,m,r,d}$, $\gamma_{t,m,r,d}$	Deposit share and ore grade	$t,\ m,\ r,\ d$	`deposit_share, ore_grade`
$\psi_{t,m,r,d,p}$	Route share within a deposit type	$t,\ m,\ r,\ d,\ p$	`route_share`
$\omega_{t,m,r,p,s}$	Route-to-stream share	$t,\ m,\ r,\ p,\ s$	`route_stream_share`
η^{mb}, η^{sm}, η^{ref}	Mining-and-beneficiation, smelting and refining yields	$t,\ m,\ r,\ p$	`route_yield.mining_beneficiation_yield / smelting_yield / refining_yield`
η^{fab}, η^{srt}, η^{rec}, η^{rem}	Fabrication, sorting, recycling and remanufacturing yields	$t,\ m,\ r,\ s$	`fabrication_yield, sorting_yield, recycling_yield, reman_yield`
$\rho^{col}_{t,m,r,s}$	Gate 1 collection rate	$t,\ m,\ r,\ s$	`collection_rate`
ρ^{rec}, ρ^{rem}, ρ^{dow}	Gate 2 recycling, remanufacturing and downcycling rates	$t,\ m,\ r,\ s$	`recycling_rate, remanufacturing_rate, downcycling_rate`
$\sigma_{t,m,r,s}$	New-scrap recovery rate	$t,\ m,\ r,\ s$	`new_scrap_recovery_rate`
$\varphi_{t,m,r,z}$	Product share of service demand	$t,\ m,\ r,\ z$	`product_share`
$I,\ O$	Inflow to and outflow from the use phase	$t,\ m,\ r,\ z,\ s$	`inflow_to_use_total, outflow_from_use`
$f_{c,m,r,s}(a)$	Retirement probability of cohort c at age a; the cohort index makes the distribution cohort-varying	$c,\ m,\ r,\ s,\ a$	`lifetime_distributions`
S^{use}, S^{2nd}	First-life and second-life in-use stock	$t,\ m,\ r,\ z,\ s$	`__first_life_stock, __second_life_stock`

Symbol	Quantity	Dimensions	Model identifier
S^{stk}, S^{res}	Refinery stockpile and strategic inventory	t, m, r, s	`refinery_stockpile_native`, `strategic_inventory_native`
G, C, U	End-of-life mass generated, collected and uncollected	t, m, r, s	`old_scrap_generated`, `old_scrap_collected`, `old_scrap_uncollected`
N, N^{sec}	New (manufacturing) scrap generated and recovered to secondary	t, m, r, s	`new_scrap_generated`, `new_scrap_to_secondary`
P, X	Primary and secondary supply used at the refining boundary	t, m, r, s	`Primary_supply`, `Secondary_supply`
D^{des}, D	Desired and real demand	t, m, r	`SD_demand_desired`, `SD_demand`
D^{srv}	Service demand at the refined boundary	t, m, r, z	`__service_demand`
Λ, Λ^{fun}	Service level and its use-boundary complement	t, m, r	`Service_level`, `Functional_service_level`
π, π_0, θ	Price, reference price and price ratio	t, m, r	`SD_price`, `price_base`, `SD_price_ratio`
κ, ν, β	Capacity envelope, flow utilisation and bottleneck pressure	t, m, r	`SD_capacity_envelope`, `SD_flow_utilization`, `SD_bottleneck_pressure`
ε, μ^{col}	Collection effort state and the collection multiplier it produces	t, m, r	`SD_collection_effort`, `Coupling_collection_multiplier`
Σ^{srv}, Σ^{cir}, Y^{res}	The three dMFA-to-SD feedback signals: service stress, circular-supply stress, reserve coverage	t, m, r	`service_stress_t`, `circular_supply_stress_t`, `strategic_stock_coverage_years_t`
ξ	A small positive constant guarding a denominator	—	`eps`

The primary chain. The upstream half of the system converts an authored ore-feed-equivalent input into refined metal resolved by stream. It runs over the deposit and route axes, which are summed out before the flow solve begins, so no downstream equation carries a d or p index. Table 8 states it. The chain is multiplicative rather than additive: each stage scales the previous stage's output by a yield, so a proportional error in any one yield propagates unattenuated to refined output.

Table 8. *Primary supply chain: source and route allocation, from ore-feed equivalent to stream-resolved refined output.*

Eq.	Equation	Interpretation
Eq-MFA-01	$M_{t,m,r,d} = Q_{t,m,r} \cdot \delta_{t,m,r,d} \cdot \gamma_{t,m,r,d}$	Ore feed apportioned to a deposit type and converted to contained metal by that deposit's grade.
Eq-MFA-02	$M_{t,m,r,d,p} = M_{t,m,r,d} \cdot \psi_{t,m,r,d,p}$	Deposit output apportioned across processing routes.
Eq-MFA-03	$R_{t,m,r,d,p} = M_{t,m,r,d,p} \cdot \eta^{mb}_{t,m,r,p} \cdot \eta^{sm}_{t,m,r,p} \cdot \eta^{ref}_{t,m,r,p}$	The three-stage upstream yield chain: mining and beneficiation, smelting, refining.
Eq-MFA-04	$R_{t,m,r,s} = \sum_{d,p} R_{t,m,r,d,p} \cdot \omega_{t,m,r,p,s}$	Stream identity is created here and only here: refined output is split across streams, and the deposit and route axes are summed out.
Eq-MFA-05	$P_{t,m,r,s} = max(0,\ R_{t,m,r,s} + T^{ref}_{t,m,r,s} + \chi_{t,m,r} \cdot T^{con}_{t,m,r,s})$	Primary availability at refining, inclusive of net refined and concentrate trade. Both trade terms are identically zero in the active baseline.
Eq-MFA-06	$Q_{t,m,r} = \frac{P^{eq}_{t,m,r}}{\eta^{ref}_{t,m,r}}$	The inverse relation by which the ore-feed-equivalent input is authored: it is derived backwards from a primary metal equivalent, not observed.

The use phase and its stock. The use phase is where the model's dynamics live, and it is inflow-driven: each year's inflow is a cohort with its own retirement distribution, and the outflow in any year is the sum over all surviving cohorts of the mass reaching its retirement age that year. Table 9 states this. Two features distinguish it from a fixed-lifetime formulation. The retirement density carries a cohort index as well as an age index, so a change in product design shows up as a different distribution for cohorts made after the change rather than as an instantaneous change to the whole stock. And the convolution is truncated at the maximum tracked age and at the start of the run, which is why the pre-1950 initialisation described in the Calibration section is necessary: without it the first decades of the reporting window would retire nothing, because nothing had yet been put into use.

Table 9. *Use-phase cohort dynamics: retirement, in-use stock, and the second-life pathway.*

Eq.	Equation	Interpretation
Eq-MFA-07	$I_{t,m,r,z,s} = D^{srv}_{t,m,r,z} \cdot \varphi_{t,m,r,z} \cdot \eta^{fab}_{t,m,r,s}$	Inflow to use: service demand distributed over products and reduced by the fabrication yield chain.
Eq-MFA-08	$O^{1}_{t,m,r,z,s} = \sum_{a=1}^{min(t-t_0, A)} I_{t-a,m,r,z,s} \cdot f_{t-a,m,r,s}(a$	First-life retirement as a cohort convolution. The retirement density is indexed by cohort year as well as age, so lifetimes vary by cohort; the sum is truncated at the maximum tracked age A and at the first year of the run.
Eq-MFA-09	$S^{use}_{t,m,r,z,s} = S^{use}_{t-1,m,r,z,s} + I_{t,m,r,z,s} - O^{1}_{t,m,r,z,}$	First-life in-use stock balance. Equivalently, the cumulative sum of net inflow from the first year of the run.
Eq-MFA-10	$I^{2}_{t,m,r,z,s} = \lambda_{t,m,r,s} \cdot O^{1}_{t,m,r,z,s}$	Second-life inflow: the share of first-life retirements diverted to reuse or refurbishment before end-of-life collection. Zero in the active baseline, where the product-recovery lever is off.
Eq-MFA-11	$O^{2}_{t,m,r,z,s} = \sum_{a=1}^{min(t-t_0, A)} I^{2}_{t-a,m,r,z,s} \cdot f^{2}_{t-a,m,r,s}(a$	Second-life retirement, convolved against the second-life retirement density.
Eq-MFA-12	$S^{2nd}_{t,m,r,z,s} = S^{2nd}_{t-1,m,r,z,s} + I^{2}_{t,m,r,z,s} - O^{2}_{t,m,r,z,}$	Second-life stock balance.
Eq-MFA-13	$S^{tot}_{t,m,r,z,s} = S^{use}_{t,m,r,z,s} + S^{2nd}_{t,m,r,z,s}$	Reported in-use stock is the sum of the two, and is the quantity compared against independent stock estimates.
Eq-MFA-14	$O_{t,m,r,z,s} = O^{1}_{t,m,r,z,s} + O^{2}_{t,m,r,z,s}$	Total retirement from the use phase.
Eq-MFA-15	$L_{t,m,r,s} = L^{ref}_{m,r,s} \cdot \Lambda^{slow}_{t,m,r,s}$	Lifetime extension is applied as a multiplier on the reference mean lifetime, which rescales the retirement density rather than introducing a separate mechanism.

End-of-life routing. Retired mass passes two gates in sequence, and Table 10 states them. The separation matters for interpretation as much as for accounting: the first gate is a collection question, answered by infrastructure and policy, and the second is a treatment question, answered by economics and material quality. Because they are separate, each can be calibrated against its own evidence, and each responds to a different driver in the endogenous layer described below.

Table 10. *End-of-life routing: the two gates, the recovery chain, and manufacturing scrap.*

Eq.	Equation	Interpretation
Eq-MFA-16	$G_{t,m,r,s} = \sum_{z} (O_{t,m,r,z,s} - I^{2}_{t,m,r,z,s})$	End-of-life mass entering the routing chain, net of mass diverted to a second life, aggregated over products.

Eq.	Equation	Interpretation
Eq-MFA-17	$C_{t,m,r,s} = G_{t,m,r,s} \cdot \rho^{col}_{t,m,r,s} \cdot \mu^{col}_{t,m,r}$	Gate 1. The exogenous collection rate is scaled by the collection multiplier returned by the behavioural layer, then bounded.
Eq-MFA-18	$U_{t,m,r,s} = G_{t,m,r,s} - C_{t,m,r,s}$	The uncollected complement, which leaves the system to losses. Landfill and dissipation are therefore not competing members of the Gate 2 simplex.
Eq-MFA-19	$\rho^{rec}_{t,m,r,s} + \rho^{rem}_{t,m,r,s} + \rho^{dow}_{t,m,r,s} = 1$	Gate 2 closure identity, conditional on collection. Enforced exactly at each solve.
Eq-MFA-20	$C^{k}_{t,m,r,s} = C_{t,m,r,s} \cdot \rho^{k}_{t,m,r,s}, \quad k \in \{rec, rem, dow\}$	Gate 2 applied: collected mass split three ways.
Eq-MFA-21	$A^{srt}_{t,m,r,s} = C^{rec}_{t,m,r,s} \cdot \eta^{srt}_{t,m,r,s}$	Sorting and preprocessing pass; the complement is rejected to residue treatment.
Eq-MFA-22	$X^{eol}_{t,m,r,s} = A^{srt}_{t,m,r,s} \cdot \eta^{rec}_{t,m,r,s}$	Secondary metal recovered from end-of-life material. The recycling yield here is a metallurgical process yield only: cross-circuit loss is accounted separately through the downcycling pathway.
Eq-MFA-23	$I^{rem}_{t,m,r,s} = C^{rem}_{t,m,r,s} \cdot \eta^{rem}_{t,m,r,s}$	Remanufactured mass returned directly to use, without passing back through fabrication.
Eq-MFA-24	$N_{t,m,r,s} = \frac{I_{t,m,r,s}}{\eta^{fab}_{t,m,r,s}} \cdot (1 - \eta^{fab}_{t,m,r,s})$	Manufacturing scrap generated upstream of use, implied by the fabrication yield.
Eq-MFA-25	$N^{sec}_{t,m,r,s} = N_{t,m,r,s} \cdot \sigma_{t,m,r,s}$	New scrap recovered to the refinery stockpile. The remainder leaves for residue treatment.

Eq.	Equation	Interpretation
Eq-MFA-26	$W_{t,m,r,s} = U_{t,m,r,s} + C^{dow}_{t,m,r,s} + (C^{rec}_{t,m,r,s} - X^{eol}_{t,m,r,s}) + (N_{t,m,r,s} - N^{sec}_{t,m,r,}$	Total mass leaving the system: uncollected, downcycled, lost in sorting and recycling, and unrecovered manufacturing scrap.

Supply balance, stockpiles and indicators. The last block of the physical layer reconciles what is available against what is demanded, buffers the difference through two stocks, and defines the quantities reported as results. Table 11 states it. The mass-balance check in Eq-MFA-32 is not a diagnostic that can be switched off: it is evaluated on all runs and raises rather than warns, so a configuration that does not close cannot produce output at all.

Table 11. *Supply balance, buffer stocks, the mass-balance condition, and the reported indicators.*

Eq.	Equation	Interpretation
Eq-MFA-27	$X_{t,m,r,s} = X^{eol}_{t,m,r,s} + N^{sec}_{t,m,r,s}$	Secondary supply comprises both end-of-life recycling and recovered manufacturing scrap. The inclusion of the second term is what makes this quantity differ from a published end-of-life recycling figure.
Eq-MFA-28	$S^{stk}_{t,m,r,s} = max(0, S^{stk}_{t-1,m,r,s} + X_{t,m,r,s} -$	Refinery stockpile: secondary material in, secondary material drawn to fabrication and diverted to the strategic reserve out, floored at zero.
Eq-MFA-29	$S^{res}_{t,m,r,s} = max(0, S^{res}_{t-1,m,r,s} + F^{res}_{t,m,r,s} -$	Strategic inventory, with fill and release flows set by the reserve intent policy. Inactive for copper and silicon in the active configuration.
Eq-MFA-30	$\Phi_{t,m,r} = \sum_{s} (P^{use}_{t,m,r,s} + X^{use}_{t,m,r,s} + E^{res}_{t,m,r}$	Fabrication throughput at the refined boundary: the quantity service level is measured on.
Eq-MFA-31	$Y^{k}_{t,m,r,s} = \frac{Y^{k+1}_{t,m,r,s}}{\eta^{k}_{t,m,r,s}}$	The backward yield chain: input required at any stage is the next stage's requirement divided by that stage's yield, applied from refining back to mining.
Eq-MFA-32	$max_{t} \lvert \varepsilon_{t} \rvert \leq \tau \cdot \Xi$	Mass-balance condition. The maximum absolute residual across all fifteen process balances must not exceed the tolerance times of a flow scale; a breach raises and halts the run.
Eq-MFA-33	$\Lambda_{t,m,r} = clip(\frac{\Phi_{t,m,r}}{max(D^{srv}_{t,m,r}, \xi)}, 0, 1)$	Service level: delivered against demand at the refined boundary, clipped to the unit interval.
Eq-MFA-34	$\Lambda^{fun}_{t,m,r} = clip(\frac{\Phi_{t,m,r} + \zeta_{m} \cdot \sum_{z,s} I^{2}_{t,m,r,z,s}}{max(D^{srv}_{t,m,r}, \xi)}, 0, 1)$	Use-boundary complement, adding a credit for the second-life bypass flow discounted by a per-material second-life ratio.

Eq.	Equation	Interpretation
		Numerically identical to Eq-MFA-33 whenever the product-recovery lever is off.
Eq-MFA-35	$\Delta_{t,m,r} = max(0, \Lambda^{*} - \Lambda_{t,m,r})$	Service deficit against a configured adequacy threshold; the basis of the resilience scalars.
Eq-MFA-36	$EoL - RR_{t,m,r} = \frac{\sum_{s} X^{eol}_{t,m,r,s}}{\sum_{s} G_{t,m,r,s}}$	End-of-life recycling rate: recovered over generated, measured post-yield. Flows returning to use without passing back through refining are invisible to it by construction.
Eq-MFA-37	$RIR_{t,m,r} = \frac{\sum_{s} X^{use}_{t,m,r,s}}{\sum_{s} (P^{use}_{t,m,r,s} + X^{use}_{t,m,r,s})}$	Recycling input rate: secondary share of refined input. Because the numerator includes recovered manufacturing scrap, this reads higher than an end-of-life-only measure.

Table 12. *Endogenous circular-loop equations (Eq-END-00 through Eq-END-07).*

Eq.	Equation	Interpretation
Eq-END-00	$\rho^{k,eff}_{t,m,r,s} = \rho^{k,anc}_{t,m,r,s} \cdot g_k(\frac{x_{t,m,r,s}}{x^{bas}_{t,m,r,s}}), \quad g_k(1) \equiv 1$	Baseline-anchored response identity, common to all three stages: an effective rate is its exogenous anchor times a response function of a driver measured relative to its own baseline. Because each response function is the identity at a driver ratio of one, a baseline run reproduces the anchor exactly, and the endogenous layer adds nothing.
Eq-END-01	$V_{t,m,r,s} = ln \frac{\rho^{rec,anc}_{t,m,r,s}}{\rho^{dow,anc}_{t,m,r,s}} + \alpha^{\pi}_{m}(\theta_{t,m,r} - 1) +$	Stage 1 utility governing the recycling-versus-downcycling split: a structural intercept set by the anchor's own ratio, plus price, quality and collected-volume deviation terms, each of which is zero at baseline.
Eq-END-02	$\rho^{rec,*}_{t,m,r,s} = \frac{1}{1+e^{-V_{t,m,r,s}}}, \quad \rho^{dow,*}_{t,m,r,s} = 1 - \rho^{rec,*}_{t,m,r,s}$	Logistic split of that utility. The two shares sum to one exactly by construction rather than by normalisation after the fact.
Eq-END-03	$\rho^{rec,eff}_{t,m,r,s} = (1 - \rho^{rem,eff}_{t,m,r,s}) \cdot \rho^{rec,*}_{t,m,r,s}$	Stage 1 applied to the Gate 2 simplex: remanufacturing is skimmed first, and the remainder is split by Eq-END-02, so Gate 2 closure holds exactly whatever the two stages do.

Eq.	Equation	Interpretation
Eq-END-04	$\Psi_{t,m,r,s} = \frac{\sum_a w(a) \cdot I_{t-a,m,r,s} \cdot f_{t-a,m,r,s}(a)}{\sum_a I_{t-a,m,r,s} \cdot f_{t-a,m,r,s}(a)}, \quad w'(a) <$	Stage 2 core suitability: an age-weighted average over the cohorts retiring this year, with a weight that declines in age, so an early retirement supplies a better remanufacturing core than a late one.
Eq-END-05	$\rho^{rem,eff}_{t,m,r,s} = \rho^{rem,anc}_{t,m,r,s} \cdot clip(1 + \gamma_m (\Psi_{t,m,r,s} -$	Stage 2 remanufacturing response. The multiplier is bounded, and mean-preserving: at run-average suitability and price it equals one, so remanufacturing follows its anchor and departs from it only as conditions depart from typical.
Eq-END-06	$Q_{t,m,r,s} = \frac{P^{use}_{t,m,r,s} \cdot Q^{pri}_m + X^{use}_{t,m,r,s} \cdot (Q_{t-1,m,r,s} - \lambda^Q_m)}{P^{use}_{t,m,r,s} + X^{use}_{t,m,r,s}}$	Stage 3 quality mixing balance: primary input enters at a fixed reference quality; secondary input enters at the previous cycle's quality, less a per-cycle loss term. With a zero loss term, the index equals the primary reference for any feed mix.
Eq-END-07	$\Delta Q_{t,m,r,s} = Q_{t,m,r,s} - Q^{bas}_{t,m,r,s} \leq 0$	Quality deviation feeding back into Eq-END-01. It cannot be positive, so degraded quality can only ever push material toward downcycling, never toward recycling.

System Dynamics formulation

Purpose and boundary. The SD layer is the model's behavioural component, deterministic and reduced-form as set out in Goal and Scope, with a background stochastic noise process layered on top of that core (below). Four mechanisms are implemented, developed in turn below: a scarcity-price-demand channel, a capacity-envelope/bottleneck loop, a collection-response loop, and a strategic-reserve intent policy.

Mechanism 1 — scarcity-price-demand channel. Coupling supplies a scarcity multiplier, which bottleneck pressure amplifies to an effective scarcity multiplier; price responds to effective scarcity; and demand responds to price, but only from a configured demand-response start year (a calendar gate distinguishing the historical/warm-start period from the reporting period in which price feedback is judged reliable). A long-run elasticity refinement (below) also ramps demand's price sensitivity upward over a decade following the short-run response. Table 13 gives the channel in full.

Table 13. *Scarcity-price-demand channel (Eq-SD-01 to Eq-SD-05b).*

Eq.	Equation	Controls and diagnostics
Eq-SD-01	$\Sigma^{eff}_{t,m,r} = \Sigma_{t,m,r} \cdot (1 + b^{\Sigma}_m \cdot \beta_{t,m,r})$	bottleneck_scarcity_gain; SD_scarcity_multiplier_effective

Eq.	Equation	Controls and diagnostics
Eq-SD-02	$\pi_{t,m,r} = \pi^{0}_{m,r} \cdot (1 + \sigma^{\pi}_{m}(\Sigma^{eff}_{t,m,r} - 1))$	price_base, price_scarcity_sensitivity
Eq-SD-03	$\theta_{t,m,r} = \frac{\pi_{t,m,r}}{max(\pi^{0}_{m,r}, \xi)}$	SD_price_ratio
Eq-SD-04	$e_{t,m} = e_{m} \cdot 1[t \geq t^{resp}_{m}]$	demand_response_start_year: a calendar gate separating the warm-start period from the reporting window in which price feedback is judged reliable
Eq-SD-05	$D_{t,m,r} = D^{des}_{t,m,r} \cdot max(0, 1 - e_{t,m}(\theta_{t,m,r}$	demand_price_elasticity; Service_demand
Eq-SD-05 b	$e^{lr}_{t,m} = e_{m} + (e^{LR}_{m} - e_{m}) \cdot min(1, \frac{t -$	demand_price_elasticity_lr, lr_ramp_years: the long-run elasticity phases in over a configured ramp rather than applying immediately

Mechanism 2 — capacity envelope and bottleneck loop. Flow capacity is set from exogenous demand scale and a capacity envelope; utilisation above one creates bottleneck pressure; capacity pressure blends physical shortage (bottleneck) and price channels; a raw capacity target responds to that pressure (expanding under shortage, retiring — more slowly, asymmetrically — under persistent low price); the target is bounded to a configured envelope range; and the realised envelope follows the target with a lag, representing investment/retirement delay rather than instantaneous adjustment. Table 14 gives the equations.

Table 14. *Capacity envelope and bottleneck loop (Eq-SD-06 to Eq-SD-12b).*

Eq.	Equation	Controls and diagnostics
Eq-SD-06	$K_{t,m,r} = max(\xi, D^{exo}_{t,m,r} \cdot max(\kappa_{t,m,r}, \xi))$	capacity_envelope
Eq-SD-07	$\nu_{t,m,r} = \frac{D^{des}_{t,m,r}}{K_{t,m,r}}$	SD_flow_utilization
Eq-SD-08	$\beta_{t,m,r} = max(0, \nu_{t,m,r} - 1)$	SD_bottleneck_pressure
Eq-SD-09	$\Pi_{t,m,r} = w_{m} \cdot \beta_{t,m,r} + (1 - w_{m}) \cdot max(0, \theta_{t,m,r} - 1)$	capacity_pressure_shortage_weight (w): blends physical shortage and price channels
Eq-SD-10	$\kappa^{*}_{t,m,r} = 1 + g^{+}_{m} \cdot \Pi_{t,m,r} - g^{-}_{m} \cdot \omega^{ret}_{m} \cdot max(0, 1 - \theta_{t,m,r})$	capacity_expansion_gain, capacity_retirement_gain, retirement_slack_weight: expansion and

Eq.	Equation	Controls and diagnostics
		retirement respond asymmetrically
Eq-SD-11	$\kappa_{t,m,r}^{tgt} = clip(\kappa_{t,m,r}^{*}, \kappa_{m}^{-}, \kappa_{m}^{+})$	capacity_envelope_min, capacity_envelope_max
Eq-SD-12	$\frac{d\kappa_{t,m,r}}{dt} = \frac{\kappa_{t,m,r}^{tgt} - \kappa_{t,m,r}}{\tau_{m}^{cap}}$	capacity_adjustment_lag_years: first-order smoothing representing investment and retirement delay; SD_capacity_envelope
Eq-SD-12b	$\hat{D}_{t,m,r} = \frac{1}{H_m} \sum_{h=0}^{H_m-1} D_{t-h,m,r}$	expectation_horizon_years: capacity responds to demand integrated over an investor perception horizon, not to the instantaneous value

Mechanism 3 — collection-response loop. Collection effort is an integrator state: it builds under price pressure (and, through a configured circular-collection gain, under circular-supply stress passed from coupling) and decays back toward a baseline effort level; bottleneck pressure throttles how much of that effort translates into realised collection response; and the resulting collection multiplier — bounded to a configured range — scales the dMFA layer's baseline exogenous collection rate. This is the model's principal direct route for circularity pressure to raise collection response without first requiring price-driven scarcity. Table 15 gives the loop in full.

Table 15. *Collection-response loop (Eq-SD-13 to Eq-SD-18).*

Eq.	Equation	Controls and diagnostics
Eq-SD-13	$\tau_{t,m,r}^{col} = \frac{1}{1 + b_{m}^{col} \cdot \beta_{t,m,r}}$	bottleneck_collection_sensitivity: physical bottleneck throttles how much effort converts into realised collection
Eq-SD-14	$B_{t,m,r} = g_{m}^{B}(max(0, \theta_{t,m,r} - 1)$	collection_effort_build_gain, coupling_circular_collection_gain: the second term is the model's direct route for circularity pressure to raise collection without first requiring price scarcity

Eq.	Equation	Controls and diagnostics
Eq-SD-15	$\Delta_{t,m,r} = g_m^D \cdot max(0, \varepsilon_{t,m,r} - \varepsilon_m^0)$	collection_effort_decay_gain, collection_effort_baseline
Eq-SD-16	$\frac{d\varepsilon_{t,m,r}}{dt} = B_{t,m,r} - \Delta_{t,m,r}$	collection_effort_initial: effort is an integrator state, so collection response is lagged rather than instantaneous
Eq-SD-17	$\mu^*_{t,m,r} = \varepsilon_{t,m,r} \cdot m^{shk}_{t,m,r} \cdot \tau^{col}_{t,m,r}$	collection_shock_mult: the scenario shock channel enters here
Eq-SD-18	$\mu^{col}_{t,m,r} = clip(\mu^*_{t,m,r}, \mu^-_m, \mu^+_m)$	collection_multiplier_min, collection_multiplier_max; Coupling_collection_multiplier, applied at Eq-MFA-17

Mechanism 4 — strategic reserve intent policy. Fill and release are two independently gated channels. Fill activates only when price and service stress are both low (avoiding pro-cyclical buying into a rising market) and then closes a configured share of the coverage gap between the target and current coverage-years signal, subject to a maximum fill rate. Release activates on emergency pressure — the larger of a service-stress overshoot or a price-ratio overshoot beyond configured release thresholds — and releases proportionally to that pressure, subject to a maximum release rate. Both channels are gated to zero whenever the reserve is disabled for the material in question. Table 16 states the policy equations.

Table 16. *Strategic reserve intent policy (Eq-SD-20 to Eq-SD-24).*

Eq.	Equation	Controls and diagnostics
Eq-SD-20	$\Gamma_{t,m,r} = max(0, Y^*_m - Y^{res}_{t,m,r})$	strategic_reserve_target_coverage_years
Eq-SD-21	$1^{fill}_{t,m,r} = 1[\theta_{t,m,r} \le \theta^{fill}_m] \cdot 1[\Sigma^{srv}_{t,m,r} \le \Sigma^{fill}_m]$	fill_price_threshold, fill_service_threshold: filling requires both a calm price and adequate service
Eq-SD-22	$\varphi^{fill}_{t,m,r} = clip(1^{fill}_{t,m,r} \cdot min(\varphi^{max}_m, g^{fill}_m \cdot \Gamma_{t,m,r}), 0, 1)$	strategic_reserve_fill_gain, max_fill_rate; identically zero where the reserve is disabled
Eq-SD-23	$\Omega_{t,m,r} = max(0, \Sigma^{srv}_{t,m,r} - \Sigma^{rel}_m, \theta_{t,m,r} - \theta^{rel}_m)$	release_service_threshold, release_price_threshold: either channel alone can trigger release
Eq-SD-24	$\varphi^{rel}_{t,m,r} = clip(min(\varphi^{rel,max}_m, g^{rel}_m \cdot \Omega_{t,m,r}), 0, 1)$	strategic_reserve_release_gain, max_release_rate

Additional demand and capacity refinements. Several further refinements sit on the four core mechanisms, implemented directly in the SD runtime rather than as separate equation cards. A long-run price elasticity (`demand_price_elasticity_lr`, roughly three times the short-run anchor) phases in over a configured ramp period (`lr_ramp_years`) rather than applying immediately. An expectation-horizon smoothing (`expectation_horizon_years`) represents investor perception lag, under which capacity expansion responds to demand integrated over that horizon rather than to the

instantaneous value. A backlog-clearance rate (`backlog_clearance_rate`) governs what fraction of accumulated unmet-demand backlog is cleared per year. A retirement-slack weight (`retirement_slack_weight`) on the idle-capacity retirement arm of the capacity-target equation tunes how quickly the envelope recovers from oversupply and helps set the characteristic period of any capacity cycle the model produces. Finally, a background stochastic noise process runs on demand and collection, active from each material's historical-dynamics start year onward. Autocorrelation (`noise_ar1_rho`) and variance (`noise_sigma_demand`, `noise_sigma_collection`) parameters are calibrated per material to reflect real-world demand and collection volatility rather than a single global default;[7] a fixed seed (`noise_seed`) keeps a given run reproducible and can be varied across runs for ensemble comparisons, and the process composes with, rather than substitutes for, separately configured scenario shock and surge multipliers. It represents background variability, not formal parameter-uncertainty propagation: the runtime executes one seeded trajectory per configuration rather than a Monte Carlo ensemble, so it does not generate confidence intervals.

Each refinement softens one idealisation in the core mechanisms; together they keep behaviour plausible at both ends of the horizon without adding feedback loops. The long-run elasticity and expectation-horizon smoothing stop demand and capacity over- or under-reacting to a single unusual year, reflecting that substitution, efficiency and investment responses build gradually rather than jumping to a new equilibrium. The backlog-clearance rate and retirement-slack weight govern how quickly the system sheds accumulated stress or idle capacity, which is what lets the same core equations produce a plausible multi-decade capacity cycle rather than a permanently oscillating or permanently stuck system. The noise process ensures that deterministic core mechanisms still produce year-to-year variability of realistic magnitude, so historical fit is not judged against an artificially smooth trajectory. None of the five change what the mechanisms compute on average; they change how quickly and smoothly the model approaches that average.

Historical warm-start and the two activation gates. The model's exogenous data starts in 1950, but several materials had already been mined, used and accumulated in stock for decades or centuries before then. Starting the simulation with zero in-use stock in 1950 would therefore understate reality for those materials and never fully catch up. To avoid this, each material is given a *burn-in*: a material-specific run-up period before 1950 in which the model builds a physically plausible starting stock rather than starting from nothing. Copper and lithium use a 20-year burn-in and cobalt a shorter 10-year one; nickel — mined and used industrially since the nineteenth century — uses a much longer 80-year burn-in starting from 1870, so that its substantial pre-1950 accumulation is actually captured rather than assumed away; silicon uses a 33-year burn-in, and because directly measured silicon demand data are not available that far back, the burn-in period is instead run against a reconstructed demand estimate for 1870–1970, built from the best available historical proxies for silicon use over that period, standing in for real observations where none exist.

Two separate on/off switches, or *gates*, control when different behaviours start, and are kept independent. The first switch (`sd_dynamics_start_year`, set per material) turns on the historical/burn-in behaviour described above, including the background noise process. The second switch (`report_start_year_bptk`) turns on price-driven demand response, and only from the start

[7] For example, cobalt carries elevated demand noise reflecting Democratic Republic of the Congo(DRC) supply concentration, and battery-linked materials carry elevated noise reflecting policy and EV-adoption uncertainty.

of the reporting window (2020) onward — before that point, price feedback stays off even while the rest of the historical dynamics are running. They are separate because the two are validated differently: burn-in is judged on whether it reproduces plausible historical stock levels, while price response is a reporting-window mechanism not meant to be tested against history. As a safety check, turning the historical warm-start off entirely is required to reproduce exactly the same results as a "cold-start" run beginning directly in the reporting window — confirming that warm-start only changes how the model is initialised, not what it computes once running.

During the burn-in period, historical demand is not left entirely to the model's internal logic, nor is it forced to match the historical data exactly — instead, it blends the two. A *tether strength* dial, set between 0 and 1, controls the blend: at 1.0 (full tether), demand exactly follows the known historical series; at 0.0, demand is computed entirely by the model's own scarcity, price and capacity mechanics — this is always the setting used in the reporting window from 2020 onward. During a handful of specific, sourced historical episodes — for example a known supply shock or price spike — the tether is loosened to an intermediate value (roughly 0.05–0.30) for that window only, letting the model's own dynamics partly show through and be tested against what actually happened, while the historical data still anchors the overall trajectory so the model cannot drift arbitrarily far from known reality. Table 17 gives the accepted per-material warm-start settings.

Table 17. *Accepted five-material SD warm-start candidate (sd_warm_calibrated). For definition and sources of the historical events chosen for the replicability gate refer to Table 45*

Material	Dynamics start	Evidence status
Copper	1950	Root-mean-square error (RMSE) gate passed; Chile supply event reproduced.
Nickel	1980	RMSE gate passed; LME and Indonesia events reproduced.
Silicon	2005	Polysilicon shortage event reproduced; RMSE gate comparator provisional (not yet independently validated).
Cobalt	1950	Event mechanisms active; RMSE gate comparator provisional (not yet independently validated).
Lithium	2010	Electric vehicle-boom supply event reproduced; RMSE gate comparator provisional (not yet independently validated).

Warm-start acceptance is judged on two independent gates — an event-window deviation score and a stock-trajectory gate — and no material clears both without qualification. Copper is the only material whose event evidence and stock comparator are both balanced. Nickel's stock comparator is independently sourced, but one of its three scored episodes does not reproduce. Cobalt, lithium and silicon reproduce all scored episodes, but the observed stock series their trajectory gate is checked against is derived from an earlier run of the model's own baseline rather than sourced independently, so a pass or fail on that gate reflects agreement with that snapshot rather than validation against observation. This is a governed and explicitly flagged limitation. The gates, the per-material results and the consequences for interpretation are set out in full in Calibration and in Assumptions and Limitations.

Coupling methodology

Purpose and boundary. Coupling defines how SD and dMFA exchange signals and how convergence between them is judged. Coupling coordinates module interaction; it does not replace either solver. The architecture is *loose-iterative*: SD and dMFA are solved sequentially, each over the full modelled time horizon, with lagged feedback signals passed between them and the pair re-solved under iteration until convergence. This is explicitly not year-by-year co-simulation — the model does not run sequentially one year at a time, with the two calculation engines (solvers) exchanging information and reconciling results during each year before moving to the next year; instead, an entire submodel trajectory is solved before control passes to the other, and successive passes refine the joint solution over the full horizon.

Inner loop (per material–region slice). Each coupling iteration: (1) builds SD inputs using the prior iteration's feedback signals; (2) runs SD over the full year vector to obtain demand, the collection multiplier, and strategic fill/release intents; (3) builds dMFA inputs from those SD outputs and runs dMFA over the full year vector; (4) computes MFA-derived feedback targets — physical ceilings and constraints — from the dMFA result; (5) smooths and, depending on configuration, gates the update of feedback signals; and (6) computes a convergence metric and decides whether to stop or continue. The three dMFA-to-SD signals introduced in Goal and Scope — service stress, circular supply stress and strategic stock coverage — are formalised in the equations below; SD returns desired demand, product shares, and strategic fill and release intents in the opposite direction. Table 18 gives the inner-loop equations.

Table 18. *Inner-loop coupling equations (Eq-CPL-01 to Eq-CPL-08). Four additional coupling equations govern the outer trade loop and are not reproduced here because the origin–destination trade layer they belong to is inactive in the configuration described.*

Eq.	Equation	Interpretation
Eq-CPL-01	$\Sigma^{srv}_{t,m,r} = clip(1 - \frac{\Phi_{t,m,r}}{max(D^{srv}_{t,m,r}, \xi)}, 0, 1)$	Refined-boundary service deficit, defined on exactly the quantities the reported service indicator uses, so the signal and the indicator cannot diverge.
Eq-CPL-02	$\Sigma^{cir}_{t,m,r} = clip(1 - \frac{\sum_s X_{t,m,r,s}}{max(\sum_s (P_{t,m,r,s} + X_{t,m,r,s}), \xi)}$	Primary-dependence pressure. It is a composition measure, not a physical shortage signal: it can be high in a well-supplied year simply because secondary supply is small.
Eq-CPL-03	$Y^{res}_{t,m,r} = \frac{\sum_s S^{res}_{t,m,r,s}}{max(D^{srv}_{t,m,r}, \xi)}$	Reserve coverage in years of demand, with a guarded denominator.
Eq-CPL-04	$\Sigma_{t,m,r} = min(1 + c^{srv}_m \cdot \Sigma^{srv}_{t,m,r} + c^{cir}_m \cdot \Sigma^{ci}_{t,r}$	The scarcity multiplier the SD layer receives, capped. The circular-supply gain is zero in the active baseline, so only the service channel drives price there.
Eq-CPL-05	$\Sigma^{(n+1)} = (1 - \alpha)\Sigma^{(n)} + \alpha\Sigma^{tgt}$	Under-relaxation of each feedback signal between coupling iterations, at the configured smoothing factor. The

Eq.	Equation	Interpretation
		strategic channel carries its own smoothing constant (Eq-CPL-06).
Eq-CPL-06	$Y^{(n+1)} = (1 - \alpha^{res}) Y^{(n)} + \alpha^{res} Y^{tgt}$	The same update applies to the reserve-coverage channel, with a separate smoothing factor.
Eq-CPL-07	$\Upsilon^{(n)} = max(\|\Delta\Sigma\|, \|\Delta\mu^{col}\|, \|\Delta\varphi\|, \|\Delta\rho$	Inner-loop convergence norm across all channels. The last two terms are the Stage 1 and Stage 2 endogenous-rate deltas and are identically zero when those stages are off.
Eq-CPL-08	$\Upsilon^{(n)} < \tau^{cpl} \quad or \quad n \geq n^{max}$	Stop criterion: the loop halts on convergence or on the iteration cap, whichever comes first.

Eq-CPL-07 includes the Stage-1 routing-split delta and Stage-2 remanufacturing-efficiency delta terms because Stages 1 and 2 of the endogenous circular-loop response are enabled by default in the active baseline (see Dynamic MFA formulation, endogenous circular-loop equations); both terms are zero whenever those stages are switched off.

OD-trade (inactive capability). A further origin–destination trade layer is implemented in code as a general capability: when enabled, it would allocate flows across regions for three commodity channels — concentrates, refined metal, and scrap — through an outer loop wrapped around the inner coupling solution described above. It is not developed further here because it is disabled in the active configuration: the canonical run entrypoint selects a trade preset that switches it off while preserving the underlying trade schema for future use. All statements elsewhere in this document about primary and secondary supply, service delivery, and resilience describe outcomes for the single pooled global region this configuration actually runs, not a multi-region allocation.

Active coupling configuration. The canonical coupling preset selected by the active run entrypoint is `loose_iterative_sd_full_endogenous_v1`, which sets `sd_coupling_mode: full` — the full SD response layer participates in all runs by default. A separate diagnostic entrypoint (`sd_coupling_mode: mfa_only`) exists to bypass SD entirely for SD-isolation studies, but it is not the default. The same preset enables the endogenous circular-loop response described above (Stages 1, 2 and 3 on, via `circular_loop_response_reman_low_v1`)[8] for all scenarios built on the canonical coupling configuration, so that endogenous routing response is a baseline property, not an opt-in scenario feature. Solver settings in the active configuration are a maximum of 30 inner-loop iterations,[9] a convergence tolerance of 0.02, time-series (not scalar-mean) feedback signal mode, and feedback updates applied in every modelled year rather than gated to reporting years only. The circular-supply-to-scarcity-price channel is explicitly set to zero gain in the active configuration (`coupling_circular_supply_stress_gain: 0.0`), while the circular-supply-to-collection-effort channel carries a substantial gain (`coupling_circular_collection_gain: 0.8`): low recycled

[8] Response-surface names are given here in a simplified form that matches the Stage 1/2/3 numbering used throughout this document. The literal key in the model's configuration retains an implementation-history label and reads circular_loop_response_stage4a_reman_low_v1; the two refer to the same surface.

[9] The cap was raised from an earlier default of 3 after nickel, cobalt and lithium were found to need 11-16 iterations to reach the convergence tolerance under demand-responsive and supply-capped conditions, at the configured signal smoothing of 0.2. Thirty leaves headroom above the observed requirement rather than setting it at the observed maximum.

share drives collection effort directly, without independently suppressing demand through scarcity pricing. The scarcity channel remains supported for authored sensitivity or legacy configurations, so this is a configuration choice within a more general capability, not a structural limitation.

Data

Documents the exogenous input data and the complete parameter surface that together define an executable model run, with the source and processing history of each.

A model run is fully specified by two families of authored input. The first is a set of *exogenous data artefacts* — annual time series held as versioned CSV files, covering demand, primary supply, processing yields, stream composition, end-of-life routing and product lifetimes. These drive the physical system: they are read at load time, joined onto the model's dimensional structure, and carried through the solve as state drivers. The second is a *parameter surface* — scalar or short time-series values held in configuration files, covering process yields, behavioural gains and elasticities, policy thresholds and solver controls. These shape how the model responds to the data rather than supplying quantities directly. This section documents both in full: what each artefact contains, how it was constructed and from which sources, what the active configuration selects from it, and which parameters are exposed for calibration.

A structural distinction runs through the whole section and is worth stating once at the outset. Several artefacts are *scenario-resolved*: they carry a scenario column and hold multiple alternative trajectories, only one of which is active in any given run. Where this is the case, the document gives the general structure first and then names the specific selection made by the configuration described here. Historical periods are never scenario-conditional —all scenarios share an identical observed-era record, and trajectories diverge only over the forward window. This is a deliberate design choice: it prevents alternative-history trajectories from contaminating the calibrated stock that the forward period inherits.

Exogenous input data

The required artefact set. Ten exogenous files are required to execute the active baseline configuration, of which six carry the physical drivers of the material system, and four are structural companions to them. All are held in a single store under the model repository, versioned alongside the code that consumes them, so that a given commit of the model always resolves to a specific, reproducible set of input values. Table 19 summarises the artefact set; each is then developed in turn. A separate observation data set exists for calibration purposes and is addressed briefly at the end of this subsection, not in Table 19.

Table 19. *Exogenous input artefacts required to execute the active baseline configuration.*

Artifact	Role	Key dimensions	Units	Rows
final_demand	Annual material-level demand forcing (state driver)	scenario × year × material × region	t material/yr	3,465
primary_ore_input	Authored primary ore-feed anchor (state driver)	year × material × region	t ore-feed-equiv./yr	1,155

Artifact	Role	Key dimensions	Units	Rows
route_yield	Ore-to-metal conversion coefficient	year × material × region × route	dimensionless (0–1)	1,155
route_stream_share	Allocation of refined output into material streams	scenario × year × material × region × route × stream	share (sums to 1)	13,860
collection_routing_rates	Two-gate end-of-life routing	scenario × year × material × stream × region	rates (0–1)	13,860
lifetime_distributions	Age-cohort retirement kernel	cohort_year × material × stream × region × product	years / dimensionless	12,705
deposit_share, ore_grade, route_share, product_share	Source/route and product allocation companions	as per primary chain	shares / grades	placeholder-valued

Temporal and spatial coverage. All artefacts span the full 1870–2100 horizon at annual resolution (231 years) and resolve the region axis to a single `Global` value. The horizon substantially exceeds the 1950–2050 study window because the model requires a historical run-up long enough to accumulate a physically plausible in-use stock before the reporting period begins (see Model Architecture, historical warm-start), and because forward extension beyond 2050 supports sensitivity work on long-lived stock that has not yet retired by the end of the reporting window. Data quality is not uniform across this span: the observed-era core is roughly 1950–2023 depending on material, the pre-1950 period is reconstructed from historical proxies and benchmark anchors, and the post-2050 tail is an authored extension rather than a sourced projection. The three regimes should not be read with equal confidence, and the limitations subsection below identifies where each applies.

General scenario capability versus the active selection. Three of the six driver artefacts are scenario-resolved, and the configuration selects among them through two run-level fields rather than by naming scenarios file by file. The *demand scenario* field jointly selects both the total-demand trajectory and the stream-composition trajectory, on the reasoning that a world with faster clean-energy deployment has both higher battery-material demand in total and a different mix of product forms — the two cannot be varied independently without producing incoherent combinations. The *circularity ambition* field separately selects the structural end-of-life routing surface. The two are deliberately independent: demand growth and recycling policy are distinct drivers that a scenario design should be able to cross. Table 20 gives the mapping and the active selection.

Table 20. *Scenario-resolved artefacts, the run-level fields that select them, and the active configuration.*

Run-level field	Artifact(s) selected	Available values	Active selection
demand_scenario	final_demand (total demand)	aps / steps / nze	steps

Run-level field	Artifact(s) selected	Available values	Active selection
demand_scenario	route_stream_share (composition)	central / low / high	low (mapped from steps)
circularity_scenario	collection_routing_rates (routing surface)	L / M / H → low / central / high	L → low
(not scenario-resolved)	primary_ore_input, route_yield, lifetime_distributions	single trajectory	single trajectory

The active baseline is therefore a coherent STEPS-L world: demand follows the IEA Stated Policies trajectory, stream composition follows the correspondingly conservative low-transition mix, and end-of-life routing follows the current-practice L rung. It is worth flagging that the library-level defaults built into the data-loading layer are different — they resolve to an APS/central combination — so the active configuration is an explicit choice made at the run level, not an inherited default. Any figure quoted elsewhere against an APS basis will not match this baseline over the forward window, though it will match exactly over the historical period, where the scenarios are identical by construction.

Final demand

The demand series supplies annual global apparent final demand for each of the five materials and is the primary forcing on the whole physical system: it is the gross annual inflow requirement before allocation into streams, lifetime cohorts, collection pathways, and primary versus secondary supply. It is deliberately held at material level rather than stream level — the split into streams is performed downstream by a separate artefact — so that demand magnitude and demand composition remain independently sourced and independently revisable. As set out in Model Architecture, the series is defined at the *refined* boundary (refined metal entering fabrication) rather than at the into-use boundary, which is the convention its underlying sources are measured on.

Construction. The series is assembled per material from a historical reconstruction, an observed-era core, a forward scenario window, and a long-run tail. Interpolation between sparse historical anchors uses Fritsch–Carlson piecewise cubic Hermite interpolation (PCHIP) rather than a standard cubic spline, because it preserves the local shape of the anchor sequence and does not overshoot between widely spaced points — an important property when anchors are decades apart and a spline could otherwise generate spurious peaks or negative values in between. Forward trajectories over 2024–2050 are taken from the IEA Global Critical Minerals Outlook 2024 via its machine-readable Critical Minerals Data Explorer workbook *[IEA-GCMO-24],*[10] with all three scenarios built from the same source by the same method (PCHIP in log space on 2024/2030/2040/2050 milestones) so that the scenario fan is internally coherent. Beyond 2050, the series continues as an anchored-logistic extension fitted per scenario on a fifteen-year window; the same anchored-logistic form, fitted on the first fifteen years of the observed record, generates the pre-observation backcast where one is required. Table 33 gives the full source register, and Table 34 binds each source to the segments it supports.

[10] International Energy Agency. (2024). *Global critical minerals outlook 2024*. IEA Publications. Accompanying Critical Minerals Data Explorer workbook.

Per-material segment provenance. Each material's series is stitched from segments sourced differently and with sharply differing confidence, so Table 21 sets out the segmentation before the per-material narrative. Labels are those recorded in the artefact's provenance record: **EB** (evidence-based — the value sits on a documented observation or a quantitative literature anchor) and **P** (prospective — an extrapolation or scenario assumption rather than an observation).

Table 21. *Final demand: segment-level provenance per material. Source keys refer to Table 33.*

Material	Segment	Source	Method	Confidence
copper	1870–1899	back-extension from [USGS-GC-Cu][11]	Anchored logistic fitted on the 1900–1914 window	P
copper	1900–2024	[USGS-GC-Cu]	Direct; world refined copper usage, digitised from USGS/ICSG historical tables	EB
copper	2025–2050	[IEA-GCMO-24]	Direct APS column; 2025 = 27.833 Mt, 2030 = 32.620 Mt, 2050 = 39.641 Mt	EB
copper	2051–2100	extension of [IEA-GCMO-24]	Anchored logistic fitted on the 2036–2050 window	P
nickel	1870–1949	[USGS-DS140][12]	PCHIP on pre-1950 anchors (1870 = 0.2, 1890 = 3, 1910 = 22, 1930 = 55, 1945 = 130 kt, primary basis). Pre-1900 values predate the statistical record and are a reconstruction from early Sudbury and New Caledonia output	P
nickel	1950–1981	[USGS-DS140], [CLIO-NI],[13] [INSG-FB-24][14]	PCHIP on Sudbury-era anchors (1950 = 126, 1960 = 383, 1970 = 648, 1980 =	EB

[11] Glöser-Chahoud, S., & U.S. Geological Survey. *World refined copper usage, 1900–2024.* [Data set digitised from ICSG and USGS historical tables]. Retrieved from: USGS website

[12] Kelly, T. D., & Matos, G. R. (Comps.). *Historical statistics for mineral and material commodities in the United States* (U.S. Geological Survey Data Series 140). U.S. Department of the Interior.

[13] Clio Infra. *Nickel production dataset* [Data set]. Utrecht University. Compiled from British Geological Survey and U.S. Geological Survey returns.

[14] International Nickel Study Group. (2024). *The world nickel factbook 2024*. INSG.

Material	Segment	Source	Method	Confidence
			680 kt, primary basis)	
nickel	1982–2023	[IMF-T71],[15] [INSG-FB-24]	Observed 1982–1988; PCHIP bridge 1989–1994; digitised INSG series 1995–2023	EB
nickel	2024–2050	[IEA-GCMO-24]	Direct APS values; already on a total-consumption basis	EB
nickel	2051–2100	extension of [IEA-GCMO-24]	Anchored logistic on the 2036–2050 window	P
cobalt	1870–1949	back-extension from [USGS-MYB-Co][16]	Anchored logistic fitted on the 1950–1964 window	P
cobalt	1950–2020	[USGS-MYB-Co], [CI-2024],[17] [DARTON-19],[18] [GULLEY-22][19]	Piecewise-linear through apparent-consumption anchors; captures the 1986 peak at 43 kt and the 1986–1995 collapse to 22 kt	EB
cobalt	2021–2023	bridge [USGS-MYB-Co] → [IEA-GCMO-24]	Geometric interpolation, 135 kt to 220.7 kt, implied CAGR ≈ 13.1%/yr	P
cobalt	2024–2050	[IEA-GCMO-24]	Direct APS column; 2030 = 332 kt, 2050 = 453 kt	EB
cobalt	2051–2100	extension of [IEA-GCMO-24]	Anchored logistic on the 2036–2050 window	P
lithium	1870–2008	[USGS-DS140], [USGS-MCS],[20] [SVERDRUP-16][21]	Piecewise-linear on commercial apparent consumption × (1 − nuclear/defence share). The	L pre-1970; M 1970–90; H 1990–2008

[15] International Monetary Fund. (1988). Nickel: World commodity balance (Table 71). In *Primary commodities: Market developments and outlook*. IMF.
[16] U.S. Geological Survey. (Annual). Cobalt. In *Minerals yearbook*; and *Statistical compendium* (1990). U.S. Department of the Interior.
[17] Cobalt Institute. (2024). *State of the cobalt market report*. Cobalt Institute.
[18] Darton Commodities Ltd. (2019). *Cobalt market review 2018–2019*. Darton Commodities.
[19] Gulley, A. L. (2022). One hundred years of cobalt production in the Democratic Republic of the Congo. *Resources Policy, 79*, 103007. doi:10.1016/j.resourpol.2022.103007

Material	Segment	Source	Method	Confidence
			defence share nets out the roughly 42 kt of lithium hydroxide produced for the United States weapons programme between 1954 and the late 1960s, which was stockpiled rather than entering civilian use	
lithium	2009–2023	[USGS-MCS]	Direct observed; 2009 = 18.7 kt, 2023 = 181 kt	EB
lithium	2024–2050	[IEA-GCMO-24]	PCHIP-log on Data Explorer 'Total demand' row 37; 2024 base 205 kt	EB
lithium	2051–2100	extension of [IEA-GCMO-24]	Anchored logistic per scenario on the 2036–2050 window	P
silicon	1870–1969	[USGS-DS140], [SILICONES-EU][22]	Production-based proxy reconstructed from steel, manufacturing, aluminium, organosilicon and polysilicon helper series rather than from a silicon demand statistic, since none exists for this period. The values are held as fixed inputs and are not recomputed at build time	P
silicon	1970–2100	[USGS-DS140], [USGS-MCS],	Regularised ferrosilicon history	EB to 2024; P thereafter

[20] U.S. Geological Survey. (1995–2026). *Mineral commodity summaries* [Annual series]. U.S. Department of the Interior.
[21] Sverdrup, H. U. (2016). Modelling global extraction, supply, price and depletion of the extractable geological resources with the LITHIUM model. *Resources, Conservation and Recycling, 114*, 112–129. doi:10.1016/j.resconrec.2016.07.002
[22] Amec Foster Wheeler Environment & Infrastructure UK Limited. (2016). *The socioeconomic impact of the silicones industry in Europe* (Regional summary). Global Silicones Council; CES – Silicones Europe.Global Silicones Council Website & Costogue, E., Ferber, R., Hasbach, W., Pellin, R., & Yaws, C. (1979). *Silicon materials outlook study for 1980–1985 calendar years* (JPL-PUB-79-110; DOE/JPL-1012-33; NASA-CR-162541). Jet Propulsion Laboratory, California Institute of Technology. https://ntrs.nasa.gov/citations/19800006234

Material	Segment	Source	Method	Confidence
		[BERNREUTER],[23] [IEA-GCMO-25][24]	on a silicon-content basis + silicon-metal application evidence + IEA guidance to 2050 + logistic extension	

Two materials required substantial historical reconstruction beyond simple interpolation, and the model's whole stock trajectory depends on both. **Nickel** is reconstructed in five stages, the most involved of any material. Its historical anchors are held on a *primary* basis *[INSG-FB-24], [IMF-T71], [USGS-DS140], [CLIO-NI]*, but the model's convention requires total apparent consumption, so a secondary-share uplift is applied to convert between the two: total equals primary divided by one minus the secondary share, with that share itself PCHIP-interpolated between authored anchors of 0.05 (1870), 0.10 (1950), 0.18 (1970), 0.25 (1990), 0.30 (2000–2023) and 0.0 (from 2035). The primary anchors and the secondary-share curve should be understood as an internally consistent *pair calibrated to reproduce benchmark totals*, not as two independently observed quantities: the benchmark's own mid-century figures are 115/360/640/700 kt primary with 25/80/150/170 kt secondary at 1950/60/70/80, whereas the model uses 126/383/648/680 kt primary with a systematically lower secondary share. The post-uplift totals nonetheless reproduce the benchmark totals to within 0.6% at every anchor year (140.0 against 140, 442.2 against 440, 790.2 against 790, 865.0 against 870), which is the quantity the construction actually targets. The taper to zero exists because the IEA forward data already represents total demand from 2024 onward *[IEA-GCMO-24]*; applying the uplift there would double-count. Stored anchors and the values the model actually uses therefore differ materially for nickel — the 1970 anchor of 648 kt becomes 790 kt after uplift, and 2020's 2,382 kt becomes 3,403 kt — so the two should never be quoted interchangeably.

Lithium required rebuilding the entire 1870–2008 period from civilian apparent-consumption anchors *[USGS-DS140], [USGS-MCS], [SVERDRUP-16]*, because the previously used logistic form was calibrated on the post-2000 battery ramp and decayed to near-zero before 1990, understating pre-battery demand by roughly 99.7% at 1950 and 79.6% at 1990. Military lithium is deliberately excluded: the US lithium-6 government stockpile programme of roughly 42 kt total mass, concentrated in 1954–1970 *[USGS-MYB-Co]*, is a strategic reservoir with different retirement dynamics from civilian stock, and including it would overstate both in-use stock and eventual end-of-life outflow. **Silicon** over 1870–1970 is anchored to a production-based proxy rather than a direct final-demand statistic *[USGS-DS140], [SILICONES-EU]*, because demand statistics do not extend that far back; the anchors run 32 kt in 1870 to 632 kt in 1950 and 1,640 kt in 1970. **Copper** required no historical repair, being anchored throughout on observed world refined usage *[USGS-GC-Cu]*. **Cobalt** was rebuilt from apparent-consumption anchors *[USGS-MYB-Co], [CI-2024], [DARTON-19]* after the earlier nickel-host-weighted proxy was found to understate consumption by a factor of 1.7–2 over 1970–2020 and to grow monotonically, missing both the 1986 consumption peak of 43 kt

[23] Bernreuter Research. (2024). *Polysilicon market analysis*; Freedonia Group. (2024). *Global silicones*; and Wacker Chemie AG annual reports.
[24] International Energy Agency. (2025). *Global critical minerals outlook 2025*. IEA Publications. Base-case total-supply projections.

[GULLEY-22] and the 1986–1995 Cold War superalloy collapse to 22 kt, which are the two features that most shape cobalt's mid-century stock.

Table 22. *Final demand in the active STEPS scenario, selected years (kt material per year). Values through 2020 are scenario-invariant.*

Material	1950	1970	1980	2020	2050 (STEPS)	2050 (APS)	2050 (NZE)
copper	3,004	7,291	9,396	24,953	37,451	39,641	41,277
nickel	140	790	865	3,403	6,337	7,230	7,612
cobalt	14	27	40	135	390	453	503
lithium	0.58	3.25	5.53	77.3	1,170	1,540	1,834
silicon	632	1,640	2,750	8,120	16,263	15,700	14,935

Silicon's scenario ordering in Table 22 is inverted relative to the other four materials, and this is a real feature of the construction rather than an error. Silicon demand is decomposed by end use before scenario logic is applied: its ferrosilicon component follows projected crude-steel output, and its alloying component follows aluminium, both of which are *higher* under a slower-transition world, while only the solar-polysilicon component rises under faster decarbonisation, and the semiconductor and silicones components are scenario-invariant. Because ferrosilicon remains by far the largest component by mass, the steel-linked effect dominates and total silicon demand is highest under STEPS and lowest under NZE. For the other four materials, the battery and clean-energy components dominate, and the conventional ordering holds.

Primary ore input

The primary ore input series is the authored annual primary ore-feed-equivalent anchor,[25] and it is the artefact that makes the model supply-constrained rather than demand-following. It is constructed *independently of demand*. The consequence is methodologically important—a primary-supply shortfall is an observable model outcome that the resilience indicators can detect and report, rather than something the input data silently prevents from ever arising. The file stores ore-feed equivalent rather than metal, and the conversion to metal runs through the route-yield surface described next. Table 23 gives each material's historical basis, forward basis and post-2040 tail.

Table 23. *Primary ore input construction by material: historical basis, forward basis, and post-2040 tail. Sources keyed to Table 33.*

[25] "Ore-feed equivalent" is a modelling construct, not a measured mine-boundary tonnage. Polymetallic ores, co-production and by-production mean that most of these metals have no unique literal ore input of their own: cobalt is overwhelmingly a by-product of copper and nickel mining, a large share of nickel comes from laterites whose ore grade is not comparable with sulfide ore, and silicon is not mined as silicon at all. Rather than allocate shared ore bodies between metals — which would require an allocation convention the model does not need — the series is derived backwards from an authored primary metal equivalent through the active refining yield, so that ore feed is by construction whatever upstream throughput is consistent with producing that quantity of refined metal. It is therefore a harmonised upstream anchor compatible with producing the metal, not a claim about physical host-ore tonnage, and it should not be compared directly against reported mine ore-throughput statistics. The consequence is that the model deliberately abstracts away from strict mine-boundary allocation: it can speak to how much refined metal the upstream system can deliver, but not to which mine or which co-product bears the burden of delivering it.

Material	Historical construction	Forward 2024–2040	Post-2040 CAGR	Ceiling
copper	Resource-panel metal-equivalent series, 1870–2023 (logistic back-extension pre-1970; calibrated on USGS/ICSG refined usage 1970–2023)	IEA GCMO 2025[26] base-case mine supply	−4.36%/yr	2.5× (not binding)
nickel	Sudbury-era bridge 1870–1912; MISO2 primary raw production 1913–2022; USGS-informed 2023 bridge	IEA GCMO 2025 base-case mine supply	+0.30%/yr	2.5×
cobalt	Resource-panel metal-equivalent base with USGS Minerals Yearbook mine-production anchors replacing 2010–2020	IEA GCMO 2025 base-case mine supply	−3.56%/yr	3× (not binding)
lithium	USGS/benchmark piecewise-linear anchors 1870–2014; USGS MCS 2026 anchors 2015–2023	IEA GCMO 2025 base-case mine supply	−1.42%/yr	4× (not binding)
silicon	Resource-panel metal-equivalent base to 1969; USGS world silicon production 1970–2023	Resource-panel support path, since the selected IEA outlook carries no silicon mining series	+0.50%/yr	3×

The post-2040 tail requires careful reading, because for three of the five materials it is a *decline*, not a capped expansion. The rule is trend-preserving: the tail inherits the direction of the 2035–2040 interval of the sourced supply path. The IEA base case is a committed-and-announced-projects projection rather than a demand-matching one, so for copper, cobalt and lithium it peaks around 2030 and is already falling by 2040 as existing operations deplete without sufficient new capacity entering; the tail continues that fall. Only nickel and silicon carry positive tails, so the per-material reserve ceilings — expressed as a multiple of the 2040 value — bind only on those two and are inert by construction for the other three. Table 24 gives the resulting trajectories.

Table 24. *Primary metal equivalent implied by the authored supply path (Mt/yr), showing the post-2040 decline for copper, cobalt and lithium.*

Material	2024	2030	2040	2050	2075	2100
copper	22.79	24.27	15.84	10.14	3.32	1.09
nickel	3.86	4.49	4.54	4.68	5.04	5.44
cobalt	0.273	0.321	0.214	0.149	0.060	0.024
lithium	0.255	0.471	0.418	0.362	0.253	0.177
silicon	9.58	10.04	10.70	11.25	12.75	14.45

This is the single most consequential assumption in the input data set: the model's forward criticality results are driven substantially by a supply path that declines steeply for copper, cobalt and lithium while demand rises. Modelled copper primary metal equivalent falls from a 2030 peak of 24.3 Mt to 15.8 Mt by 2040 and roughly 1.1 Mt by 2100 — a mining industry one-twentieth of its 2030 size — against copper demand rising past 37 Mt. The 2024–2040 portion has been verified value-for-value against the source workbook and is the IEA's own base case; the post-2040

[26] IEA. (2026). Critical Minerals Dataset. Retrieved from: IEA Website

continuation is an authored extrapolation of that direction. Nothing in the physical model prevents it: the supply-anchored design deliberately lets the shortfall appear as an observable service deficit rather than eliminating it with a demand-following floor. Results beyond roughly 2050 should therefore be read as the consequence of an assumed supply trajectory rather than as a projection of scarcity in their own right.

One further property of the sourced window determines what kind of constraint the model represents. The supply anchor is taken from the *mining* series, not the refining series. In the source workbook, the two diverge sharply over the same period: copper mine supply falls 34.7% between 2030 and 2040 (24,271 kt to 15,836 kt) while copper refining capacity falls only 7.5% (33,666 kt to 31,143 kt). The model's binding constraint is therefore an ore-availability constraint rather than a processing-capacity one — an important distinction when interpreting service deficits, since the two imply different policy responses and the model as configured can only speak to the former.

Silicon is a further exception on the sourced side: its post-2040 tail and 3× ceiling are authored assumptions rather than an IEA-backed mine-supply projection. Silicon's supply path also sits persistently below its demand path across the whole forward window — a gap of roughly 0.7 Mt in 2024 widening to about 4.5 Mt by 2050 — so silicon service deficits are among the more prominent model outputs and rest on the weakest of the five supply constructions. Silicon's historical construction is the least conventional of the five, because it is not a mine-production series in the sense the other four are. Silicon is not mined as silicon: it is produced by carbothermic reduction of quartz into two commodity products, ferrosilicon and silicon metal, from which every silicon, solar-grade and semiconductor-grade polysilicon — is subsequently derived. Consequently, there is no silicon equivalent of the copper or nickel mine-output statistic, and the model's primary supply series is built from world *production* of those two commodities instead.

The series is calibrated in two segments against different evidence. From 1970 onward, it is anchored on the USGS world silicon production record: the historical statistics series carries combined ferrosilicon and silicon-metal output from the mid-1960s, and the annual commodity summaries supply world totals for the recent years. All values are placed on a silicon-content basis, which is the step that makes the series internally consistent — ferrosilicon is reported as gross alloy weight and must be divided by its grade (75% for the standard grade) to yield contained silicon, while silicon metal is already 98–99% silicon and passes through essentially unchanged. Two adjustments are applied to the recent record. Recent USGS world totals exclude United States production, which is withheld as proprietary, so the published totals are incomplete; a correction adds the withheld United States tonnage back from the separate industry survey that reports it, raising the 2023 and 2024 world totals to 9.67 and 10.29 Mt against published figures of 9.0 and 9.7 Mt. The forward direction from 2025 is then aligned to the IEA outlook's silicon demand track *[IEA-GCMO-25]*, which ties the supply path to the scenario framework the other four materials use.

Before 1970, no world silicon production statistics existed, and that segment is a proxy reconstruction. Its construction is not what the phrase 'silicon content' would suggest: it does *not* multiply physical host-material tonnages by average silicon contents. It works the other way round: each of the four silicon application streams is anchored on its *observed 2020 mass* and projected backwards along a normalised activity index, so the helper series sets the *shape* of the historical path while the level is fixed by an observation at the anchor year. The specific weakness of this approach is

the assumption that the silicon intensity of steel and aluminium production was constant over 150 years, which it was not — ferrosilicon use per tonne of steel rose substantially through the twentieth century as deoxidation practice standardised, so the method probably overstates early silicon demand. No per-tonne silicon content factor enters the pre-1970 chain at any point. Table 25 gives the four streams, the driver behind each, and its source.

Table 25. *Helper series behind the pre-1970 silicon reconstruction. Each application stream is backcast from its observed 2020 mass along the normalised index shown; no silicon-content factor is applied.*

Silicon stream	Backcast driver	Coverage	Source
Ferrosilicon (iron and steel)	Normalised global iron-and-steel gross-addition-to-stock index. Cast iron is not a separate driver — the underlying series covers iron and steel together	1870–2022	MISO2 global iron_steel end-use flow series [MISO2-GAS][27]
Silicon metal to aluminium alloys	Normalised global aluminium gross-addition-to-stock index, cross-checked against a USGS aluminium-alloy end-use index normalised to 2003 = 1.0	1890–2022 (aluminium index); 1975–2003 (USGS cross-check)	MISO2 global aluminium end-use flow series [MISO2-GAS]; USGS silicon end-use worksheet [USGS-SI-USE][28]
Organosilicon chemicals (silicones, siloxanes, silanes)	Manufacturing activity proxy combined with a delayed-adoption ramp, since commercial silicone production begins only in the 1940s	1870–2019	Workbook construction on the manufacturing index; adoption ramp is authored [USGS-SI-USE]
Polysilicon (total)	Manufacturing activity proxy with a delayed-adoption ramp, since semiconductor-grade production begins in the 1950s	1870–2019	Workbook construction; adoption ramp is authored
Solar share within polysilicon	Delayed-adoption proxy schedule, replaced by observed values where they exist: 0.7% in 1977 rising to 18.8% by 1985	1977–1985 observed; proxy elsewhere	NASA/JPL flat-plate solar array project, Table 4.3a median forecast [NASA-JPL-80][29]

The application split into which these streams resolve was originally set from a published industry breakdown of metallic silicon demand — aluminium alloys 45% (midpoint of a stated 40–50% range), the organic silicon market 30%, polysilicon 20%, and a 5% refractory and other residual closing the split *[SILICONES-EU]*. That 2011 split has since been superseded for the reporting period by an annual construction driven by photovoltaic module production and polysilicon intensity, and is retained only for the historical segment and for traceability. Two independent later breakdowns bound it: a European criticality assessment putting chemical applications at 54%, aluminium alloys 38%, solar 6% and electronics 2% for the late 2010s, and a French geological-survey deck putting silicones and

[27] Wiedenhofer, D., et al. (2024). *MISO2: Global material stocks and flows database — Gross additions to stock by end use, iron and steel (1870–2022) and aluminium (1890–2022).* [Data set]. doi:10.5281/zenodo.12794253
[28] U.S. Geological Survey. *Silicon end-use statistics, 1975–2003* [Worksheet within Data Series 140]. U.S. Department of the Interior.
[29] Jet Propulsion Laboratory. (1980). *Low-cost solar array project: Silicon material task report* (NTRS 19800006234). National Aeronautics and Space Administration.

silanes at 35%, aluminium alloys with iron and steel at 41%, solar cells 18% and other 6% for 2018 *[SCRREEN2],*[30] *[BRGM-19].*[31] These are used as validation bounds rather than as inputs.

Route yield

The route-yield surface supplies the coefficient that converts authored ore-feed input into primary metal-equivalent availability. It carries three yield columns — mining/beneficiation, smelting, and refining — so that the schema can support genuinely stage-resolved processing in future, but in the active configuration the first two are set to 1 for each row, and the entire ore-to-metal conversion is carried in the refining column. That column should therefore be read as a consolidated ore-to-metal conversion proxy, not as a laboratory refining efficiency: it embeds ore grade, beneficiation recovery, and smelting losses alongside refining yield proper. Multiplying the three columns together as though they were independent process stages would be incorrect in this configuration. Table 26 gives the deployed coefficient by material and year.

Table 26. *Consolidated ore-to-metal conversion coefficient by material and year (refining_yield, per cent).*

Material	1870	1950	1970	1980	2000	2020	2050	2100
copper	1.251	0.985	0.919	0.900	0.833	0.782	0.682	0.516
nickel	1.364	1.364	1.364	1.157	1.386	1.381	1.413	1.413
cobalt	0.292	0.292	0.292	0.279	0.312	0.324	0.340	0.340
lithium	3.664	3.664	3.664	3.664	5.964	6.092	6.755	6.755
silicon	2.144	2.144	2.777	3.464	3.558	4.868	4.957	4.957

The coefficients are derived per material from global ore-and-commodity production statistics *[IRP-GMF]*[32], generally as the ratio of commodity output to ore input over a supported window (roughly 1970–2019), smoothed with a five-year centred rolling mean to damp reporting noise, then extended forward and backwards. The extension rules differ by material and encode a physical claim in each case. **Copper** is the only material whose yield continues to decline over the long run rather than holding flat, representing secular ore-grade dilution partly offset by beneficiation improvement. The coefficient falls by roughly 60% across the horizon, from 1.25% in 1870 to 0.52% in 2100 — a consequential assumption about future copper mining. **Nickel, cobalt and lithium** hold constant at their 2024 value from 2025 onward. **Lithium** additionally uses a brine-versus-spodumene route decomposition over 2000–2019, weighted by production-share anchors that shift from 60% brine in 2000 to 33% by 2019, before reverting to a single global figure. **Silicon** uses a model-calibrated proxy rather than a statistical ratio, because published silica production statistics cover broad quartzite uses well outside silicon-metal production and would give a misleading denominator. For silicon, the physical relation is a feed-to-output conversion, not a geological ore grade. Nickel's flatness is a known simplification: sulfide, laterite, ferronickel and matte routes have materially different ore bases and recoveries that a single canonical route cannot distinguish.

[30] SCRREEN2 Consortium. (2023). *Critical raw materials factsheet: Silicon metal*. European Commission Horizon 2020.
[31] Bureau de Recherches Géologiques et Minières. (2019). *Silicon metal criticality assessment*. BRGM.
[32] United Nations Environment Programme, International Resource Panel. (n.d.). *Global material flows database* [Data set]. UNEP.

Route stream share

The stream-share artefact allocates each material's refined output to its named material streams and bridges material-level demand with stream-level stock, lifetime, and collection dynamics. It is required because the model is stream-resolved: copper wire rod, lithium battery grade, cobalt superalloy metal and silicon solar-grade polysilicon have lifetimes differing by an order of magnitude, recyclability ranging from near-total to effectively zero, and entirely different retirement dynamics — none of which a single material-level demand series can express. The taxonomy is deliberately *product-form* based rather than end-use-sector based, on the reasoning that lifetime and recyclability are properties of the material form itself rather than of the sector consuming it; a tonne of stainless-grade nickel behaves the same way whether it ends up in a building or a railcar.

Stream activation. Shares are hard-clamped to zero before a stream's activation year and the displaced mass redirected to a designated fallback stream. This is a physical-plausibility constraint, not data smoothing: without it the model would route cobalt into superalloys in 1870 (superalloys date from 1907) and nickel into stainless steel before stainless steel existed. The activation years are tied to specific documented technology introductions, and the fallback streams are chosen to represent what the material was actually used for beforehand. Table 27 gives the activation year and fallback assignment for each of the twenty streams classified in Goal and Scope, Table 2.

Table 27. *Stream activation years and pre-activation fallback assignment. Stream classification is given in Goal and Scope for copper,[33] nickel,[34] [35] cobalt,[36] lithium,[37] and silicon,[38] Table 2. Activation years are indicative dates for when each material form came into industrial use, drawn from general histories of the relevant technologies. They are rounded and set the earliest year a stream may receive mass; they are not estimates of production or use in that year. Fallback streams are active from the start of the model window by construction, so their activation year is a model requirement rather than a historical claim.*

Material	Stream	Activation	Basis for activation year
copper	wire_rod_grade	1870	Copper wire predates the model window (telegraph era)
copper	alloy_grade (fallback)	1870	Brass and bronze in use since antiquity

[33] Some of the trends surrounding activation years and events are detailed in: AD Chandler, et al. (2005). Inventing the Electronic Century: The Epic Story of the Consumer Electronics and Computer Industries. Cambridge, MA: Harvard University Press; BC Blake-Coleman. (1992). Copper Wire and Electrical Conductors: The Shaping of a Technology. Chur: Harwood Academic Publishers. ISBN 978-3-7186-5200-6.

[34] Some of the trends surrounding activation years and events are detailed in:American Chemical Society (2022). Raney Nickel: A Life-Changing Catalyst. Washington, DC: American Chemical Society, National Historic Chemical Landmarks Program. Retrieved from: ACS; Ananikov, V. P. (2015). Nickel: The “spirited horse” of transition metal catalysis. *ACS Catalysis*, *5*(3), 1964–1971. doi:10.1021/acscatal.5b00072

[35] Some of the trends surrounding activation years and events are detailed in: Preece, G. *Harry Brearley and the History of Stainless Steel, 1913–2013*. Hawley Collection Trust. Retrieved from HCT

[36] Some of the trends surrounding activation years and events are detailed in: CT Sims, et al. eds. (1987). Superalloys II: High-Temperature Materials for Aerospace and Industrial Power. New York: John Wiley & Sons. Retrieved from: Internet Archive.

[37] Some of the trends surrounding activation years and events are detailed in: Yoshino, A. (2012). The birth of the lithium-ion battery. *Angewandte Chemie International Edition*, *51*(24), 5798–5800. doi:10.1002/anie.201105006; Wietelmann, U., & Klett, J. (2018). 200 years of lithium and 100 years of Organolithium Chemistry. *Zeitschrift Für Anorganische Und Allgemeine Chemie*, *644*(4), 194–204. doi:10.1002/zaac.201700394

[38] Some of the trends surrounding activation years and events are detailed in: Tangstad, M. (2013). Ferrosilicon and silicon technology. In M. Gasik (Ed.), *Handbook of ferroalloys: Theory and technology* (1st edition, pp. 179–220). Butterworth-Heinemann. Retrieved from: Google Books; Green, M. A. (2005). Silicon Photovoltaic Modules: A brief history of the first 50 years. *Progress in Photovoltaics: Research and Applications*, *13*(5), 447–455. doi:10.1002/pip.612

Material	Stream	Activation	Basis for activation year
copper	tube_sheet_plate_grade	1870	Plumbing and roofing applications established
copper	electronic_grade	1920	First commercial electronics
nickel	class_i_refined_metal (fallback)	1870	Coinage, plating and copper-nickel alloys
nickel	class_ii_stainless_feed	1913	Brearley's invention of stainless steel[39]
nickel	chemicals_catalysts	1900	Sabatier catalysis; Raney nickel commercialised 1926
nickel	battery_grade_sulfate	1990	NiCd/NiMH commercialisation, pre-dating NMC
cobalt	chemicals_pigments_catalysts (fallback)	1870	Historic pigment, glass and ceramic use
cobalt	superalloy_grade_metal	1907	Invention of Stellite; industrial superalloys from 1930s
cobalt	hardmetal_grade_powder	1923	Invention of cemented carbide
cobalt	magnets_grade_metal	1932	AlNiCo commercialisation
cobalt	battery_grade_sulfate	1991	Lithium-ion commercialisation
lithium	industrial_grade (fallback)	1870	Long-running industrial glass, grease and flux use
lithium	specialty_organolithium_metal	1900	Organolithium chemistry from the early 20th century
lithium	battery_grade	1991	First commercial lithium-ion battery
silicon	ferroalloy_grade (fallback)	1900	Commercial ferrosilicon from ~1907 (Moissan furnace 1898)
silicon	chemical_alloying_grade	1900	Aluminium-silicon alloys, early 20th century
silicon	semiconductor_grade_polysilicon	1950	Onset of the transistor era
silicon	solar_grade_polysilicon	1955	First silicon photovoltaic cell (1954)

Activation years are modelling proxies for the emergence of a commercially or technologically distinct demand stream. They do not necessarily represent the first scientific discovery of the material or technology.

Anchors and interpolation. Shares are authored as anchor years — between 10 and 17 per material — with annual values interpolated between them and each year's vector renormalised to sum to exactly one. Historical anchors over roughly 1950–2020 are transcribed from published end-use split tables — principally the USGS commodity end-use series and historical statistics *[USGS-MCS], [USGS-MYB-Co], [USGS-DS140]* — and cross-checked against industry-association statistics: cobalt's 2020 composition against the Cobalt Institute market review *[CI-2024]*, nickel's against the INSG factbook *[INSG-FB-24]* and the anthropogenic nickel cycle study *[RECK-08],*[40] copper's product-form split against industry semi-fabricated production data *[ICSG-ICA],*[41] lithiums against the global lithium

[39] Preece, G. *Harry Brearley and the History of Stainless Steel, 1913–2013*. Hawley Collection Trust. Retrieved from HCT
[40] Reck, B. K., Müller, D. B., Rostkowski, K., & Graedel, T. E. (2008). Anthropogenic nickel cycle: Insights into use, trade, and recycling. *Environmental Science & Technology, 42*(9), 3394–3400. doi:10.1021/es072108l
[41] International Copper Study Group & International Copper Association. (2020–2024). *Copper market data and semi-fabricated production statistics* [Data set].

flow study *[SUN-17],*[42] and cobalt's pre-2010 shares against the tracked-use inventory in *[HARPER-12].*[43] Forward anchors from 2025 follow directional guidance in *[IEA-GCMO-25]*. Two interpolation modes are used. Most streams use straight-line interpolation, appropriate for slow structural trends. The three battery streams — cobalt, nickel and lithium battery grades — instead use logistic (sigmoid) interpolation with a steepness parameter of 10, rescaled so the blend runs exactly from 0 to 1 across each anchor interval, because technology adoption is empirically S-shaped: cobalt's battery share moved through roughly 5%, 16%, 25%, 35%, 37%, 49% and 74% over 1990–2020, which a straight line between endpoints would badly misrepresent in the middle years.

The cobalt superalloy calibration. One anchor set departs deliberately from its source. The end-use split compiled from *[HARPER-12], [USGS-MYB-Co], [CI-2024]* gives cobalt superalloy and hardmetal shares of 50%/15% at 1950, 45%/18% at 1960, 42%/18% at 1970 and 38%/18% at 1980. Run through the lifetime and retirement machinery, these produce far more in-use superalloy stock than the independent inventory reports *[ZENG-22],*[44] which puts superalloy stock at 66 kt in 2019 against a total cobalt in-use stock of roughly 470 kt. Historical superalloy shares were therefore scaled by a factor of 0.6, with the freed share reallocated exactly to hardmetal, moving the 1950 split to 30%/35% and the 2020 superalloy share from 8% to 4.8%. Both share sets have been verified against the source tables; the pre-calibration excess is the model's own diagnostic and has not been independently reproduced here. The implication is that published end-use tables may overcount superalloy use relative to physically tracked stock, consistent with allocating on apparent consumption rather than tracked use. This is a documented calibration decision against a stock benchmark, not a data error, and three dedicated validation gates now guard against it being silently reverted. Table 28 gives the resulting 2020 composition for each material, together with the 2050 battery and electrification shares across scenarios.

Table 28. *Stream composition anchors at 2020 (share of material output), and 2050 battery/electrification stream shares across scenarios.*

Material	2020 stream composition	Key forward stream	2050 low	2050 central	2050 high
copper	wire rod 0.65 · alloy 0.10 · tube/sheet/plate 0.17 · electronic 0.08	wire_rod_grade	66%	68%	72%
nickel	stainless 0.68 · class I 0.25 · battery 0.05 · catalysts 0.02	battery_grade_sulfate	18%	25%	38%
cobalt	battery 0.74 · superalloy 0.048 · hardmetal 0.072 · magnets 0.04 · chemicals 0.10	battery_grade_sulfate	55%	70%	80%
lithium	battery 0.65 · industrial 0.25 · specialty 0.10	battery_grade	82%	85%	92%
silicon	ferroalloy 0.65 · chemical/alloying 0.245 · solar 0.097 · semiconductor 0.008	solar_grade_polysilicon	11.5 %	17.5%	20.0%

[42] Sun, X., Hao, H., Zhao, F., & Liu, Z. (2017). Global lithium flow 1994–2015: Implications for improving resource efficiency and security. *Environmental Science & Technology, 52*(5), 2827–2834. doi:10.1021/acs.est.7b06092

[43] Harper, E. M., Kavlak, G., & Graedel, T. E. (2012). Tracking the metal of the goblins: Cobalt's cycle of use. *Environmental Science & Technology, 46*(2), 1079–1086. doi:10.1021/es201874e

[44] Zeng, A., Chen, W., Rasmussen, K. D., Zhu, X., Lundhaug, M., Müller, D. B., ... & Liu, G. (2022). Battery technology and recycling alone will not save the electric mobility transition from future cobalt shortages. Nature communications, 13(1), 1341. *[As compiled in the project's cobalt calibration benchmark; the underlying publication was not re-verified directly for this document.]*

Forward trajectories are explicitly non-flat beyond 2050, and in one case non-monotonic: cobalt's battery share peaks around 2030 in the central trajectory and then *declines* toward 2100 as lithium-iron-phosphate and sodium-ion chemistries displace cobalt-bearing cathodes.

This is a substantive modelling assumption regarding battery chemistry substitution rather than a saturation artefact. As a result, the model shifts future cobalt demand away from batteries, which have historically been among the most economically attractive end-of-life cobalt recovery streams. Copper's wire-rod share continues rising under electrification, silicon's polysilicon share rises with photovoltaic deployment, lithium's battery share saturates near 85–92%, and nickel's battery share grows modestly.

Aggregation caveat. A twenty-stream taxonomy is necessarily an aggregation: several streams contain sub-uses with sharply different physics, so their blended parameters are weighted averages rather than descriptions of any actual product. Nickel's class I refined metal stream is roughly a quarter superalloys with high recovery and roughly a third electroplating with effectively zero recovery — which is why its aggregate recycling rate is only moderate despite containing a high-recovery component. Cobalt's chemicals stream mixes short-lived refinery catalysts, dissipative pigments bound permanently in fired ceramic, and a long tail of miscellaneous uses. Lithium's specialty stream is the most extreme case, and it carries a scope inconsistency as well as an aggregation one. In the benchmark end-use split from which the stream shares are taken, the stream is roughly 87% nuclear and defence metal at its 1960 Cold War peak and effectively none of it today, having been replaced by organolithium reagents and aluminium-lithium alloys. The model does not carry that defence material: as set out above, the lithium demand series is civilian apparent consumption, with the United States lithium-6 stockpile programme netted out as a separate reservoir. The stream shares, however, are transcribed from a benchmark split that includes it, so the 1960 specialty share of 54% is set by a demand component the model excludes and is then applied to civilian demand instead. Mid-century lithium is therefore over-routed into a stream whose blended lifetime and recyclability profile is itself anchored on a sub-mix dominated by material outside the modelled system. Both the routing and the retirement profile should be read as unreliable across roughly 1955–1985; the tonnages involved are small beside the battery era, but the inconsistency is real and is recorded here as unresolved rather than reconciled.

Collection and routing rates

The collection-routing artefact supplies the rates for the two-gate end-of-life schema defined in Model Architecture. It is stream-specific because recycling physics vary sharply and non-negotiably across streams: stainless-grade nickel has high functional recovery, lithium industrial uses are largely dissipative, silicon ferroalloy recovery as functional silicon is essentially zero, and battery-grade streams improve rapidly over time from a very low pre-2020 base.

How the two gates are derived. The two gates are not authored independently. Anchors are authored per stream as *unconditional* shares of total end-of-life outflow — recycling, remanufacturing and downcycling fractions of everything retiring — and linearly interpolated to annual values. The collection rate is then *derived* as the sum of those three shares, clamped to the unit interval, and each routing rate is re-expressed as conditional on collection by dividing it by that sum. The gate-2 identity therefore closes to exactly one by construction rather than by calibration.

Before a stream's activation year, the collection rate is clamped to zero, and the conditional distribution is set to an inert placeholder, so the entire end-of-life share of a not-yet-existing stream routes to system losses.

Why the schema is built this way. The earlier design used a single gate, with recycling, remanufacturing, downcycling and disposal as four mutually exclusive shares of total end-of-life outflow closing to one. Four connected reasons drove the move to two gates. **Different causes:** whether a retired product is collected depends on collection infrastructure, product form, informal-sector economics, scrap price and regulation, whereas what happens to it afterwards depends on sorting technology, contamination tolerance and refinery configuration. These are distinct processes with distinct evidence bases, and one simplex forces a single set of numbers to carry both. **Only the first stage is behaviourally responsive:** the System Dynamics layer modulates collection through an effort integrator applied as a multiplier on the exogenous rate, which needs a single scalar in the unit interval that can be scaled in isolation. Under one gate, raising collection would require scaling one share and renormalising the rest, so collecting more would silently change the fate of what was collected — better kerbside collection does not alter refinery metallurgy. That multiply-and-renormalisation option was considered and rejected precisely because renormalisation distorts which pathway absorbs the slack.

The second stage has its own causal ordering: remanufacturing is decided first, as a product-level modelling assumption about core condition applying to a small share; the remainder goes to destructive recovery and splits between functional recycling and downcycling on a quality and economics threshold. Downcycling is thus the *residual* of collected material that fails functional recovery, not an independent third lever. Conditional shares make that hierarchy explicit and give the endogenous quality-cascade and recycling-versus-downcycling responses a well-defined quantity to act on without touching how much was collected. **Finally, it matches how the evidence is reported:** policy targets and waste-authority statistics address collection, while metallurgical recovery yields are reported by process step, and the headline end-of-life recycling rate is their product. A single-gate artefact can only be compared against that product; the two-gate form lets each stage be validated against its own literature.

The functional-recycling convention. The recycling rate carries a specific and restrictive meaning: it is the share that is collected, processed, *and returned to the same material's functional supply pool*. Material that is collected and enters a metal circuit but ends up in a matrix from which the target material cannot be functionally recovered is counted as downcycling instead—never as recycling. This distinction is the reason the downcycling pathway exists as a separate accounting channel rather than being folded into either recycling or disposal, and the physical cases it captures are specific and well documented: nickel from stainless steel entering carbon-steel electric-arc furnace circuits, where it ends up below 0.05% as a tramp element; silicon in aluminium-silicon castings recovered with the aluminium but not as functional silicon; lithium in aluminium-lithium alloys oxidised to slag during secondary aluminium smelting; cobalt and nickel from batteries sent to pyrometallurgical copper smelters rather than hydrometallurgical refineries, reporting as stainless alloying element rather than battery-grade sulfate; and copper diluting into steel scrap at the tramp-element threshold. Treating these flows as recycling could overstate circularity; treating them as disposal could misstate their physical fate. The convention is applied uniformly across all five materials.

Reading downcycling against the model's closed stream structure. The word *downcycling* ordinarily implies a cascade: material that is no longer fit for its original purpose descends to a less demanding one and continues to do useful work there. That reading does not transfer to this model, and the difference is structural rather than a matter of emphasis. Each material stream here is a *closed loop*. Streams are established at refining and carried on each downstream flow, so recovered secondary material re-enters refining still labelled with the stream it came from, and can only displace primary input to that same stream. There is no intra-material substitutability: secondary copper recovered from the electronic-grade stream cannot supply the wire-rod stream, and surplus secondary nickel in the stainless stream cannot relieve a deficit in the battery-sulfate stream. Where a stream's secondary supply exceeds its own requirement, the surplus is recorded as such and left unused rather than reallocated across the metal.

Two consequences follow. First, there is no lower rung inside the metal for downcycled material to fall to — the model has no mechanism by which material failing the specification of one stream is re-credited to a less demanding stream of the same metal. Downcycled mass therefore leaves the modelled metal cycle outright: it is routed to a terminal node and passes out of the system boundary in the year it is generated. In supply terms, it is indistinguishable from an outright loss. What separates it from the uncollected residual is not its consequence for the metal but its physical fate and its reportability — it is collected, it is present in another material's circuit, and it can be quantified against the literature that measures it, whereas dissipated and landfilled mass cannot. The distinction is one of accounting and evidence, not a difference in how much of the material remains available to the metal cycle.

Second, and more consequentially for interpretation, the stream-level recycling rates in this artefact are *stream-closed* functional recovery rates and should be read as such. In the physical system, some cross-stream recovery does occur legitimately: high-purity electronic-grade copper scrap is routinely remelted into wire rod, and cathode-grade nickel can serve lower-specification alloy uses. Because the model cannot route mass sideways, such flows must be represented either as recycling within the originating stream or as downcycling out of the system; the anchors take the second reading wherever functional recovery *as the same stream* is not demonstrated. The material-level recycling rates are consequently conservative relative to a quality-cascade formulation of the same system, and the gap is widest for the metals with the most heterogeneous stream sets — copper and nickel — not for the dissipative streams, where no cascade exists. Representing genuine cascades would require either inter-stream substitutability or an explicit quality axis; the model defines the latter but leaves it dormant, making this a deferred capability rather than a data defect.

Table 29. *Material-level effective end-of-life rates at 2020 (stream-share weighted) against the literature benchmarks the build validates them on.*[45]

[45] These material-level rates are weighted differently from published ones and are not like-for-like with them. Each rate in the underlying artifact is a share of the end-of-life outflow *of one stream*; stating a single figure for a whole metal requires weighting those stream rates, and the weights used here are the composition shares of material *entering* the streams in that year. Published end-of-life recycling rates are normally weighted the other way, on the mass actually *retiring* that year, which reflects the composition of products made decades earlier rather than the composition being made now. The two diverge whenever a metal's stream mix is shifting, and for cobalt, nickel and lithium it is shifting fast: cobalt's inflow was roughly three-quarters battery-grade by 2020, while the cobalt retiring that year was dominated by superalloys and hardmetals installed long before batteries mattered. The figures here are internally consistent with the model's construction but should not be set beside headline literature rates without converting to a common basis. Reconciling the two weightings is an open item.

Material	Functional recycling	Downcycling	Total collected	Benchmark target	Source basis
copper	45.9%	3.3%	52.3%	45% ± 5%	Global copper flow study [GLOSER-13][46]
nickel	57.3%	11.4%	71.2%	57% ± 7% functional; 14% ± 5% downcycled	Global nickel cycle study [RECK-08]
cobalt	14.3%	16.5%	31.2%	16% ± 6%	Cobalt Institute / JRC material system analysis
lithium	0.5%	17.4%	18.0%	<2% ± 2% functional	[WB-2020];[47] USGS classifies as insignificant
silicon	0.5%	7.4%	8.0%	<2% ± 2% functional	USGS classifies as insignificant

Table 29's values are computed from the deployed artefacts by weighting each stream's rates on its 2020 composition share, and are scenario-invariant at 2020 as the construction requires. Two points follow. First, each functional-recycling rate sits inside its benchmark tolerance band, so the build-time gates pass — but cobalt at 14.3% sits in the lower half of its 16% ± 6% band, a direct consequence of the superalloy share calibration discussed above: superalloys carry the highest recovery rate of any cobalt stream, so cutting their share necessarily lowers the material-level functional rate. Second, downcycling is far more material than the functional rates alone suggest for the two materials usually described as unrecycled. Lithium's functional recycling is negligible at 0.5%, but 17.4% of its end-of-life mass is collected into another material's circuit — predominantly lithium in aluminium-lithium alloys lost to slag during secondary aluminium smelting. Silicon shows the same pattern at 7.4%. Reporting these materials as simply 'not recycled' would misstate where the mass actually goes.

Remanufacturing evidence status. The remanufacturing rates are the weakest-evidenced values in the data set. Recycling anchors are almost entirely evidence-based, tied to quantitative literature values. Downcycling anchors are evidence-based for the well-studied cases (nickel in stainless, copper tramp, silicon in aluminium) and assumption-based for photovoltaic and semiconductor streams. Remanufacturing anchors are predominantly low-evidence placeholders set by assumption-based parameters within physically plausible bounds, because no benchmark in the literature quantifies remanufacturing rates at the material-stream level in a form that could be used directly. Since remanufacturing is one of the flows that circularity indicators are most sensitive to — and, as noted in Model Architecture, one that the refined-boundary resilience indicators cannot see at all — this evidence gap propagates into any conclusion that rests on remanufacturing volumes specifically.

Table 30 sets that unevenness out cell by cell. The table gives the confidence label carried in the provenance record for each of the sixty stream-and-pathway combinations, together with the evidence each stream's recycling and downcycling anchors rest on. Table 35 later summarises the

[46] Glöser, S., Soulier, M., & Tercero Espinoza, L. A. (2013). Dynamic analysis of global copper flows: Global stocks, postconsumer material flows, recycling indicators, and uncertainty evaluation. *Environmental Science & Technology, 47*(12), 6564–6572. doi:10.1021/es400069b
[47] Hund, K., La Porta, D., Fabregas, T. P., Laing, T., & Drexhage, J. (2020). *Minerals for climate action: The mineral intensity of the clean energy transition*. World Bank Group.

same sixty cells as counts; this is the underlying detail, and it is the table to consult before relying on any individual stream's rates.

Table 30. *Per-stream evidence basis and confidence label for the collection-routing anchors. Labels are EB (evidence-based), P (prospective) and L (low-evidence), as recorded per material, stream and pathway in the artefact's provenance record. Table 35 gives the aggregate distribution of these sixty cells.*

Material	Stream	Recycling	Remanufacturing.	Downcycling	Evidence basis for the recycling and downcycling anchors
copper	wire_rod_grade	EB	L	EB	Building and grid wire; system-level 45% ± 5% end-of-life rate with stream ranges of 50–65% for building and transport and 80–90% for electrical infrastructure. Downcycling is copper dilution into steel scrap at the tramp-element threshold [GLOSER-13], [ICSG-ICA]
copper	alloy_grade	EB	L	EB	Brass and bronze scrap circuits; mid-range of the building and transport band [GLOSER-13], [ICSG-FB-20][48]
copper	tube_sheet_plate_grade	EB	L	EB	HVAC, plumbing and roofing recovery; building and transport band [GLOSER-13], [ICSG-ICA]
copper	electronic_grade	EB	L	EB	Printed-circuit and e-waste recovery, 15–35% band; elevated downcycling from shredded mixed scrap [GLOSER-13], [ICSG-ICA]
nickel	battery_grade_sulfate	EB	P	EB	Nickel-manganese-cobalt and nickel-cobalt-aluminium recovery of 5–15% pre-2020. Downcycling covers battery material routed to pyrometallurgical smelters, reporting as stainless alloying element rather than battery-grade sulfate [RECK-08], [JRC-MSA][49]. Forward path on recovery mandates [EU-2023/1542][50]
nickel	class_i_refined_metal	EB	L	EB	Blend of roughly 80% superalloy recovery with near-zero electroplating recovery; old-scrap efficiency of 56.2% [RECK-08], [USGS-C1196Z][51]
nickel	class_ii_stainless_feed	EB	L	EB	Stainless scrap circuit at 75–85%. The 14-percentage-point dilution of

[48] International Copper Study Group. (2020). *The world copper factbook 2020*; and ICSG. (2019). *Scrap survey*. ICSG.
[49] European Commission, Joint Research Centre. (2020). *Material system analysis of cobalt and nickel*. Publications Office of the European Union.
[50] European Parliament and Council. (2023). *Regulation (EU) 2023/1542 of 12 July 2023 concerning batteries and waste batteries*. Official Journal of the European Union, L 191.
[51] Sibley, S. F. (2009). *Nickel recycling in the United States in 2004* (U.S. Geological Survey Circular 1196-Z). U.S. Department of the Interior.

Material	Stream	Recycling	Remanufacturing.	Downcycling	Evidence basis for the recycling and downcycling anchors
					stainless scrap into carbon-steel electric-arc circuits is the canonical downcycling anchor for the whole artefact [RECK-08]
nickel	chemicals_catalysts	EB	EB	EB	Catalyst reclamation at roughly 70%; no non-functional carry-over pathway identified, so downcycling is zero on a physical argument [RECK-08], [JRC-MSA]
cobalt	battery_grade_sulfate	EB	P	EB	Below 10% recovery pre-2020, and below 5% for electric-vehicle cells. Downcycling covers black mass sent to copper and nickel pyrometallurgy rather than hydrometallurgical cobalt refining [CI-2024], [JRC-MSA]. Forward path on recovery mandates [EU-2023/1542]
cobalt	superalloy_grade_metal	EB	L	EB	Roughly 80% functional recovery, the highest of any cobalt stream and the reason the superalloy share calibration moves the material-level rate so much [CI-2024], [JRC-MSA]
cobalt	hardmetal_grade_powder	EB	L	EB	Cemented-carbide recovery at 30–40%: solid tool scrap is reclaimed by zinc-sinter or oxidation routes at roughly 93% cobalt recovery, while grinding sludge is largely landfilled [CI-2024]
cobalt	magnets_grade_metal	EB	L	EB	0–5% recovery; samarium-cobalt magnets cannot be separated economically from neodymium-iron-boron in mixed magnet scrap [CI-2024]
cobalt	chemicals_pigments_catalysts	EB	EB	EB	Aggregate of roughly 70% catalyst reclamation with fully dissipative pigments and driers, weighted on the sub-use mix [CI-2024], [JRC-MSA]
lithium	battery_grade	EB	P	EB	Below 1% global end-of-life recovery pre-2020, corroborated by the USGS 'insignificant' classification held across each edition [WB-2020], [USGS-MCS]. Forward path on recovery mandates [EU-2023/1542]
lithium	industrial_grade	EB	EB	EB	Greases, glass and ceramics are dissipative; lithium is chemically bound into the silicate matrix and cannot be recovered without dissolving the product, so all

Material	Stream	Recycling	Remanufacturing.	Downcycling	Evidence basis for the recycling and downcycling anchors
					pathways are zero [WB-2020], [USGS-MCS]
lithium	specialty_organolithium_metal	L	L	EB	Downcycling is anchored on lithium oxidised to slag and dross during secondary aluminium smelting of aluminium-lithium alloys [LIU-12][52]. The organolithium reagent sub-pool is consumed in synthesis and carries no quantitative recovery anchor, which is why the recycling label is low-evidence
silicon	ferroalloy_grade	EB	EB	EB	Silicon is oxidised to silica and reported to slag during steel and iron deoxidation; 70–80% is consumed in the process itself and the remainder is diluted in steel at 0.1–4%. No recovery pathway exists, so all three pathways are zero on a physical argument [USGS-MCS], [TAKIGUCHI-11][53]
silicon	chemical_alloying_grade	EB	L	EB	Silicon in aluminium-silicon castings is recovered along with the aluminium but persists in the secondary alloy rather than as functional silicon — the largest single downcycling flow in the silicon system [LIU-12], [TAKIGUCHI-11]
silicon	solar_grade_polysilicon	EB	P	L	Module recovery was negligible in 2020 and is emerging at 5–15% thereafter [IEA-PVPS-12][54], [IRENA-16][55]. No quantitative benchmark exists for silicon carried non-functionally in recovered module glass, so the downcycling anchor is assumption-based
silicon	semiconductor_grade_polysilicon	EB	L	L	Wafer kerf loss of 40–50% is prompt scrap arising before service and is accounted for at fabrication, not here; device-grade end-of-life recovery is negligible [TAKIGUCHI-11]. No anchor exists

[52] Liu, G., Bangs, C. E., & Müller, D. B. (2013). Stock dynamics and emission pathways of the global aluminium cycle. *Nature Climate Change, 3*, 338–342. doi:10.1038/nclimate1698; and Løvik, A. N., Modaresi, R., & Müller, D. B. (2014). Long-term strategies for increased recycling of automotive aluminum and its alloying elements. *Environmental Science & Technology, 48*(8), 4257–4265. doi:10.1021/es405604g

[53] Takiguchi, H. (2011). Global flow analysis of crystalline silicon. In *Crystalline silicon — Properties and uses*. IntechOpen. doi:10.5772/23573; and Takiguchi, H., & Morita, K. (2009). Sustainability of silicon feedstock for a low-carbon society. *Sustainability Science*, *4*(1), 117–131.doi:10.1007/s11625-009-0069-1; and Williams, E. (2003). Forecasting material and economic flows in the global production chain for silicon. *Technological Forecasting and Social Change*, *70*(4), 341–357.doi:10.1016/S0040-1625(02)00201-9

[54] International Energy Agency, Photovoltaic Power Systems Programme. (2024). *End-of-life management of photovoltaic panels* (IEA-PVPS Task 12). IEA-PVPS.

[55] International Renewable Energy Agency & IEA-PVPS. (2016). *End-of-life management: Solar photovoltaic panels*. IRENA.

Material	Stream	Recycling	Remanufacturing.	Downcycling	Evidence basis for the recycling and downcycling anchors
					for silicon embedded in discarded chips

Battery-stream scenario construction. Seventeen of the twenty streams have their low and high scenario trajectories derived mechanically from the same anchor logic that produces the central trajectory. The three battery streams are the exception: their central trajectory uses a shared collection curve with metal-specific recycling yields, so their low and high trajectories are authored directly, grounded on the EU Batteries Regulation recovery targets *[EU-2023/1542]* and IEA outlook guidance *[IEA-GCMO-25]*. Within each authored scenario, recycling and remanufacturing are set as the raw anchor divided by the scenario collection rate, and downcycling takes the residual. Table 31 gives the collection trajectories.

Table 31. *Authored battery-stream collection trajectories by scenario (share of end-of-life outflow collected). Applies jointly to the cobalt, nickel and lithium battery streams.*

Scenario	2025	2030	2040	2050	2075	2100
low (active baseline)	0.35	0.40	0.46	0.50	0.53	0.55
central	0.35	0.45	0.57	0.65	0.70	0.75
high	0.35	0.55	0.70	0.80	0.85	0.88

The scenarios alter collection and the recycling/downcycling split; they do not directly prescribe substantial growth in remanufacturing, which is represented through the endogenous response mechanism. The differences between the low, central and high trajectories are differences in *collection* — how much end-of-life material is captured — and in how the captured material splits between recycling and downcycling. They deliberately hold remanufacturing almost flat. That is not because the scenarios assume remanufacturing does not grow; it is because remanufacturing ambition is represented elsewhere, in the endogenous Stage 2 response surface described in Model Architecture, which raises remanufacturing in response to model conditions rather than by scenario assumption. A reader who compares the low and high rows of this artefact and concludes that the model's ambitious circularity scenario involves little additional remanufacturing would be drawing the wrong conclusion from the right numbers.

Product lifetime distributions

The lifetime artefact supplies the age-cohort retirement kernel: for each inflow cohort entering use in a given year, it specifies the probability density of that cohort's retirement timing. Integrated over the accumulated stock, this kernel determines the annual end-of-life outflow available to collection and recycling. It is the single most consequential artefact for secondary supply timing, because it alone determines *when* material placed into use decades ago becomes available for recovery — and therefore how much secondary supply a given recycling rate can deliver in any particular year.

Distribution families. Three distribution families are used, chosen per stream on physical grounds rather than statistical fit. A *normal* distribution is used where a stream maps predominantly to products with a well-defined, roughly symmetric service life governed by design codes or standards

— buildings, vehicles, photovoltaic modules. A *lognormal* is used where the stream is dominated by consumable or dissipative uses with a long right tail: cutting inserts consumed in months, catalysts regenerated over varying cycles, consumer electronics with a minority of long-lived industrial units. The lognormal is parameterised so that the stored mean is the arithmetic mean of the lifetime itself, not the log-space mean, with the shape parameter stored separately. A *two-component normal mixture* is used for the five streams whose retirement is genuinely bimodal, where a single unimodal distribution cannot simultaneously reproduce both the long-run in-use stock and the early end-of-life outflow. Those five are copper wire rod, nickel class II stainless feed, and all three battery streams. Fifteen streams hold their parameters constant across all cohort years; five vary by cohort year to represent documented technology transitions.[56] Table 32 gives the deployed parameterisation for each stream.

Table 32. *Product lifetime distribution by material stream, in the deployed configuration.*

Material / stream	Distribution	Parameters	Physical basis
copper / wire_rod_grade	mixture (cohort-varying)	long normal(55, 15) + short normal(15, 6); short-mode weight 0.75 → 0.375 over 1950–2000, holding implied mean 25 → 40 y	Cable insulation shift from impregnated paper (~25 y) to cross-linked polyethylene (~40 y), plus faster-turnover equipment wiring
copper / alloy_grade	normal	mean 18 y, sd 7	Plumbing fittings, automotive thermal hardware, bearings — medium-lived industrial pool
copper / tube_sheet_plate_grade	normal	mean 24 y, sd 9	HVAC tube (15–25 y) against long-lived plumbing and roofing (40–80 y)
copper / electronic_grade	lognormal	mean 8 y, sigma 0.4	Consumer electronics peak with an industrial and high-reliability right tail
nickel / class_i_refined_metal	normal	mean 20 y, sd 6	Blend of electroplating (8–20 y), aerospace superalloy (20–30 y) and specialty alloys
nickel / class_ii_stainless_feed	mixture	long normal(55, 18) weight 0.85 + short normal(15, 5) weight 0.15; inflow-weighted mean 49 y	Genuinely bimodal: long-lived buildings and process plant against short-lived transport, consumer and medical goods
nickel / battery_grade_sulfate	mixture (cohort-varying)	long normal(13, 3) + short normal(3, 1.5); short-mode weight 0.35 → 0.30 by 2035	EV and stationary packs against consumer-electronics and early packs

[56] Cobalt is not uniformly cohort-invariant: its battery-sulfate stream varies by cohort, while the superalloy, hardmetal, magnet and chemical streams are held constant across the whole 1870–2100 range. For those four, the deployed means — 15, 3, 15 and 3.5 years respectively — are hand-tuned calibration settings that supersede the literature-derived parameterisation the project's own generator still encodes: a 15-to-25-year cohort ramp over 1960–2000 for superalloys, following documented gains in commercial turbine on-wing life *[REED-06]*, , and means of 5, 20 and 11 years for hardmetal, magnets and chemicals. Only the means differ; the dispersion parameters are unchanged. The generator is frozen and does not reproduce the deployed file, which is recorded as the authority with the divergence held open as a deferred reconciliation. No derivation is recorded for the deployed means, so they should be read as settings calibrated against observed stock rather than as sourced lifetime estimates.

Material / stream	Distribution	Parameters	Physical basis
nickel / chemicals_catalysts	lognormal	mean 5 y, sigma 0.6	Short-cycle refinery hydrotreating and hydrogenation catalysts
cobalt / battery_grade_sulfate	mixture (cohort-varying)	long normal(12, 3) + short normal(3, 0.95 rising to 1.5); short-mode weight 0.944 through 2000 → 0.833 (2010) → 0.711 (2020) → 0.333 from 2035	Same bimodal battery structure, with cobalt-specific effective mean
cobalt / superalloy_grade_metal	normal	mean 15 y, sd 7, constant across all cohort years	Doubling of commercial turbine on-wing life with thermal barrier coatings and single-crystal blades
cobalt / hardmetal_grade_powder	lognormal	mean 3 y, sigma 0.7	Cemented-carbide inserts consumed in months to years; specialty dies extend the tail to 10–15 y
cobalt / magnets_grade_metal	normal	mean 15 y, sd 6	Samarium-cobalt in aerospace actuators, high-reliability motors and medical devices
cobalt / chemicals_pigments_catalysts	lognormal	mean 3.5 y, sigma 0.841	Mixture of ~1–5 y refinery catalysts, ~30 y pigments bound in fired ceramic, and miscellaneous chemicals
lithium / battery_grade	mixture (cohort-varying)	long normal(12, 3) + short mode itself cohort-varying: normal(3, 1.5) through 2010, shortening to normal(1.5, 0.75) for 2011–2015 cohorts, then ramping back to normal(3, 1.5) by 2020; short-mode weight 1.0 through 2017 → 0.722 (2020) → 0.333 from 2035	Transition from portable consumer electronics to EV and stationary storage
lithium / industrial_grade	lognormal (step at 2000)	mean 3 y before 2000, 4 y from 2000; sigma 0.6	Use mix shifting from long-lived glass and ceramics toward greases, desiccants and smelter flux
lithium / specialty_organolithium_metal	lognormal	mean 2 y, sigma 0.6	Polymerisation initiators and pharmaceutical synthesis consumed per campaign
silicon / ferroalloy_grade	normal	mean 65 y, sd 20	Silicon locked into the steel matrix; construction-dominated (~70%) steel product lives

Material / stream	Distribution	Parameters	Physical basis
silicon / chemical_alloying_grade	normal	mean 20 y, sd 6	Aluminium-silicon castings (~51%) with construction silicone sealants (~38%) at 20–30 y
silicon / solar_grade_polysilicon	lognormal → normal (step at 2010)	lognormal(6, 0.4) before 2010; normal(27, 5) from 2010	Structural shift from niche and off-grid modules to utility-scale photovoltaics
silicon / semiconductor_grade_polysilicon	lognormal	mean 6 y, sigma 0.4	Consumer electronics peak with industrial and high-reliability tail

Two of these parameterisations were set by calibration against observed stock rather than by literature value alone, and both materially change model behaviour. Cobalt's hardmetal stream is the clearest case: representing cemented-carbide cobalt on a machinery-equipment lifetime anchor of around 35 years treats the cobalt as locked in the body of the machine, when physically it sits in an insert that is consumed and replaced repeatedly. The deployed distribution is short and strongly right-skewed, with an arithmetic mean of three years, which brings modelled pre-battery cobalt stock into line with independent estimates *[HEL-19],*[57] *[ZENG-22]*; the same reasoning drives the short means on cobalt's chemicals and magnets streams. Silicon's post-2010 photovoltaic lifetime was widened from a narrow 25-year design-life assumption to a mean of 27 years with a substantially wider spread, to capture both the empirical economic-life consensus of 27–30 years *[JORDAN-13],*[58] *[FhISE-22]*[59] and the well-documented early-failure tail *[IRENA-16]* that a design-tolerance-width distribution excludes by construction.

Calibration data (inactive capability). Five further observation series — in-use stock, refining output, semis output, manufacturing output, and secondary recycling output, each covering all five materials — are present in the repository for use as calibration targets, together with an automated optimiser that can fit selected parameters against them. This capability is not exercised in the active configuration: a systematic fit against these series was carried out for copper and nickel earlier in the model's development, but the fitted values were subsequently reverted in favour of the literature-anchored per-stream defaults documented throughout this section, and the fitting workflow is retained only as a legacy/diagnostic entrypoint. The observation series therefore play no role in the results this document describes and are not developed further here (see Calibration).

Sources, provenance, validation, and data status

Consolidated source register. Table 33 lists all external sources bound to a value in the input data set, with the short key used in citations throughout this section and the specific anchors each source supports. Table 34 then binds sources to artefacts, so that the evidence base behind any single artefact can be read off directly. Each entry is a published source that can be obtained and checked independently. Where an anchor was assembled by the modelling team from several published

[57] Helbig, C., Bradshaw, A. M., Wietschel, L., Thorenz, A., & Tuma, A. (2018). Supply risks associated with lithium-ion battery materials. Journal of Cleaner Production, 172. doi:274-286. 10.1016/j.jclepro.2017.10.122
[58] Jordan, D. C., & Kurtz, S. R. (2013). Photovoltaic degradation rates — An analytical review. *Progress in Photovoltaics: Research and Applications, 21*(1), 12–29. doi:10.1002/pip.1182
[59] Fraunhofer Institute for Solar Energy Systems. (2022). *Photovoltaics report*. Fraunhofer ISE.

series rather than taken from one of them, the register names the underlying publications and the section text states the processing applied, so that no value rests on a citation a reader cannot follow through to its origin.

A small number of series are *fixed inputs* rather than quantities the data pipeline recomputes on each build, and this is stated wherever it applies rather than left implicit. Three cases arise. Silicon's pre-1970 proxy reconstruction and the battery streams' central collection trajectory were built by procedures that the current generating scripts do not reproduce; the deployed values are the authority and are stored directly. The four historical primary-supply base series were extracted once from the resource-panel metal-equivalent record and stored, so that the extraction does not have to be repeated on each build. In all three cases, the underlying sources are cited in Table 33 in the ordinary way, and the values themselves have been checked against those sources where the sources are machine-readable. What is fixed is the derivation step, not the provenance: a reader can verify what the numbers rest on, but cannot regenerate them from the scripts alone. Re-deriving the generators is a recorded open item.

Table 33. *Consolidated source register for the exogenous input data set. Citation keys used in this chapter resolve here; the Calibration and Scenarios chapters carry their own registers for the evidence specific to each, so a key should be looked up in the register belonging to the chapter that cites it.*

Key	Source	Anchors supported
[IEA-GCMO-24]	International Energy Agency. "Critical Minerals Data Explorer." 2024 edition. Paris: International Energy Agency, 2024, and its Critical Minerals Data Explorer workbook. International Energy Agency. *Critical Minerals Dataset: 2024 Edition*. Microsoft Excel workbook. Paris: International Energy Agency, 2024	Forward demand 2024–2050 for all five materials under STEPS/APS/NZE; battery design-life target of 10 years
[IEA-GCMO-25]	International Energy Agency. Global Critical Minerals Outlook 2025. Paris: International Energy Agency, 2025. (base case, total supply)	Mine-supply milestones 2024–2040 for copper, cobalt, lithium and nickel; forward stream-share and recycling scale-up direction
[IEA-WEO-24][60]	International Energy Agency. *World Energy Outlook 2024*. Paris: International Energy Agency, 2024	Crude-steel and aluminium trajectories underlying silicon's scenario decomposition
[IEA-PVPS-12]	IEA PVPS Task 12 (2024 update), PV End-of-Life Management.[61]	Photovoltaic field lifetime and waste-volume projections; emerging module recycling at 5–15%
[USGS-MCS]	U.S. Geological Survey. *Mineral Commodity Summaries 2026*. Version 1.3, May 2026. Reston, VA: U.S. Geological Survey, 2026. https://doi.org/10.3133/mcs2026. (2026 edition and the 1995–2024 series)	Lithium observed demand 2009–2023; lithium mine production 2015–2024; cobalt 2010–2020; end-use sector time series; 'insignificant' recycling classifications for Li and Si

[60] International Energy Agency. (2024). *World energy outlook 2024*. IEA Publications.
[61] S Weckend. (2016) End-of-Life Management: Solar Photovoltaic Panels, IRENA and IEA-PVPS Task 12, Report T12-06:2016. Retrieved from: IRENA. K Wambach et al. (2024). Advances in Module Recycling - Literature Review and Update to Empirical Life Cycle Inventory Data and Patent Review IEA PVPS Task 12, Report IEA-PVPS T12-28:2024. Retrieved from: IEA.

Key	Source	Anchors supported
[USGS-MYB-Co]	USGS, Minerals Yearbook — Cobalt (annual) and Statistical Compendium (1990) Ewing, Samantha M. “Cobalt.” In 2023 Minerals Yearbook. Reston, VA: U.S. Geological Survey, 2026	Cobalt apparent consumption anchors 1950–2020; pre-1990 end-use shares; US lithium-6 stockpile documentation
[SUN-19]	Sun, X., Hao, H., Liu, Z., Zhao, F., & Song, J. (2019). Tracing global cobalt flow: 1995–2015. Resources, Conservation and Recycling, 149, 45-55.	Benchmark for cobalt flows 1995-2015
[USGS-GC-Cu]	Glöser-Chahoud, S., & U.S. Geological Survey. *World refined copper usage, 1900–2024.* [Data set digitised from ICSG and USGS historical tables]. Source	Copper demand 1900–2024; 1960 = 4,738 kt, 2020 = 24,953 kt, 2024 = 27,353 kt
[GLOSER-13]	Glöser, S., Soulier, M., Tercero Espinoza, L. (2013), Dynamic analysis of global copper flows. Global stocks, postconsumer material flows, recycling indicators, and uncertainty evaluation, Environ. Sci. Technol. 47, 6564–6574.	Copper end-of-life recycling rate 45% ± 5%, stable since ~1960; copper stock-saturation lifetime anchor
[RECK-08]	Reck, Barbara K., Daniel B. Müller, Katherine Rostkowski, and T. E. Graedel. “Anthropogenic Nickel Cycle: Insights into Use, Trade, and Recycling.” Environmental Science & Technology 42, no. 9 (2008): 3394–3400. https://doi.org/10.1021/es072108l	Nickel functional recycling 57%, downcycling to carbon steel 14%, dissipation 29%; nickel end-use split at 2000
[CULLEN-17][62]	Cullen, Jonathan M., Julian M. Allwood, and Margarita D. Bambach. “Mapping the Global Flow of Steel: From Steelmaking to End-Use Goods.” Environmental Science & Technology 46, no. 24 (2012): 13048–13055. https://doi.org/10.1021/es302433p.	Host-product lifetime anchors (construction 50 y, vehicles 20 y, machinery 35 y, computers 10 y)
[HARPER-12]	Harper, E.M. et al. (2012), Tracking the metal of the Goblins: Cobalt’s cycle of use. Environ. Sci. Technol. 46, 1079–1086	Pre-2010 cobalt end-use mass shares; the in-use stock benchmark against which superalloy shares were calibrated
[CI-2024]	Benchmark Mineral Intelligence. Cobalt Market Report 2024. Prepared for the Cobalt Institute. Guildford, UK: Cobalt Institute, 2025.	Cobalt 2020 composition (battery 74%, superalloy 8%, hardmetal 4%, magnets 4%, chemicals 10%); stream-specific recovery rates
[INSG-FB-24]	International Nickel Study Group, World Nickel Factbook 2024. Retrieved from: International Nickel Study Group website	Nickel apparent consumption 1995–2023 (digitised); nickel 2020 composition (stainless 68%, battery 5%)
[IMF-T71]	International Monetary Fund, “Nickel: World Commodity Balance, 1982–88,” table 75 in *Primary Commodities: Market Developments and Outlook* (Washington, DC: International Monetary Fund, 1989)	Observed nickel apparent consumption 1982–1988

[62] Cullen, J. M., Allwood, J. M., & Bambach, M. D. (2012). Mapping the global flow of steel: From steelmaking to end-use goods. *Environmental Science & Technology, 46*(24), 13048–13055. doi:10.1021/es302433p

Key	Source	Anchors supported
[ICSG-ICA]	ICSG World Copper Factbook (2025), Retrieved from: ICSG website	Copper semi-fabricated production breakdown by product form; copper recycled input rate 28–35%
[MISO2-Ni][63]	MISO2 global nickel data, primary raw production flows (alloys-in-steel and non-alloy use)	Nickel primary raw production 1913–2022
[IRP-GMF]	UNEP International Resource Panel, Global Material Flows Database — global ore and commodity production series, converted to a primary metal equivalent basis. Retrieved from: Resource Panel website	Copper, cobalt and silicon historical supply base; nickel 2023 bridge; silicon 2024–2040 support path
[SUN-17]	Sun, X., Hao, H., Zhao, F., & Liu, Z. (2017). Tracing global lithium flow: A trade-linked material flow analysis. Resources, Conservation and Recycling, 124, 50-61.	Lithium end-use split 1994–2015 and battery share trajectory
[HEL-19]	Helbig, C., Bradshaw, A. M., Wietschel, L., Thorenz, A., & Tuma, A. (2018). Supply risks associated with lithium-ion battery materials. Journal of Cleaner Production, 172, 274-286.	Lithium battery share history; pre-battery cobalt stock estimate used in the hardmetal cross-check
[WB-2020]	World Bank. Minerals for Climate Action Data. Microsoft Excel workbook. Washington, DC: World Bank, 2020; Hund, Kirsten, Daniele La Porta, Thao P. Fabregas, Tim Laing, and John Drexhage. Minerals for Climate Action: The Mineral Intensity of the Clean Energy Transition. Washington, DC: World Bank, 2020	Global lithium battery end-of-life recycling below 1% pre-2020
[EU-2023/1542]	Regulation (EU) 2023/1542 concerning batteries and waste batteries[64]	Recovery targets: copper/ nickel /cobalt 90% (2027) rising to 95% (2031); Lithium 50% rising to 80%; underpins the authored battery collection trajectories
[JRC-MSA]	European Commission Joint Research Centre, Material System Analysis of five battery-related raw materials: Cobalt, Lithium, Manganese, Natural Graphite, Nickel in 2020 (EUR 30103 EN)	Stream-specific material system breakdowns for cobalt and nickel
[JORDAN-13]	Jordan, D.C., Kurtz, S.R. (2013), Progress in Photovoltaics 21(1), 12–29	Photovoltaic degradation rate 0.7%/yr; economic lifetime 25–30 y
[IRENA-16]	International Renewable Energy Agency and International Energy Agency Photovoltaic Power Systems Programme. End-of-Life Management: Solar Photovoltaic Panels. IEA-PVPS Report T12-06:2016. Abu Dhabi: International Renewable Energy Agency, 2016	Regular-loss (~30 y) and early-loss (~30% of modules below 25 y) retirement scenarios

[63] Wiedenhofer, D., Baumgart, A., Matej, S., Virág, D., Kalt, G., Lanau, M., Tanikawa, H., Sun, L., Cao, Z., Streeck, J., Wieland, H., Schug, F., Mayer, A., Wiedenhofer, N., Krausmann, F., & Haberl, H. *MISO2: Global material stocks and flows database — nickel primary raw production.* Retrieved from: https://onlinelibrary.wiley.com/doi/10.1111/jiec.13575 [Data set]

[64] EC, Regulation (EU) 2023/1542 of the European Parliament and of the Council of 12 July 2023 concerning batteries and waste batteries, amending Directive 2008/98/EC and Regulation (EU) 2019/1020 and repealing Directive 2006/66/EC, OJ L 191/1.Retrieved from: EUR-Lex website

Key	Source	Anchors supported
[FhISE-22]	Philipps, Simon, and Werner Warmuth. Photovoltaics Report. Freiburg: Fraunhofer Institute for Solar Energy Systems ISE, September 22, 2022.	25–30 y crystalline-silicon module field lifetime
[CIGRE-07][65]	CIGRE Working Group B1.07. Statistics of AC Underground Cables in Power Networks. Technical Brochure 338. Paris: CIGRE, 2007. Retrieved from: CIGRE website	Cable service life: impregnated-paper ~25 y, cross-linked polyethylene ~40 y
[REED-06][66]	Reed, R.C. (2006), The Superalloys: Fundamentals and Applications, Cambridge University Press	Cobalt superalloy turbine on-wing time evolution
[HUISMAN-12][6] [67]	Huisman, J. et al. (2012), Quantitative Assessment of WEEE in Europe, United Nations University / StEP	Electronics lifetime distributions for printed circuit boards and consumer electronics
[ALLABY-17][68]	Allaby, R. (2017), Stainless Steel in Architecture, Building & Construction, Nickel Institute	Stainless structural lifetime 40–60 y
[HELBIG-22][69]	Helbig, C.,et al . (2021), An overview of indicator choice and normalization in raw material supply risk assessments *Resources, 10*(8),; C Poncelet, et al. (2022), Losses and lifetimes of metals in the economy Nature Sustainability 5 — supplementary information and workbook	Consolidated registry underlying the per-stream MFA yield and collection parameters
[GULLEY-22]	Gulley, A.L. (2022), One hundred years of cobalt production from the Democratic Republic of the Congo (DRC), Resources Policy	1986 cobalt consumption peak of 43 kt
[DARTON-19]	Darton Commodities Limited. Cobalt Market Review 2019. Guildford, UK: Darton Commodities Limited, 2019.	Confirms 2019 cobalt apparent consumption of 122.6 kt
[ZENG-22]	Zeng et al. 2022, global cobalt in-use stock inventory by end-use stream (as compiled in the cobalt benchmark)	Cobalt in-use stock 2019: superalloys 66 kt, batteries 263 kt, carbides 52 kt, of ~470 kt total; stream mean lifetimes
[GODOY-20][70]	Godoy León, M.F., Dewulf, J. (2020), cobalt flows and recycling	Cobalt secondary share and end-of-life context

[65] CIGRE Working Group B1.07. (2007). *Statistics of AC underground cables in power networks* (Technical Brochure 338). CIGRE. See also IEC 60502 and IEEE Std 1407.
[66] Reed, R. C. (2006). *The superalloys: Fundamentals and applications*. Cambridge University Press.
[67] Huisman, J., van der Maesen, M., Eijsbouts, R. J. J., Wang, F., Baldé, C. P., & Wielenga, C. A. (2012). *The Dutch WEEE flows: 2011 update — Quantitative assessment of WEEE in Europe*. United Nations University / Solving the E-waste Problem (StEP) Initiative.
[68] Allaby, R. (2017). *Stainless steel in architecture, building and construction*. Nickel Institute.
[69] Charpentier Poncelet, A., Helbig, C., Loubet, P., Beylot, A., Muller, S., Villeneuve, J., Laratte, B., Thorenz, A., Tuma, A., & Sonnemann, G. (2022). Losses and lifetimes of metals in the economy. *Nature Sustainability, 5*, 717–726. doi:10.1038/s41893-022-00895-8 — supplementary information and workbook. Note: the model's own parameter documentation records this source under a conflated citation, giving the author list and title of a different paper (Helbig, C., Bruckler, M., Thorenz, A., & Tuma, A. (2021), An overview of indicator choice and normalization in raw material supply risk assessments, *Resources, 10*(8), 79) against this article's journal, year and DOI. The digital object identifier and the supplementary workbook filenames recorded there resolve to the article cited here, which is the one the parameters are drawn from.
[70] Godoy León, M. F., Blengini, G. A., & Dewulf, J. (2020). Cobalt in end-of-life products in the EU, where does it end up? The MaTrace approach. *Resources, Conservation and Recycling, 158*, 104842. doi:10.1016/j.resconrec.2020.104842

Key	Source	Anchors supported
[JRC-BATT-20][71]	JRC Material System Analysis — batteries (2020)	Cobalt battery-stream recovery context
[DELOITTE-15][72]	BIO by Deloitte France (2015), study on cobalt flows	Pre-2015 cobalt end-use and recovery shares
[NRCAN-24][73]	Natural Resources Canada (2024), cobalt facts	2020 cobalt end-use split corroboration
[SFA-21][74]	SFA (Oxford) (2021), cobalt market analysis	2020 cobalt end-use split corroboration
[USGS-DS140][75]	USGS, Data Series 140 — Historical Statistics for Mineral and Material Commodities (per-commodity workbooks: copper from 1850, nickel from ~1900, cobalt from ~1907, lithium from 1900, silicon from 1932)	Long-run production backbone for the pre-1970 segments of each material; nickel Sudbury-era and pre-1950 anchors; lithium historical production; silicon world production 1964–2019
[USGS-DS896][76]	USGS, Data Series 896 — Historical Global Statistics for Mineral and Material Commodities	Global counterpart to DS-140 where the national series alone is insufficient
[USGS-C1196Z][77]	USGS Circular 1196-Z (2009), Nickel Recycling in the United States in 2004	Nickel old-scrap recycling efficiency of 56.2%; the functional end-of-life recycling convention adopted throughout
[CLIO-NI][78]	Clio Infra (Utrecht University), nickel production dataset, compiled from BGS and USGS returns	Country-level nickel production from 1950; corroborates the Sudbury-era consumption anchors
[ICSG-FB-20][79]	ICSG, The World Copper Factbook 2020, and the ICSG Scrap Survey (2019)	Copper mine production 1960–2019; new-scrap 5.2 Mt and old-scrap 3.3 Mt split; SX-EW share of mine output
[GORDON-06][80]	Gordon, R.B., Bertram, M., Graedel, T.E. (2006), PNAS 103(5), 1209–1214; and Kapur, A., Graedel, T.E. (2006), Environ. Sci. Technol. 40, 3135–3141	Global copper above-ground stock of ~350 Mt at 2000, the earlier of the two copper stock anchors
[SVERDRUP-16][81]	Sverdrup, H.U. (2016), Resources, Conservation & Recycling 114, 112–129	Modelled historical lithium production 1900–2016, used to support the pre-1970 lithium reconstruction

[71] European Commission, Joint Research Centre. (2020). *Material system analysis: Batteries*. Publications Office of the European Union.
[72] BIO by Deloitte. (2015). *Study on data for a raw material system analysis: Cobalt flows*. European Commission. Retrieved from: European Commission website
[73] Natural Resources Canada. (2024). *Cobalt facts*. Government of Canada.
[74] SFA (Oxford). (2021). *Cobalt market analysis*. SFA (Oxford) Ltd.
[75] Kelly, Thomas D., Grecia R. Matos, David A. Buckingham, Carl A. DiFrancesco, Kenneth E. Porter, Cyrus Berry, Melissa Crane, Thomas Goonan, and John Sznopek. Historical Statistics for Mineral and Material Commodities in the United States. U.S. Geological Survey Data Series 140. Reston, VA: U.S. Geological Survey. https://doi.org/10.3133/ds140.
[76] U.S. Geological Survey. *Historical global statistics for mineral and material commodities* (Data Series 896). U.S. Department of the Interior.
[77] Goonan, Thomas G. "Nickel Recycling in the United States in 2004." In Flow Studies for Recycling Metal Commodities in the United States, edited by Scott F. Sibley, Z1–Z30. U.S. Geological Survey Circular 1196-Z. Reston, VA: U.S. Geological Survey, 2009. https://doi.org/10.3133/cir1196Z.
[78] Klein Goldewijk, Kees, and Jonathan Fink-Jensen. *Nickel Production*. Version 1.1. Clio Infra, IISH Data Collection, 2015. hdl:10622/ZGRJQY.
[79] International Copper Study Group. *The World Copper Factbook 2020.* Lisbon: International Copper Study Group, 2020.
[80] Gordon, R. B., Bertram, M., & Graedel, T. E. (2006). Metal stocks and sustainability. *Proceedings of the National Academy of Sciences, 103*(5), 1209–1214. doi:10.1073/pnas.0509498103; and Kapur, A., & Graedel, T. E. (2006). Copper mines above and below the ground. *Environmental Science & Technology, 40*(10), 3135–3141. doi:10.1021/es0626887
[81] Sverdrup, Harald Ulrik. "Modelling Global Extraction, Supply, Price and Depletion of the Extractable Geological Resources with the LITHIUM Model." Resources, Conservation and Recycling 114 (2016): 112–129. https://doi.org/10.1016/j.resconrec.2016.07.002.

Key	Source	Anchors supported
[SILICONES-EU][82]	Silicones Europe, socio-economic evaluation of silicones (executive summary); NASA technical report 19800006234 on silicon material production	Silicon end-use shares underlying the pre-1970 proxy reconstruction: metallic silicon to aluminium alloys 40–50%, and the historical organosilicon and polysilicon helper series
[TAKIGUCHI-11]	Takiguchi, H., Morita, K. (2009), J. Japan Inst. Metals; Takiguchi, H. (2011), Global Flow Analysis of Crystalline Silicon, IntechOpen, doi:10.5772/23573; Williams, E. (2003). Forecasting material and economic flows in the global production chain for silicon. Technological Forecasting and Social Change, 70, 341–357. https://doi.org/10.1016/S0040-1625(02)00201-9	Silicon stream structure and process yields: metallurgical-grade input per unit polysilicon, wafer kerf loss, silicone process losses
[BERNREUTER][83]	Bernreuter Research, Polysilicon Market Analysis (2024); Freedonia Group, Global Silicones (2024); Wacker Chemie AG annual reports	Polysilicon output and capacity by year; silicone production volumes underpinning the silicon application split
[MISO2-GAS][84] [85]	MISO2 global material stocks and flows database, gross-addition-to-stock end-use series for iron and steel (1870–2022) and aluminium (1890–2022), doi:10.5281/zenodo.12794253	Normalised activity indices driving the pre-1970 backcast of ferrosilicon and of silicon in aluminium alloys; downstream sector mapping for both host materials
[USGS-SI-USE][86]	USGS, Data Series 140 annual silicon end-use worksheet (1975–2003)	Aluminium-alloy and miscellaneous end-use indices normalised to 2003 = 1.0, used to cross-check the aluminium backcast; miscellaneous covers silicones, silanes, fumed silica and, from 1991, aluminium alloys
[NASA-JPL-80][87]	NASA / Jet Propulsion Laboratory, low-cost solar array project report (NTRS 19800006234), Table 4.3a median forecast N8716395	Observed solar share within polysilicon 1977–1985, rising from 0.7% to 18.8%
[SCRREEN2]	SCRREEN2 critical raw materials factsheet, silicon metal (2023 update), Table 2	Late-2010s global silicon-metal end-use validation bounds: chemical applications 54%, aluminium alloys 38%, solar 6%, electronics 2%

[82] NASA, (1979). Silicon Materials Outlook Study for 1980-85, Photovoltaics Program Distribution Category Technology Development and Applications Lead Center, Retrieved from: NASA website

[83] Bernreuter Research. Polysilicon Market Analysis 2024, www.scribd.com/document/963875678/polysilicon-market-outlook-2024; Würzburg: Bernreuter Research, 2024; The Freedonia Group. Global Silicones. Cleveland, OH: Freedonia Group, 2024, Bemreuter Research website; Wacker Chemie AG. Annual Report, 2025., Retrieved from: Wacker website.

[84] Wiedenhofer, Dominik, Jan Streeck, Hanspeter Wieland, Benedikt Grammer, André Baumgart, Barbara Plank, Christoph Helbig, Stefan Pauliuk, Helmut Haberl, and Fridolin Krausmann. "From Extraction to End-Uses and Waste Management: Modeling Economy-Wide Material Cycles and Stock Dynamics around the World." *Journal of Industrial Ecology* (2024): 1–17. https://doi.org/10.1111/jiec.13575

[85] Wiedenhofer, D., J. Streeck, H. Wieland, B. Grammer, A. Baumgart, and B. Plank. "The MAT_STOCKS Database: Economy-wide Material Flows and Material Stock Dynamics Around the World". In Journal of Industrial Ecology, version 1.0.0, vol. 28. no. 6. Zenodo, July 22, 2024. https://doi.org/10.5281/zenodo.12794253.

[86] U.S. Geological Survey. "Silicon." In Historical Statistics for Mineral and Material Commodities in the United States. Data Series 140. Microsoft Excel workbook. Reston, VA: U.S. Geological Survey, [workbook update year]. https://doi.org/10.3133/ds140.

[87] P. D. Maycock. (1987). Low-Cost Solar Array Project (1974-1986), The Jet Propulsion Laboratory, Photovoltaic Energy Systems, Inc. Retrieved from: NASA website.

Key	Source	Anchors supported
[BRGM-19]	BRGM, silicon metal criticality assessment (2019, public deck)	2018 world silicon-metal consumption validation bounds: silicones and silanes 35%, aluminium alloys with iron and steel 41%, solar cells 18%, other 6%
[PAULIUK-13][88]	Pauliuk, S., Wang, T., Müller, D.B. (2013), Resources, Conservation & Recycling 71, 22–30; Wang, T., Müller, D.B., Graedel, T.E. (2007), Environ. Sci. Technol. 41(14), 5120–5129	Global in-use steel stock, the host-stock basis for estimating silicon held in steel
[LIU-12]	Liu, G., Bangs, C.E., Müller, D.B. (2012), Nature Climate Change 3, 338–342; Løvik, A.N., Modaresi, R., Müller, D.B. (2014), Environ. Sci. Technol. 48(8), 4257–4265	Global in-use aluminium stock and alloying-element recovery, the basis for silicon-in-aluminium carry-over and its downcycling treatment
[PINDYCK-94][89]	Pindyck, R.S. (1994), on inventories and the short-run dynamics of commodity prices; Boer, L., Pescatori, A., Stuermer, M. (2021), IMF working paper on energy-transition metals	Global price-scarcity sensitivity anchor of 0.55 for deep exchange-traded markets
[DAHL-20][90]	Dahl, C.A. (2020), Mineral and Energy Demand Studies database; Stuermer, M. (2017) and Fally, T. & Sayre, J. (2018) on long-run mineral commodity demand	Short-run demand price elasticity of 0.10 and long-run of 0.30
[STERMAN-89][91]	Sterman, J.D. (1989), on misperceptions of feedback in dynamic decision making, Management Science	Investor perception lag of two to three years in the capacity expectation channel
[SCRAP-ELAST][92]	Fisher, F.M., Cootner, P.H., Baily, M.N. (1972), econometric model of the world copper industry; Slade, M.E. (1980) and Fu, X. et al. (2017) on secondary-metal supply response	Scrap-supply price elasticity band of 0.10–0.29 bounding the collection-effort build and decay gains

Source-quality caveats. Reading the benchmark compilations against their own citation lists surfaces three limitations worth stating explicitly. First, the compilations draw on a mix of source tiers: alongside government statistical series and peer-reviewed papers, they cite industry newsletters, consultancy market reports, and general-reference encyclopaedia entries. The resulting anchors are generally supported across multiple sources. However, the underlying evidence base is not consistently peer-reviewed and should therefore not be presented as if it were. Second, portions of the silicon end-use split are derived by the modelling team through constrained mass-balance allocation rather than being sourced directly from external datasets. This is a legitimate approach that was necessary because published silicon end-use data do not provide the level of

[88] Pauliuk, S., Wang, T., & Müller, D. B. (2013). Steel all over the world: Estimating in-use stocks of iron for 200 countries. *Resources, Conservation and Recycling, 71*, 22–30. doi:10.1016/j.resconrec.2012.11.008; and Wang, T., Müller, D. B., & Graedel, T. E. (2007). Forging the anthropogenic iron cycle. *Environmental Science & Technology, 41*(14), 5120–5129. doi.10.1021/es062761t

[89] Pindyck, R. S. (1994). Inventories and the short-run dynamics of commodity prices. *RAND Journal of Economics, 25*(1), 141–159; and Boer, L., Pescatori, A., & Stuermer, M. (2021). *Energy transition metals* (IMF Working Paper No. 2021/243). International Monetary Fund.

[90] Dahl, C. A. (2020). *Mineral and energy demand studies database*. Colorado School of Mines; with Stuermer, M. (2017), *Journal of the European Economic Association*, and Fally, T., & Sayre, J. (2018), *NBER Working Paper 24434*, on long-run mineral commodity demand.

[91] Sterman, J. D. (1989). Modeling managerial behavior: Misperceptions of feedback in a dynamic decision making experiment. *Management Science, 35*(3), 321–339. doi:10.1287/mnsc.35.3.321

[92] Fisher, F. M., Cootner, P. H., & Baily, M. N. (1972). An econometric model of the world copper industry. *Bell Journal of Economics and Management Science, 3*(2), 568–609; with Slade, M. E. (1980) and Fu, X., et al. (2017) on secondary-metal supply response.

stream-specific disaggregation required by the model but is one that makes silicon composition an authored construct rather than a sourced observation. Third, the cobalt in-use stock benchmark used to justify the superalloy calibration comes from a different study than the end-use shares it overrides; the calibration therefore reconciles two independent sources against each other rather than correcting an internal inconsistency in one.

Table 34. *Source-to-artefact binding: which sources support each input artefact.*

Artifact	Primary sources	Secondary / validation sources
final_demand	[IEA-GCMO-24], [USGS-GC-Cu], [USGS-MCS], [USGS-MYB-Co], [INSG-FB-24], [IMF-T71]	[USGS-DS140], [CLIO-NI], [SVERDRUP-16], [GULLEY-22], [DARTON-19], [SILICONES-EU], [BERNREUTER]
primary_ore_input	[IEA-GCMO-25], [MISO2-Ni], [USGS-MCS], [IRP-GMF]	[USGS-DS140], [USGS-MYB-Co], [BERNREUTER]
route_yield	[IRP-GMF] (commodity-to-ore ratios)	None — no provenance record exists for this artefact
route_stream_share	[CI-2024], [INSG-FB-24], [ICSG-ICA], [SUN-17], [IEA-GCMO-25], [HARPER-12]	[USGS-MCS], [USGS-MYB-Co], [USGS-DS140], [RECK-08], [ZENG-22], [TAKIGUCHI-11]
collection_routing_rates	[GLOSER-13], [RECK-08], [CI-2024], [WB-2020], [USGS-MCS], [EU-2023/1542]	[ICSG-ICA], [IEA-GCMO-25], [JRC-MSA], [USGS-C1196Z], [LIU-12]
lifetime_distributions	[CULLEN-17], [ALLABY-17], [CIGRE-05], [REED-06], , [JORDAN-13], [IRENA-16]	[FhISE-22], [IEA-PVPS-12], [HUISMAN-12], [IEA-GCMO-24], [SUN-17], [USGS-MCS]
Per-stream MFA parameters	[HELBIG-22] registry and the primary sources it traces to	[GLOSER-13], [RECK-08], [ICSG-ICA], [EU-2023/1542], [IEA-PVPS-12]

Provenance records. Five of the six driver artefacts carry a machine-readable provenance record alongside the data file itself, recording the sources bound to each anchor, per-value confidence labels, and the parameters of the construction. Confidence is recorded at the level of individual material-stream-pathway combinations rather than as a single label per file, which is what makes the uneven evidence base traceable rather than merely acknowledged — the difference between an evidence-based recycling anchor and an assumption-based remanufacturing anchor in the same row is visible in the record. The route-yield artefact has no separate provenance record; its lineage is documented in prose only, which is a known gap in the artefact set. Table 35 gives the distribution across the sixty stream-pathway cells of the routing artefact, which is the artefact where the evidence base is most uneven.

Table 35. *Confidence label distribution across the 60 stream × pathway cells of the collection-routing artefact.*

Pathway	Evidence-based (EB)	Prospective (P)	Low-evidence (L)	Interpretation
recycling	19 of 20	0	1 of 20	Nearly all anchored on quantitative literature values
downcycling	18 of 20	0	2 of 20	Evidence-based except photovoltaic and semiconductor carry-over

Pathway	Evidence-based (EB)	Prospective (P)	Low-evidence (L)	Interpretation
remanufacturing	4 of 20	4 of 20	12 of 20	Assumption-based parameters constrained by physically plausible ranges, reflecting the absence of published benchmarks for remanufacturing rates at the material-stream level.
All pathways	41 of 60	4 of 60	15 of 60	The evidence gap is concentrated almost entirely in remanufacturing

Validation gates. Three classes of automated checks run at artefact build time, all enforcing rather than advisory: a failure aborts the build. **Closure checks** verify that each share vector sums to one within tight numerical tolerance, for each year and each scenario: stream shares close to within 1e-12, routing rates within 1e-9. **Activation checks** verify that each stream's share and collection rate are exactly zero before its activation year and that displaced mass reaches the designated fallback. Two streams carry different activation years in the composition and routing artefacts — copper electronic grade at 1920 against 1900, and nickel class I refined metal at 1870 against 1880 — because the two encode different criteria: the year a stream's share first becomes non-zero, against the year its collection becomes physically meaningful. The difference is deliberate, but it means Table 27's activation years apply to composition and do not transfer to the routing artefact for those two streams. **Benchmark reconciliation checks** recompute stream-share-weighted effective rates and compare them against the literature anchors in Table 29, plus twelve equivalent gates on stream composition. Separately, the twenty material-stream cells are asserted identically across the three stream-keyed artefacts, so a stream added to one file but not another cannot pass silently.

Data status and known gaps. Four limitations in the input data are load-bearing for how results should be read, over and above the per-artefact caveats already given. First, the source and route allocation companions — deposit share, ore grade, route share, and product share — are present and structurally live but carry placeholder values in the active configuration: deposit type and processing route both collapse to a single canonical value with unit share and unit grade, and the product split is a single canonical category at unit share. This is a stated temporary simplification, not an inactive allocation layer, but it means the model cannot currently distinguish sulfide from laterite nickel or brine from spodumene lithium even though the schema supports it. The stream mapping is not affected: route-stream shares are differentiated and benchmark-anchored, as set out above. Second, all datasets are global: regional variation in stream mix, collection infrastructure and processing route is not represented, which can be important because stainless composition, cobalt processing and battery manufacturing are each strongly region-concentrated. Third, the generator-parity gap described above means two of the six driver artefacts cannot be regenerated from the shipped scripts. Fourth, the post-2050 portion of each forward series is an authored tail assumption, not a sourced projection, and its geometry is in one respect counter-intuitive: cobalt's scenario fan partially reconverges rather than widening to 2100, with the moderate scenario exceeding the most ambitious by the end of the century. Post-2050 fan geometry is illustrative rather than a sourced result.

Parameters and variables

Where the exogenous artefacts supply quantities, the parameter surface governs how the model transforms them. Parameters are held in layered configuration files rather than in the data store, and resolve through an explicit precedence chain: a shared authored baseline supplies fallback values for each parameter; per-material or per-stream slices narrow those to specific cohorts; and the run entrypoint may override either. Where a parameter also has an exogenous time series — collection rate being the important case — the CSV takes precedence over the parameter value entirely, and the parameter surface exposes a separate multiplier for calibration instead. Reading a parameter's active value therefore requires tracing this chain rather than reading the first file in which the name appears.

The executable parameter schema. The active contract fixes a named *parameter schema* that acts as an executable allowlist: it enumerates exactly which parameters the MFA layer will accept, and a configuration setting any parameter outside that list is rejected rather than silently ignored. The schema distinguishes two nested sets. The *active* set is each parameter the model reads. The *calibratable* set is the strictly smaller subset that an automated calibration routine is permitted to adjust — a deliberate restriction that prevents calibration from tuning parameters whose values encode a policy choice, a scenario lever, or a physical constant rather than an uncertain empirical quantity. This is a structural allowlist, not a record of calibration actually being run: as noted above, the active configuration uses literature-anchored values throughout rather than fitted ones (see Calibration data). Table 36 gives both sets in full.

Table 36. *MFA parameter schema: the 27 active parameters, of which 7 are exposed to calibration.*

Parameter	Role	Calibratable
semis_fabrication_yield	Yield of refined metal into semi-finished products	yes
finished_product_manufacturing_yield	Yield of semis into finished products	yes
sorting_yield	Recovery rate at the end-of-life sorting stage	yes
recycling_yield	Metal-to-metal process yield of the functional recycling loop	yes
reman_yield	Fraction of collected product eligible for direct remanufacture	yes
product_lifetime_multiplier	Scalar on the baseline lifetime distribution	yes
collection_rate_multiplier	Proportional scaling of the exogenous collection trajectory	yes
collection_rate	Aggregate fallback only; the exogenous CSV supersedes it per stream	no — superseded by CSV
recycling_rate / remanufacturing_rate / downcycling_rate	Gate-2 conditional routing shares; supplied per stream and year by the exogenous routing artefact	no — supplied by CSV
product_lifetime_multiplier_factor	Scenario lifetime-extension lever, composing multiplicatively	no — scenario lever

Parameter	Role	Calibratable
new_scrap_recovery_rate	Share of fabrication and manufacturing scrap recovered to refinery feed	no
refinery_stockpile_release_rate	Fraction of the refinery stockpile released per period	no
primary_supply_elasticity	Price responsiveness of the primary supply cap	no
warm_start	Historical burn-in enablement, length and rule	no
strategic_reserve_enabled	Master switch for the strategic reserve policy	no — policy
strategic_reserve_target_coverage_years	Reserve size target, expressed as years of demand cover	no — policy
strategic_reserve_fill_gain / _release_gain	Proportional response rates for filling and releasing	no — policy
strategic_reserve_max_fill_rate / _max_release_rate	Per-period caps on reserve build and draw-down	no — policy
strategic_reserve_fill_price_threshold / _release_price_threshold	Price ratios gating fill and release	no — policy
strategic_reserve_fill_service_threshold / _release_service_threshold	Service-stress levels gating fill and release	no — policy
mass_balance_tolerance	Numerical tolerance for the mass-balance closure check	no — solver

The exclusion of collection rate from the calibratable set is worth drawing out, because it is a structural consequence of the two-gate schema rather than a stylistic preference. Collection rate is supplied per stream and per year by the exogenous routing artefact, carrying real structural detail — activation years, scenario-dependent trajectories, stream-specific ceilings. Enabling calibration to overwrite it with a single fitted scalar would destroy that structure. The multiplier scales the whole authored trajectory proportionally instead, so calibration can adjust the level while preserving the shape and the per-stream heterogeneity underneath. The same logic applies to the lifetime-extension lever, which composes multiplicatively onto per-stream calibrated lifetime multipliers rather than replacing them: a scenario that raises product lifetimes must not silently discard the calibration.

Per-stream MFA parameter values. Nineteen slices narrow the global fallbacks to specific material-stream cohorts. The evidence base is a consolidated registry built from a published supply-risk assessment supplement and the primary sources it traces back to, supplemented where current evidence has clearly moved beyond the registry's historical window — battery collection rates in particular were raised above their historical anchors to reflect regulatory mandates and rapid recycling capacity build-out. Yields are decomposed into a semis-fabrication step and a finished-product-manufacturing step whose product should reconcile with the registry's combined fabrication-and-manufacturing figure. Table 37 gives the deployed values.

Table 37. *Per-stream MFA parameter values in the deployed configuration. Blank entries inherit the global fallback.*

Material / stream	recycling_yield	reman_yield	semis_fab_yield	finished_mfg_yield
copper / wire_rod_grade	0.93	0.20	0.96	0.88
copper / alloy_grade	0.96	0.30	0.93	0.82
copper / tube_sheet_plate_grade	0.87	0.30	0.93	0.82
copper / electronic_grade	0.825	0.05	0.88	0.72
nickel / class_i_refined_metal	0.90	0.08	0.91	0.86
nickel / class_ii_stainless_feed	0.98	0.20	0.97	0.94
nickel / battery_grade_sulfate	0.92	0.04	0.90	0.85
nickel / chemicals_catalysts	0.90	0.08	0.91	0.86
cobalt / battery_grade_sulfate	0.88	0.03	0.94	0.90
cobalt / superalloy + hardmetal	0.90	0.15	0.94	0.92
cobalt / magnets_grade_metal	0.85	0.08	0.93	0.90
cobalt / chemicals_pigments_catalysts	0.80	0.02	0.87	0.85
lithium / battery_grade	0.85	0.04	0.95	0.95
lithium / industrial_grade	0.80	0.02	time series	1.00
lithium / specialty_organolithium_metal	0.80	0.02	time series	1.00
silicon / ferroalloy_grade	0.85	0.03	0.95	1.00
silicon / chemical_alloying_grade	0.85	0.03	0.938	1.00
silicon / solar_grade_polysilicon	0.85	0.05	0.60	0.95
silicon / semiconductor_grade_polysilicon	0.20	0.02	0.50	0.85

Evidence base for the MFA parameter surface. Each per-stream yield and recovery value traces to a single consolidated evidence registry built from the supplementary material of a published supply-risk assessment *[HELBIG-22]*, which the registry records at page level: lithium parameters from page 3 of the supplement, silicon from page 4, cobalt from page 6, and nickel and copper from page 7. The registry holds 281 normalised data points, each traced to a page in the supplement or a row in its workbook, with cross-check sheets verifying that end-use shares sum to one and that production yields and end-of-life rates reconcile against independent anchors. Table 38 gives the registry's metal-level values alongside the deployed stream-level values, so the transformation between them is visible.

Table 38. *Metal-level evidence-registry anchors against deployed stream-level parameters. Registry values are from [HELBIG-22] at the pages indicated.*

Metal (registry page)	Registry anchor	Registry value	Deployed stream-level treatment
copper (p. 7)	Production yield / fab+manuf yield / new scrap recovery / remelting yield / EoL-RR	0.835 / 0.824 / 0.923 / 1.00 / 0.48	Fab+manuf decomposed into semis × finished stages, 0.96×0.88 (wire rod) down to 0.88×0.72 (electronic); new scrap recovery 0.923 applied at material level; remelting becomes stream recycling_yield 0.825–0.96
nickel (p. 7)	Production yield / fab+manuf / new scrap recovery / remelting / average collection / EoL-RR	0.788 / 0.86 / 0.84 / 1.00 / 0.63 / 0.65	New scrap recovery 0.844 applied at material level; recycling_yield 0.90–0.98 by stream; sector collection yields (construction 0.87, batteries 0.29, electronics 0.29, chemicals 0) inform the routing CSV rather than the parameter surface
cobalt (p. 6)	Production yield / fab+manuf / new scrap recovery / remelting / collection by application	0.44 / 0.94 / 0.05 / 0.95 / superalloy 0.80, carbide 0.50, CE battery 0.35, EV battery 0.90, magnet 0.10	New scrap recovery 0.047 at material level; recycling_yield 0.80–0.90 by stream; application collection yields inform routing anchors
lithium (p. 3)	Battery fab yield / other-application fab yield / overall fab+manuf / new scrap recovery / LIB collection / LIB EoL recycling	0.80 / 0.90 / 0.84 / 0 / 0.06 / 0.05	Battery stream 0.95×0.95; industrial and specialty streams carry a time-varying semis yield; new scrap recovery 0.50 for battery grade, 0 elsewhere
silicon (p. 4)	Production yield from quartz / fab+manuf / new scrap recovery / remelting / collection	0.80 / 0.90 / 0 / 0 / 0	Fab yields fall far below the registry figure for polysilicon (0.60 solar, 0.50 semiconductor) to represent ingot and wafering losses; new scrap recovery 0 retained; recycling_yield set to a non-zero process-only value with cross-circuit loss moved to downcycling

Three systematic departures from the registry. The deployed surface is not a direct transcription of the registry; the differences are systematic. **First, decomposition.** The registry reports a single combined fabrication-and-manufacturing yield per metal; the model splits this into a semis stage and a finished-product stage, and differentiates both by stream. Weighted by stream shares, the product of the two should reconcile with the registry's metal-level figure — the stated cross-check, though it is not an automated gate. **Second, re-anchoring of recycling yield.** The registry's remelting yield bundles cross-circuit mass loss together with metallurgical loss. Because cross-circuit loss is now accounted for separately through the downcycling pathway, the deployed recycling yields were re-anchored to the metal-to-metal step alone and are correspondingly higher than the values they replaced — silicon's dispersed streams moved from 0.40 to 0.85 and solar polysilicon from 0.60 to 0.85. Quoting these against remelting yields without accounting for that boundary change would suggest an unexplained improvement in recycling efficiency that has not occurred. Table 37 gives the deployed values, and they are the ones the model executes. **Third, sector-to-stream mapping.** The registry reports collection yields by end-use sector (construction, transport, electronics, chemicals); the model needs them by material stream. The mapping is the modelling team's own weighting of sectors onto streams, documented per slice but not independently sourced. Collection rate is not a

parameter of this surface: it comes per stream and year from the routing artefact, so the registry's collection yields inform that artefact's anchors rather than this surface.

Two patterns in Table 37 carry physical meaning worth stating. The silicon polysilicon streams have the lowest fabrication yields in the model, 0.60 for solar-grade silicon and 0.50 for semiconductor-grade silicon, and reflect the material-intensive nature of ingot pulling and wafer production. Diamond-wire kerf losses alone account for roughly 40% of the ingot, while Czochralski crystal growth and wafering combined can result in overall yields of approximately 50% for semiconductor-grade silicon. These are not inefficiencies that could plausibly be engineered away within the model horizon; they are close to the process physics. Separately, new-scrap recovery differs by nearly two orders of magnitude across materials — 0.923 for copper and 0.844 for nickel, against 0.047 for cobalt and zero for silicon and non-battery lithium — reflecting that copper and nickel fabrication scrap is a routinely traded, economically valuable feedstock while cobalt and silicon process residues largely are not. Recycling yields here are pure metallurgical process yields: the cross-circuit loss that was previously folded into them has been moved out to the downcycling channel in the routing artefact, so these values are not comparable with recycling yields quoted on the older convention.

System Dynamics parameter surface. The behavioural layer carries a separate parameter surface governing price formation, demand response, capacity adjustment, collection effort and background variability. All are calibrated gains, lags and elasticities rather than structural parameters of an explicit decision model, consistent with the reduced-form design described in Model Architecture. Their evidence base is of a different character from the MFA surface: where MFA parameters derive from a single consolidated materials registry, SD parameters are anchored on a mix of econometric literature and specific historical price episodes, with the global values anchored on copper and nickel as the well-documented markets and per-material slices encoding divergence from that anchor. Table 39 gives the global surface with its documented basis; Table 40 gives the historical episodes each material's slice is calibrated against.

Table 39. *Global System Dynamics parameter surface: deployed value, what it governs, and documented basis. Global values are copper/nickel-anchored.*

Parameter	Value	Governs	Documented basis
price_base	1.0	Price index normalisation	Definitional
price_scarcity_sensitivity	0.55	Strength of the scarcity-to-price channel	Copper/nickel anchor for deep exchange-traded markets; commodity price-formation literature: Pindyck, R. S. (1994). Inventories and the short-run dynamics of commodity prices. The RAND Journal of Economics, 25(1), 141–159; Boer, L., Pescatori, A., & Stuermer, M. (2021). Energy Transition Metals (IMF Working Paper No. 2021/243). International Monetary Fund. DOI: 10.5089/9781513599373.001
demand_price_elasticity	0.10	Short-run demand response to price	Global short-run anchor from a published elasticity survey (Dahl 2020)

Parameter	Value	Governs	Documented basis
demand_price_elasticity_lr	0.30	Long-run demand response	Set at 3× the short-run value on long-run commodity demand evidence (Stuermer 2017; Fally & Sayre 2018)
lr_ramp_years	10.0	Phase-in period for long-run elasticity	One decade, as the time constant of the short-to-long-run transition
capacity_envelope_initial	1.0	Starting value of the capacity envelope	Definitional: capacity begins at parity with demand
capacity_envelope_min	0.80	Floor on the capacity envelope	Observed world copper mine utilisation of 81–85% (ICSG 2020/2022)
capacity_envelope_max	1.55	Ceiling on the capacity envelope	Raised from 1.45 for prospective headroom, on IEA guidance that copper and nickel supply expand ~50% by 2035 [IEA-GCMO-25]
capacity_expansion_gain	0.26	Rate of capacity expansion under pressure	Full-model calibration against historical price episodes; per-material values in Table 40
capacity_retirement_gain	0.04	Rate of retirement under idle conditions	Set asymmetrically far below expansion gain because mines retire slowly
retirement_slack_weight	1.5	Weight on the idle-capacity retirement arm	Raised from 0.5 to produce a copper capacity cycle of roughly ten years
expectation_horizon_years	2.0	Investor perception lag on demand	Empirical anchor of 2–3 years from the system-dynamics literature on investment behaviour (Sterman 1989)
backlog_clearance_rate	0.30	Unmet-demand backlog cleared per year	Implies an average three-year drawdown of accumulated backlog
capacity_adjustment_lag_years	4.5	Decision-to-availability lag	Weighted-average fallback across materials; per-material values in Table 40
capacity_pressure_shortage_weight	0.78	Shortage vs price weight in capacity target	Calibration outcome; weights physical shortage above price signal
bottleneck_scarcity_gain	0.15	Amplification of scarcity by bottleneck pressure	Copper/nickel anchor; raised substantially for thin markets
bottleneck_collection_sensitivity	0.08	Collection throttling under bottleneck	For the purposes of this research, the assumption is derived from underlying physical parameters and system constraints rather than from externally anchored empirical data
collection_effort_build_gain	0.0625	Build-up rate of collection effort	Implies a build/decay ratio of 0.25, mid-range of the 0.10–0.29 scrap-supply elasticity band (Fisher 1972; Slade 1980; Fu et al. 2017)

Parameter	Value	Governs	Documented basis
collection_effort_decay_gain	0.25	Decay rate of collection effort	Set with build gain to hold the elasticity ratio above
collection_multiplier_min / _max	0.88 / 1.10	Bounds on realised collection response	Copper/nickel anchor; policy-driven uplifts applied per material
noise_ar1_rho	0.6	Autocorrelation of background noise	Persistence of industrial demand deviations
noise_sigma_demand / _collection	0.04 / 0.02	Noise amplitude on demand and collection	4% and 2% annual standard deviation; thin markets override demand noise upward
noise_seed	42	Reproducibility seed	Arbitrary fixed seed; varied across runs for ensemble comparison

Calibration against historical price episodes. The per-material SD slices are not fitted to a time series. Each is tuned so that the model, driven by its own physical service-stress signal, reproduces the approximate magnitude of one documented historical price episode for that material — a targeted plausibility calibration rather than a statistical fit. Table 40 records the episode used to derive each material's parameters, together with the target and the achieved response. This provides transparency into the SD layer by documenting the basis for parameter selection and enabling readers to understand how performance was evaluated.

Table 40. *Historical price episodes each material's SD slice is calibrated against.*

Material	Calibration episode	Observed price response	Model target and outcome
copper	London Metal Exchange structural shortage, 2020–2021[93]	Roughly 2× price increase	Service stress ≈ 0.25 with stress gain 7.0 gives scarcity multiplier ≈ 2.7 and price ratio ≈ 2.0; utilisation trough ≈ 0.81, consistent with observed 81–85%
nickel	London Metal Exchange short squeeze, March 2022[94]	Roughly 5× intraday	Not reproducible from physical flows: modelled service stress is ≈ 0.02, so even the maximum feasible stress gain of 20 yields only ≈ 1.3×. The multiplier cap is set to 6.0 to admit the historical magnitude, but the fundamental cycle is driven by a demand-growth bottleneck instead
cobalt	Democratic Republic of the Congo(DRC) supply concentration	Roughly 3.5× price increase	Service stress ≈ 0.195 with stress gain 7.0 gives price ratio ≈ 3.5×; utilisation trough ≈ 0.80; recovery from overshoot in ≈ 4 years, matching an observed 4–6 year cobalt cycle

[93] A Home. (2021). A year of walking on the wild side for industrial metals. Retrieved from: Reuters; A Home. (2021). Falling stocks pose problems for London Metal Exchange. Retrieved from: Reuters.
[94] Eco3. (2022). Why Nickel Is the Most Fragile Base Metal: The 2022 LME Squeeze as Structural Evidence. Retrieved from: Eco 3; Xinpeng Sun. (2026). Looking Back at the LME Nickel Crisis (Part I): A Crisis the Exchange Could Barely Intervene In. Retrieved from: SSRN.

Material	Calibration episode	Observed price response	Model target and outcome
	episode, 2017–2018[95]		
lithium	Lithium carbonate equivalent peak, 2022[96]	5–7× sustained (excluding speculative overshoot); ~10× at peak	Service stress ≈ 0.49 with stress gain 7.0 gives scarcity multiplier ≈ 4.4 and price ≈ 6×; multiplier cap raised to 12 (15 for battery grade) to admit the historical peak
silicon	Solar-grade polysilicon crisis, 2021–2022 (Chinese energy rationing)[97]	Roughly 5×, from about $6 to $32/kg; 2008 spike 8–10×	Service stress ≈ 0.19 with stress gain 22 gives ≈ 4–5× for solar grade; utilisation trough ≈ 0.74; recovery in ≈ 5 years, matching post-glut Chinese capacity rationalisation

The March 2022 nickel episode was principally an LME short squeeze involving concentrated exchange-traded and OTC short positions, margin pressure and impaired liquidity, rather than a price movement commensurate with underlying physical-market scarcity. Consistent with this distinction, the model's physical-flow layer generates service stress of only approximately 0.02 for the period and does not reproduce the exchange-driven price spike.[98] The configuration therefore does not attempt to reproduce it: the parameters are set to the maximum that remains defensible against the physical signal, and the residual gap is recorded rather than closed by raising gains further. Any reading of nickel price dynamics from this model should be understood as representing a slow fundamental supply-demand cycle, not exchange-driven price behaviour.

Capacity-adjustment lags are among the most empirically grounded parameters in the behavioural surface because they can be calibrated against documented project-development and construction timelines. Lead times for mining projects are extensively documented and commonly exceed 15-18 years from discovery to production, with nickel projects generally exhibiting particularly long development horizons.[99] [100] [101] Copper's five years follows an IEA construction-phase estimate of four to five years. Nickel's 6.5 years reflects the IEA finding that nickel has the longest lead time of any of these metals, at roughly eighteen years from discovery to production, of which the construction phase is the portion the model represents *[IEA-GCMO-25]*. Lithium's and silicon's 2.5 years are the build times for chemical-conversion and polysilicon plants, respectively, both of which are processing

[95] Magnus Ericson and Mikael Höök. (2023). "Cobalt: corporate concentration 1975–2018." *Mineral Economics* 36: 483–501. Retrieved from: Springer; IGF. (2026). Commodity Profile: Cobalt. Retrieved from: IISD.

[96] USGS. (2025). 2022 Minerals Yearbook: Lithium. Retrieved from: USGS.

[97] Johannes, Bernreuter. (2023). The Crash of the Polysilicon Price Was Merely a Matter of Time." *Bernreuter Research*. Retrieved from: Bernreuter; Vincent Shaw. (2026). China Adopts Tougher Mandatory Energy Limits for Polysilicon Plants. *pv magazine.* Retrieved from: PV Magazine; Jiashu He. Regulatory Impacts on the Development of Solar Industry in China: A Policy Analysis. *Advances in Economics, Management and Political Sciences* 157 (2025): 1–9. Retrieved from: EWA.

[98] Oliver Wyman, Independent Review of Events in the Nickel Market in March 2022: Final Report (January 2023). Retrieved from: LME; Financial Conduct Authority, Final Notice: The London Metal Exchange (RIE), March 19, 2025, FCA final notice. Retrieved from: FCA; Financial Conduct Authority, Bank of England, and Prudential Regulation Authority, Joint Statement from UK Financial Regulation Authorities on London Metal Exchange and LME Clear, April 4, 2022, joint regulatory statement. Retrieved from: FCA.

[99] IEA. (2022). Promoting Exploration, Production and Innovation. In Introducing the Critical Minerals Policy Tracker. Paris: International Energy Agency, 2022. Retrieved from: IEA

[100] Paul Manalo. (2024). Average Lead Time Almost 18 Years for Mines Started in 2020–23. S&P Global Market Intelligence, April 10, 2024. Retrieved from: S&P

[101] Paul Manalo. (2023). Discovery to Production Averages 15.7 Years for 127 Mines. S&P Global Market Intelligence, June 6, 2023. Retrieved from: S&P.

rather than mining assets and take two to three years. Cobalt's four years is the weakest of the four: cobalt is overwhelmingly a copper and nickel co-product, so its expansion is governed by host-metal project economics rather than by a cobalt project pipeline, and the value is a modelling assumption informed by that dependence rather than a directly measured lead time. The paired `expectation_horizon_years` of three years for nickel, against a global two, encodes the same slow-capital reasoning: investors in long-lead assets integrate the price signal over a longer window.

Collection ceilings are anchored on the recycling evidence already used for the routing artefact. Nickel's ceiling of 1.20 — the highest of the material-level values — reflects the well-developed stainless scrap circuit documented in the global nickel cycle work *[RECK-08]*. Lithium's 1.60 is the largest headroom in the model and is explicitly prospective rather than observed: it represents policy-driven collection capacity that does not yet exist, sized against the battery collection and recovery mandates *[EU-2023/1542]*, and is paired with a deliberately low build gain on the reasoning that lithium collection will be driven by regulation rather than by scrap price. Silicon's 1.05 reflects essentially absent end-of-life recovery, with its build gain set to zero for the same reason — photovoltaic and semiconductor silicon recycling is nascent and not price-responsive. Cobalt's 1.10 holds at the copper/nickel default and does not presume a large policy-driven recovery uplift in the baseline. Copper's electronic-grade stream carries a slightly softer ceiling of 1.15 and a reduced build gain, on the reasoning that the printed-circuit-board and e-waste circuit responds more to policy and informal-sector dynamics than to the wire-scrap price loop; these two are the only entries in the slice surface that Table 41's columns do not display, which is why that row appears empty.

Price sensitivities and envelope ceilings are calibration outcomes disciplined by market structure, not independently sourced parameters, and should be read that way. The global price-scarcity sensitivity of 0.55 is a copper/nickel anchor drawn from the commodity price-dynamics literature *[PINDYCK-94]*; copper takes no override and therefore runs at that anchor. Each override can move away from it in the direction market depth implies: cobalt to 0.95 on DRC supply concentration, lithium to 1.10 and its battery-grade sub-market to 1.20 as the thinnest market in the set, nickel's Class I sub-market to 0.85 because the 2022 London Metal Exchange short squeeze[102] was a Class I phenomenon rather than a nickel-wide one, and silicon's solar and semiconductor grades to 0.95. Silicon's metallurgical grades move the other way, to 0.55, deliberately returning to the deep-market anchor because ferrosilicon and chemical-alloying silicon are abundant, glass- and steel-grade commodities whose behaviour resembles copper's far more than it does solar-grade polysilicon's. Envelope ceilings follow announced supply expansion: lithium's 2.20 on IEA guidance that lithium supply roughly triples by 2030 under an announced-projects view, silicon's 2.00 on polysilicon capacity build-out, nickel's 1.60 on cumulative envelope growth to 2050, and cobalt's 1.45 as a conservative reading of co-product expansion headroom *[IEA-GCMO-25]*.

Demand elasticities and noise amplitudes carry the weakest per-material justification. The global short-run elasticity of 0.10 and long-run 0.30 come from metals-demand econometrics *[DAHL-20]*; the two per-material departures are reasoned rather than estimated. Lithium is set to 0.22 in both the short and long run, on the argument that substitution towards lithium-iron-phosphate and sodium-ion chemistries is already visible in the short-run response and so no further long-run widening is warranted. Silicon is set to 0.10 in both, on the argument that photovoltaic demand has

[102] Eco3. (2022). Why Nickel Is the Most Fragile Base Metal: The 2022 LME Squeeze as Structural Evidence. Retrieved from: Eco 3; Xinpeng Sun. (2026). Looking Back at the LME Nickel Crisis (Part I): A Crisis the Exchange Could Barely Intervene In. Retrieved from: SSRN.

no material substitute for silicon and is therefore near-inelastic even over long horizons. Cobalt's short-run 0.10 is held at the global default explicitly to avoid assuming rapid demand substitution. Cobalt's demand-noise amplitude and lithium's 0.08 are modelling assumptions about supply concentration and policy-driven demand volatility rather than measured quantities. Two collection-effort build gains — cobalt's and lithium's, both reduced to 0.04 in 2026 — were retuned downward to keep the implied elasticity below the conservative upper bound of the scrap-supply band while freeing response budget for the endogenous circular-loop layer. That change is an internal coherence adjustment between two competing response channels, not a revision of evidence, and it is the clearest example in this surface of a value set by model-behaviour considerations rather than by an external anchor.

The per-material overrides in Table 41 carry most of the model's behavioural differentiation, and they encode a consistent economic reading of the five markets. Copper anchors the low end: deep, liquid and well supplied by scrap, it takes no price-sensitivity override, running at the global copper/nickel value of 0.55 with a long five-year capacity lag. Lithium sits at the opposite end with a price sensitivity of 1.10, rising to 1.20 for battery grade, the shortest capacity lag at 2.5 years, and much the widest capacity envelope at 2.20 — a thin market in which conversion capacity can be added comparatively quickly and prices move sharply. Cobalt carries elevated demand noise at 0.10 against a global 0.04, reflecting the greater susceptibility of cobalt supply to production disruptions and institutional shocks arising from its high geographic concentration, and the highest retirement-slack weight at 3.5, making it the fastest of the five to shed idle capacity and therefore the quickest to work off an oversupply — recovery from overshoot in roughly four years, consistent with the observed four-to-six-year cobalt cycle. Copper is the slowest on that measure at 1.5, reflecting mines that idle reluctantly and a cycle closer to ten years. Silicon's polysilicon slice carries by far the highest service-stress coupling gain at 22, and nickel's Class I slice the next highest at 20, so unmet service in those streams feeds back into scarcity pricing far more aggressively than elsewhere. They are high for opposite reasons: silicon's because the stream is genuinely supply-constrained and the gain reproduces an observed price episode; nickel's because the modelled physical signal is almost absent and 20 is the largest value still yielding a defensible response. Nickel's gain compensates for a weak signal rather than transmitting a strong one.

Table 41. *Per-material System Dynamics overrides. Blank entries inherit the global value from Table 39.*

Slice	Price sensitivity	Capacity lag (y)	Expansion gain	Envelope max	Service-stress gain
copper (base)	—	5.0	0.15	—	7.0
copper (electronic grade)	—	—	—	—	—
nickel (supply inertia)	—	6.5	0.15	1.60	—
nickel (class I)	0.85	—	—	—	20.0
cobalt (thin market)	0.95	4.0	0.12	1.45	7.0
lithium (chemical conversion)	1.10	2.5	0.20	2.20	7.0
lithium (battery grade)	1.20	—	—	—	—
silicon (polysilicon)	0.80	2.5	0.22	2.00	22.0

Slice	Price sensitivity	Capacity lag (y)	Expansion gain	Envelope max	Service-stress gain
silicon (solar grade)	0.95	—	—	—	—
silicon (metallurgical grade)	0.55	—	—	—	—

Coupling control parameters. A final small surface governs the exchange between the two layers. The service-stress channel carries a substantial gain of 3.5, so unmet service feeds back into scarcity pricing strongly. The circular-supply channel is split deliberately: its gain into scarcity pricing is set to exactly zero while its gain into collection effort is set to 0.8, meaning that a low recycled share in the active baseline drives collection effort directly without also suppressing demand through a price signal. Signal smoothing under-relaxes the iteration at 0.2 for the main channels and 0.35 for the strategic channel by default, though the run entrypoint lowers the strategic value to 0.25 on the reasoning that physical stock is slow-moving and a lower smoothing factor damps oscillation in the nickel coverage signal between iterations. Two caps bound the feedback: the service-stress signal itself is capped at 0.6, and the resulting scarcity multiplier at 3.0 globally — though the per-material overrides raise that multiplier cap substantially for the thin markets, to 15 for lithium battery grade and 12 for both lithium conversion and silicon solar grade. Table 42 gives the full surface.

Table 42. *Coupling control parameters governing signal exchange between the SD and dMFA layers.*

Parameter	Value	Effect
coupling_service_stress_gain	3.5	Strength of the unmet-service to scarcity-price channel
coupling_circular_supply_stress_gain	0.0	Circular-supply to scarcity-price channel, deliberately disabled
coupling_circular_collection_gain	0.8	Circular-supply to collection-effort channel
coupling_signal_smoothing	0.2	Under-relaxation factor on the main signal channels
coupling_signal_smoothing_strategic	0.25 (run-level; 0.35 default)	Under-relaxation on the strategic-reserve channel
service_stress_signal_cap	0.6	Ceiling on the service-stress signal itself
coupling_stress_multiplier_cap	3.0 (up to 15 per material)	Ceiling on the resulting scarcity multiplier
max_iter	30	Maximum inner-loop iterations before the cap binds
convergence_tol	0.02	Inner-loop convergence tolerance
feedback_signal_mode	time_series	Signals exchanged as full trajectories, not scalar means
feedback_on_report_years_only	false	Feedback applied in each modelled year

Calibration

Documents how the model's adjustable quantities were set and against what evidence: the initialisation of pre-study-period stock, the activation of behavioural dynamics against historical episodes, the tuning of the behavioural parameter surface, and the shipped automated-fitting capability that is not exercised in the deployed configuration.

What is calibrated, and what is not

The word *calibration* carries more than one meaning in this model, and collapsing them produces a misleading account. No single fitting procedure sets the model's parameters. Instead, four distinct activities operate on different parts of the model, judged against different evidence, and carrying different degrees of confidence. Only one of the four is an automated statistical fit, and that one is not exercised in the deployed configuration. Setting the four apart at the outset is the necessary precondition for reading the rest of this section — and for reading the model's outputs with the right degree of trust.

The first activity is **initialisation**: setting how much material was already in use before the study period begins, so that the model does not start from an empty world. The second is **historical activation**: deciding when each material's behavioural dynamics switch on, and testing the resulting behaviour against documented historical episodes. The third is **behavioural tuning**: setting the gains, lags and elasticities of the System Dynamics layer so that its response to stress is of a defensible magnitude. The fourth is **automated parameter fitting**: running an optimiser over a bounded parameter set to minimise deviation from observed series. Table 43 sets out the four, and the remainder of this section develops each in turn.

Two boundaries constrain all four. The first is the executable parameter allowlist. The active contract fixes a named parameter schema enumerating exactly which parameters the material-flow layer will accept, and within that a strictly smaller *calibratable* subset that an automated routine is permitted to adjust — twenty-seven active parameters, of which seven are exposed to calibration (given in full as Table 36 in the Data section). The exclusion is deliberate and not merely conservative: it keeps fitting away from any parameter whose value encodes a policy choice, a scenario lever, or a physical constant rather than an uncertain empirical quantity. A configuration that sets a parameter outside the active list is rejected. The second boundary separates observation from state. The observed series held in the repository for calibration purposes — in-use stock, refining output, semis output, manufacturing output and secondary recycling output — are read for scoring only. None of them is ever used to drive the model's state.

Table 43. *Calibration activities in the deployed model: what each adjusts, what it is judged against, and its status.*

Activity	Layer	What is adjusted	Judged against	Status in the active baseline
Initialisation	dMFA	Per-material synthetic burn-in length; new-scrap	Independent in-use stock estimates at a benchmark year	Active. Per-material profile deployed.

Activity	Layer	What is adjusted	Judged against	Status in the active baseline
		recovery rate; stockpile release rate		
Historical activation	SD	Per-material dynamics start year; per-episode tether strength and supply multiplier	Documented historical supply and price episodes, scored warm against control	Active. Eleven sourced event windows; two materials publication-grade, three provisional.
Behavioural tuning	SD and coupling	Price sensitivities, elasticities, capacity and collection gains, reserve thresholds, supply elasticity	Magnitude of one documented price episode per material; econometric literature bounds	Active. Global anchors with per-material overrides.
Automated parameter fitting	dMFA	Seven per-stream levers across four streams, for copper and nickel only	Observed in-use stock and secondary recycling series, on a train/validation split	Shipped and runnable, but not exercised. Fitted values were promoted and then reverted; literature priors were deployed.

A fifth activity sits adjacent to these and is documented elsewhere in this report rather than here. The exogenous input artifacts — demand, supply, stream composition, routing rates, lifetimes — are themselves constructed against literature anchors and guarded by build-time reconciliation gates that abort the build if a stream-share-weighted effective rate drifts outside its tolerance band. That anchoring work is described per artefact in the Data section, where the artefacts themselves are documented, and is not repeated here. The key distinction is that anchoring sets the *inputs* the model consumes, whereas the four activities in Table 43 set how the model *behaves* once those inputs are supplied.

Initialisation: Seeding the pre-study-period stock

The model's reporting-relevant horizon begins in 1950, but copper, nickel, and, to a lesser degree, cobalt and silicon had been mined, fabricated, and accumulated in use for decades or centuries before that date. A model that began in 1950 with zero in-use stock would understate the material base at the start of the run and — because in-use stock retires on lifetimes of ten to sixty years — would carry that understatement forward as an understatement of end-of-life availability for most of the study period. Each recycling and circularity indicator the model reports depends on getting this initial condition approximately right.

The mechanism used is a *synthetic burn-in*: a spin-up period ahead of the study horizon in which the model is run forward on its own machinery to accumulate a starting stock, rather than being handed one. The concept, and the reason it is preferred to a directly imposed initial stock, are described in the Model Architecture section. What matters here is how long each material's burn-in runs, and how that length was chosen.

A shared template supplies the defaults — a twenty-year burn-in, disabled unless a material opts in, using a repeat-anchor-year rule that holds the earliest available demand year constant across the spin-up. All five materials opt in, and three override the default length. Each override was set by

running the model at candidate lengths and comparing the resulting stock at a benchmark year against an independent estimate, then keeping the length that was reconciled. Table 44 gives the deployed profile.

Table 44. *Per-material initialisation profile in the deployed configuration, with the basis for each burn-in length.*

Material	Burn-in years	Basis
Copper	20 (default)	The shared default is retained; copper's long documented industrial history is carried by the demand series itself rather than by a lengthened spin-up
Nickel	80	Set so that integration starts at 1870, the first year for which a real nickel demand figure exists. The result saturates rather than converging: pre-1905 demand is small enough that further extension adds nothing
Cobalt	10	Calibrated to the 82 kt 1950 stock benchmark based on a top-down inflow-convolution model, anchored at 2019 = 471 kt; a ten-year spin-up returns 54 kt, a better fit than the 20 year default
Lithium	20 (default)	The shared default is retained; lithium's pre-battery civilian demand is small enough that spin-up length is not the binding constraint on its 1950 stock
Silicon	33	Calibrated to the 12,000 kt 1950 stock benchmark after a lifetime revision raised the uncalibrated figure to 16,350 kt; a thirty-three-year spin-up returns 12,409 kt

The lithium row shows the parameter resolution chain working as designed because the material-level and stream-level values disagree by construction rather than by error. The shared baseline sets a global fallback new-scrap recovery rate of 0.75; the initialisation profile narrows lithium to 0.0 at material level; and the per-stream surface then raises the battery-grade stream alone back to 0.50. The effective value therefore depends on which stream is being asked about — zero for industrial and specialty lithium, where fabrication scrap is recovered on-site inside the fabrication yield and never leaves as a separate flow, and 0.50 for battery-grade lithium, where a genuine external recovery loop exists. Reading the material-level value alone would misstate the model.

Silicon: a worked example of the anchor method. Silicon's burn-in was reduced from eighty years to thirty-three, and the reason illustrates how tightly initialisation and lifetime parameterisation are coupled. A revision to the ferroalloy-grade lifetime distribution — widening it from a normal distribution with a mean of fifty years and a standard deviation of fifteen to a mean of sixty-five and a standard deviation of twenty — lengthened the residence time of silicon in its dominant stream. Because a longer lifetime means less retirement, the same burn-in demand history accumulated substantially more stock: the uncalibrated 1950 figure rose from 12,200 kt to 16,350 kt, an increase of roughly 36%, against a benchmark of 12,000 kt. The burn-in length was therefore re-derived rather than left at its inherited value, and a thirty-three-year spin-up returned a 1950 stock of 12,409 kt, within 3.4% of the benchmark. The general point is that burn-in length is not an independent parameter: it is conditional on the lifetime distributions in force, and a change to those distributions invalidates a previously calibrated length.

Nickel: A documented limit rather than a fit. Nickel's eighty-year burn-in is the longest in the set and was chosen so that integration begins in 1870, the earliest year for which a real nickel demand figure exists — 271 t, rising to 79 kt by 1940. The inherited twenty-year default truncated the spin-up at 1930 and captured only the tail of a genuine accumulation. Lengthening it raised the modelled 1950 stock from 1,714 kt to 1,812 kt, and then stopped: pre-1905 demand is small enough that any further extension adds essentially nothing, so the result saturates rather than converging on a target. This matters because it is a bounded negative result. An independent top-down estimate places mid-century nickel stock substantially higher, and burn-in alone cannot reach it — the ceiling is approximately 1,812 kt regardless of length. The residual gap is therefore a question about the reconstructed demand history or the lifetime distributions, not about initialisation, and it is recorded as such rather than closed by extending the spin-up until the numbers agree. Related caution applies to the corresponding top-down nickel stock figure itself, which the project's own reference work finds is not internally consistent with the apparent-consumption history the model is built on.

Cobalt. Cobalt takes the opposite adjustment — a ten-year burn-in, shorter than the shared default — on the finding that a shorter spin-up reconciles better with its early stock anchors. It also carries the one structural override in the table: the refinery stockpile release rate is raised from its operational default of 0.25 to 1.0, removing a throughput cap so that recovered secondary cobalt can feed same-year supply rather than queueing in the buffer. This is a modelling assumption about cobalt's supply chain rather than a fitted quantity, and it is stated here because it changes secondary-supply timing in a way a reader comparing cobalt against the other four materials would otherwise not see.

One point about that configuration file is worth stating, as it may be misinterpreted. The same per-material blocks that carry the burn-in settings also carry two parameters that are not initialisation settings at all: the new-scrap recovery rate and, for cobalt, the refinery stockpile release rate. They sit there because the initialisation profile was made the owner of the material-level overrides that had previously been authored inline on the shared parameter surface, not because they apply only during spin-up. Both are ordinary material-flow parameters that act in each year of the run, and both are documented with the rest of that surface in the Data section rather than repeated here.

Historical activation: Testing behaviour against documented episodes

Initialisation seeds the physical stock. A separate question is when the behavioural layer should start responding, and whether its response is credible in size. The System Dynamics layer is activated per material at a configured start year, ahead of the reporting window, so capacity, backlog, and collection-effort states are already populated by the time results are reported, rather than starting from rest. During that historical period, the model's demand is softly tethered to the exogenous demand series — the tether mechanism, and the two independent activation gates that separate historical dynamics from reporting-window price response are described in the Model Architecture section and not repeated here.

What is documented here is the evidence base behind that activation. Eleven historical episodes across the five materials are held in a source-backed registry, each carrying its window, its citation, a physical interpretation, a scoring direction, a tether value and, where the episode is a physical supply constraint, a supply multiplier that cuts effective flow capacity for the duration. Table 45 gives the register in full. These are not scenario shocks: they are historical facts the model is asked to reproduce, and the register exists so a reader can check both whether the chosen episodes are representative and whether reproducing them is the right test.

Table 45. *Historical event register underlying the System Dynamics warm-start. Tether value is the strength of the anchor to observed demand during the window; the supply multiplier cuts effective flow capacity where the episode is a physical constraint.*

Material	Episode	Window	Scoring direction	Tether	Supply mult.	Source
Copper	Chilean nationalisation and 1970s supply disruption	1971–1978	Supply shock	0.30	0.80	[CODELCO][103]
Copper	China copper supercycle	2003–2008	Scarcity spike	0.20	—	[USGS-SIR-16],[104] [USGS-OFR-08][105]
Nickel	LME nickel boom on record-low inventories	1987–1989	Supply shock	0.20	0.88	[IMF-PC-89][106]
Nickel	China boom and global financial crisis crash	2006–2009	Demand spike (not scored)	0.20	—	[USGS-OFR-08]
Nickel	Indonesian nickel ore export ban	2014–2016	Supply shock	0.30	0.88	[USGS-FS-16],[107] [USGS-MCS][108]
Cobalt	Shaba / Kolwezi supply crisis	1978–1979	Supply shock	0.05	0.35	[NIOSH-79],[109] [CAMPBELL-19][110]
Cobalt	2007–2009 commodity cycle	2007–2009	Price proxy spike	0.25	—	[USGS-MYB-Co][111]
Cobalt	DRC and electric-vehicle demand spike	2016–2018	Scarcity spike	0.15	—	[USGS-MYB-Co]
Lithium	Electric-vehicle-driven lithium boom	2016–2018	Supply shock	0.20	0.80	[USGS-MCS], [USGS-MYB-Li][112]
Lithium	Oversupply and EV-sales correction	2019	Not scored (oversupply)	0.20	—	[USGS-MYB-Li]
Silicon	Photovoltaic polysilicon shortage and rebalance	2007–2011	Supply shock	0.10	0.75	[IEA-POLY-20][113]

[103] Corporación Nacional del Cobre de Chile. (n.d.). *Chile y el cobre, una historia entrelazada*; and *41 años de la nacionalización del cobre*. Codelco.
[104] U.S. Geological Survey. (2016). *Copper metal and scrap use and trade patterns* (Scientific Investigations Report 2016-5075). U.S. Department of the Interior.
[105] U.S. Geological Survey. (2008). *Factors that influence the price of metals* (Open-File Report 2008-1356). U.S. Department of the Interior.
[106] International Monetary Fund. (1989). *Primary commodities: Market developments and outlook*. IMF.
[107] U.S. Geological Survey. (2016). *Resource nationalism in Indonesia* (Fact Sheet 2016-3072). U.S. Department of the Interior.
[108] U.S. Geological Survey. (1995–2026). *Mineral commodity summaries* [Annual series]. U.S. Department of the Interior.
[109] U.S. Bureau of Mines. (1979). *Impact of the recent cobalt supply situation*. U.S. Department of the Interior.
[110] Campbell, G. A. (2019). *The cobalt market revisited*. Michigan Technological University.
[111] U.S. Geological Survey. (Annual). Cobalt. In *Minerals yearbook*; and *Statistical compendium* (1990). U.S. Department of the Interior.
[112] U.S. Geological Survey. (2018, 2019). Lithium. In *Minerals yearbook*. U.S. Department of the Interior.
[113] International Energy Agency. (2020). *Share of photovoltaics in polysilicon demand and polysilicon price, 1975–2010* [Data set]. IEA Publications.

How the episodes are scored. The comparison is not between the warm-started model and a cold-started one. It compares two warm-started runs—a calibrated run carrying the event windows and a control run using the identical per-material activation years but with the windows removed. Both therefore accumulate the same capacity-envelope state during spin-up, which isolates the effect of the event perturbation from the effect of having been warmed up at all. Each episode is scored on the diagnostic appropriate to its physics rather than on a single generic measure: a supply shock should raise flow utilisation, a demand spike or drop should move real demand in the corresponding direction, a scarcity spike should raise effective scarcity or bottleneck pressure or the service deficit, and a price-proxy spike should raise the price ratio.

Why some episodes carry a supply multiplier, and others do not. Loosening the tether alone produces no measurable difference between the calibrated and control runs whenever the capacity envelope is at or above unity, because with no accumulated backlog the model's freely computed demand equals the exogenous series and the blend is therefore inert. An episode that represents a genuine physical constraint is consequently given a supply multiplier below one, which cuts effective flow capacity for the window and raises flow utilisation in the calibrated run relative to the control. Episodes without a multiplier score no deviation, and that zero is retained as an honest mechanistic result rather than suppressed or engineered away. A reader should understand the five rows in Table 45 without a supply multiplier as documented references rather than as passed tests.

Two episodes are deliberately not scored. The nickel window spanning 2006 to 2009 covers both the China-driven boom and the financial-crisis crash within a single interval, and no single capacity constraint can represent a period that ran hard in both directions; imposing one would misstate the physics to obtain a score. The lithium window in 2019 is a supply glut combined with weaker-than-expected vehicle sales, and the model currently has no lever that represents supply excess — a multiplier above one would raise flow capacity and therefore lower flow utilisation, the wrong direction for any supply-shock or scarcity scoring. Both windows are retained in the register for documentation, explicitly flagged as unscored, and left for a future mechanism rather than forced into an existing one.

Structural invariant. One property must hold regardless of any of the above: turning the historical warm-start off entirely must reproduce exactly the results of a cold run beginning directly in the reporting window. This confirms that warm-start changes only how the model is initialised and not what it computes once running, and it is verified by test rather than asserted. If a future change to the warm-start machinery broke this equivalence, warm-start would have become a silent modifier of reported results rather than an initialisation device.

A second, independent gate: does the resulting stock trajectory still look like history? Everything above scores individual episodes — a warm run against a control run, on a behavioural diagnostic appropriate to that one event. Accepting a material's warm-start configuration additionally requires a broader check that sits outside the event register entirely. The behavioural demand path the warm-start produces feeds directly into the material-flow layer's accounting, and its cumulative effect on in-use stock across the full calibration window is checked separately: the warm-started run's stock trajectory is compared by root-mean-square error against the observed stock series, and against a cold run with no warm-start at all, with the warm-started error required to stay within 5% of the cold run's error (or within a small floor set by the scale of the observed series, whichever is

larger). This differs from event scoring in two respects: it is judged against a cold baseline rather than the dynamics-matched control, and it tests the physical consequence of the whole warm-start path rather than any single episode. A material's warm-start is only accepted once both gates are satisfied: the stock trajectory passes its error tolerance, and at least one event reproduces in the documented direction.

Status by material. Acceptance rests on two independent gates, and it is worth being precise about which materials clear which, because no material clears both without qualification. On the event-window gate — scored against a control run carrying the same per-material activation years with the event windows removed — copper reproduces two of two scored episodes, cobalt three of three, lithium one of one and silicon one of one, while nickel reproduces two of three: its 2006–2009 China-boom and financial-crisis demand spike moves in the wrong direction, by 0.4%. Lithium carries a second documented episode, the 2019 oversupply and EV-sales correction, which is recorded but left unscored because an oversupply correction has no unambiguous expected sign on the available diagnostics. On the stock-trajectory gate, all five materials pass, but the gate discriminates weakly: warm and cold-start runs differ in calibration-window stock RMSE by at most 0.4% against a tolerance of 5%, so a pass records non-falsification under a weak test rather than agreement with an independent series. For cobalt, lithium and silicon the test is weaker still, because the observed stock series the gate is checked against is itself derived from an earlier run of the model's own baseline rather than sourced independently; a pass or fail therefore reflects agreement with that snapshot, not validation against observation. Copper is accordingly the only material sound on both axes; nickel has an independently sourced comparator but partial event evidence; cobalt, lithium and silicon have complete event evidence on a synthetic comparator. This limitation is governed and explicitly flagged rather than left implicit. Replacing the three synthetic series with external historical data, rerunning the gate, and resolving nickel's unreproduced demand episode are recorded as outstanding scientific work.

Behavioural tuning of the System Dynamics surface

The behavioural layer carries a surface of gains, lags, elasticities and thresholds that determine how strongly the model responds to scarcity and how quickly it adjusts. These are not structural parameters of an explicit decision model and cannot be measured directly; they are reduced-form coefficients, and the question is whether their values produce responses of a defensible magnitude. Two complementary approaches are used, and both are documented in the Data section where the parameter surface itself is tabulated: global anchors are drawn from the econometric literature on commodity price dynamics and demand response, and per-material overrides are set so the model reproduces the approximate magnitude of one documented historical price episode for that material (Table 40). This is a targeted plausibility calibration rather than a statistical fit, and the section does not repeat that material here.

The order of tuning, and why it matters. What is worth recording is the sequence in which the surface is tuned, because for a system of coupled feedback loops the order is not arbitrary. Baseline stability and convergence are established first: a model that has not converged cannot be tuned, because any observed response may be an artefact of the solver rather than of the mechanism. The loop is then subjected to realistic stress, since a parameter governing response magnitude is unidentifiable in a run where nothing happens. Bottleneck intensity is tuned next, then capacity

response speed, then collection response, and strategic-reserve intents last. The ordering runs from the mechanisms that shape the response signal toward those that consume it, so that each is tuned against an upstream signal already settled. Two working principles accompany it: the historical phase is kept deliberately conservative, and broad simultaneous sweeps across strongly coupled controls are avoided, because a coupled system will offer many parameter combinations that fit equally well while implying quite different dynamics.

A related consequence deserves stating plainly. Because these coefficients are tuned against episode magnitudes rather than fitted to series, and because several of them can compensate for one another, they should be read as a jointly plausible configuration rather than as independently identified quantities. The identifiability limitation is recorded in Table 61 under Assumptions and Limitations.

Run-level per-material calibration. A further set of calibrated values sits not in the shared parameter surface but in the canonical run entrypoint itself, applied as per-material slice overrides. Two families are involved. The first governs the strategic reserve: rather than enabling or disabling the mechanism wholesale, the configuration enables it and then calibrates its behaviour per material, including neutralising it entirely for two materials by setting their target coverage to zero. The second sets the primary-supply elasticity, which governs how strongly primary supply responds to price. Table 46 gives both. The distinction between them is one of evidence: the reserve settings are modelling assumptions about each material's market and policy context, whereas the supply elasticities are empirical estimates taken directly from a published econometric study of short-run global supply response.

Table 46. *Per-material calibrated overrides applied at the canonical run entrypoint.*

Material	Strategic reserve settings	Primary supply elasticity	Basis for the elasticity
Copper	Target coverage 0.0 — mechanism neutralised	0.10	Short-run global supply elasticity, reported range 0.08–0.11 [SHOJAEDDINI-25][114]
Nickel	Target coverage 0.1; release price threshold 1.25	0.31	Short-run global supply elasticity [SHOJAEDDINI-25]
Cobalt	Fill price threshold 1.3; release price threshold 1.4; release service threshold 0.30	0.16	Short-run global supply elasticity; reflects artisanal-mining and inventory response rather than a near-zero constraint [SHOJAEDDINI-25]
Lithium	Fill price threshold 1.15; fill service threshold 0.15; release service threshold 0.20	0.33	Short-run global supply elasticity [SHOJAEDDINI-25]
Silicon	Target coverage 0.0 — mechanism neutralised	0.28	Short-run global supply elasticity for silicon metal and ferrosilicon [SHOJAEDDINI-25]

The elasticity values match the model at the refinery and aggregate level, where the supply cap binds in the material-flow layer, so the econometric estimate and the modelled quantity refer to the same

[114] Shojaeddini, E., Alonso, E., & Nassar, N. T. (2025). Understanding market sensitivity: Estimation of supply and demand elasticities for non-fuel minerals. *Mineral Economics, 38*, 985–996. doi:10.1007/s13563-025-00537-3

thing. Three of the five values were revised upward from initial design priors following that econometric pass, and all five are now recorded as grounded rather than assumed. Endpoint configurations holding the elasticity at zero and at one are retained as sensitivity bounds.

The conservative cobalt configuration. Cobalt carries the most restrained behavioural settings of the five materials — a demand price elasticity of 0.10, a capacity envelope ceiling of 1.45, a collection multiplier ceiling of 1.10, and a service-stress coupling gain of 7.0. This combination was selected deliberately over more responsive alternatives. Global cobalt mine supply is highly concentrated geographically, with a substantial share originating in the Democratic Republic of the Congo,[115] and this concentration is reflected in the model through stronger responses to supply-side shocks;[116] the accepted stance was to constrain the response rather than to maximise the fit. Readers comparing cobalt results against the other four materials should be aware that part of the difference in behaviour is this deliberate choice rather than an emergent property of cobalt's physical system.

Automated parameter fitting: A shipped capability, not the deployed method

The model ships a complete automated calibration facility, and the deployed configuration does not use it. Because the capability is real, runnable, and documented in the repository, and because it was exercised and at one point promoted into the default configuration, describing it accurately—including why it is no longer active—matters more than omitting it would. This subsection sets out what the facility does and its status. It should not be read as a description of how the deployed model's parameters were set; that is the subject of the preceding three subsections.

How the facility is organised. Two layers work together, and their separation is the design's main safeguard. A calibration specification is a complete optimisation profile: it names the modelled target variables and the observed sources they are scored against, any material and region filters, the training and validation windows, the optimiser settings, and per-parameter initial values, bounds and constraints. The parameter schema, separately, is the executable allowlist described earlier; it overlays no additional parameters of its own. Keeping optimisation-specific content out of the parameter contract prevents stale or overly broad specifications from changing the search space without explicit review, including by introducing unintended parameters, retaining obsolete constraints, or excluding valid configurations.. Calibration is never implicit: it runs only when a complete specification is passed explicitly at the command line; otherwise, the canonical and material-specific run configurations share the same default-neutral contract.

Scope of the shipped specifications. The material-flow-side workflow is intentionally narrow. It covers copper and nickel only, is multi-objective only, and targets two quantities: in-use stock and total secondary recycling output, each filtered to the single material at the global region and each requiring at least twenty years of observations before a material-region pair is admitted for scoring. Each specification tunes seven levers — sorting yield, recycling yield, remanufacturing yield, semis fabrication yield, finished-product manufacturing yield, product lifetime multiplier, and collection rate multiplier — across four material streams, giving twenty-eight adjustable parameters per

[115] USGS. (2022). Minerals Yearbook. Retrieved from: USGS; USGS, 2022 Minerals Yearbook: Congo Retrieved from: USGS.
[116] USGS. (2022). Minerals Yearbook. Retrieved from: USGS; USGS, 2022 Minerals Yearbook: Congo Retrieved from: USGS.

material. A third specification exists for the origin-destination trade layer, which the active baseline disables; consistent with the treatment of trade throughout this document, it is noted as an opt-in workflow and not developed further.

The exclusion of the collection rate itself from the adjustable set is structural rather than stylistic, and follows from the two-gate routing schema. Collection rate is supplied per stream and per year by an exogenous artefact carrying real detail — activation years, scenario-dependent trajectories, stream-specific ceilings. Permitting an optimiser to overwrite that with a single fitted scalar would destroy the structure. The multiplier scales the authored trajectory proportionally instead, so a fit can adjust the level while preserving the shape and the per-stream heterogeneity beneath it.

Scoring and optimisation. This is a different use of the same observed series from the trajectory-level gate in the historical activation subsection above: that gate holds the parameter surface fixed and tests whether a material's warm-start should be enabled at all; this facility holds warm-start fixed and searches the parameter surface itself. Deviation is measured per material-region pair as root-mean-square error, mean absolute error, and a normalised root-mean-square error expressing the first relative to the mean observed value. Pairs are aggregated into a single figure as a weighted mean, weighted on mean observed stock so that a material-region pair holding a large share of the observed mass is not outweighed by a small one. The horizon is split into training and validation segments derived from the model's own time configuration, with the final twenty years reserved for validation. The optimiser searches on the training score but selects on the validation score, falling back to the training score only if validation selection is unavailable — a split that guards against a parameter set that fits the fitted period and generalises poorly. The search itself is two-stage: a differential-evolution global pass of one hundred and sixty evaluations at a population size of eight and a fixed seed, followed by a Nelder–Mead local refinement of up to eighty evaluations, with early stopping once improvement falls below a threshold for three consecutive rounds. The fixed seed makes the search reproducible.

Why the fitted values are not deployed. The copper and nickel specifications were run, and their results were promoted into the canonical parameter surface. That promotion was subsequently reverted, and the surface returned to its documented literature-prior values, where it remains. The reason was not a technical failure of the fit. Benchmark-facing comparison runs showed that the fitted values improved agreement on some targets while weakening it on others—specifically on the end-of-life and circularity side, which is the part of the model's output the report depends on most. Fitting the model to stock and secondary-recycling series was, in other words, achievable at the cost of end-of-life recovery rates that no longer reconciled with independent recycling-rate benchmarks. Given that trade-off, the documented literature-prior surface was preferred as the canonical default until a revised promotion decision was made, and the calibration entry points were retained as legacy workflows.

Two consequences follow for how to read this document. First, each per-stream material-flow parameter reported in the Data section is a literature-anchored value, not a fitted one, and the confidence labels attached there are labels about the strength of that literature rather than about goodness of fit. Second, the model's agreement with observed stock is therefore a genuine out-of-sample property of the parameterisation rather than a fitted outcome — which is a stronger claim than a fit would have supported, and is stated here because the distinction is easy to lose.

Sources for the calibration and initialisation evidence base

The sources below support the settings documented in this section. They are listed separately from the consolidated exogenous-data register given in the Data section, which covers the sources behind the input artefacts themselves. Where a source appears in both, the key is the same. Entries drawn from parameter provenance recorded in the model's configuration rather than from a formal bibliography are marked as such, so that a reader can tell which attributions have been checked against the publication and which are carried forward from the configuration record. Table 47 lists them.

Table 47. *Source register for the calibration, initialisation and warm-start evidence base.*

Key	Source	Settings supported
[HELBIG-20][117]	Helbig et al. (2020), as recorded in the model's parameter provenance comments; the bibliographic entry is held in the configuration rather than in a formal reference list and has not been independently verified for this document	Per-material new-scrap recovery rates for copper, nickel, cobalt and lithium
[GLOSER-17][118]	Glöser-Chahoud, S. (2017), doctoral thesis, Karlsruhe Institute of Technology	Copper new-scrap recovery rate of 0.923
[RECK-12][119]	Reck and Rotter (2012), as recorded in the model's parameter provenance comments; see also [RECK-08][120] in the Data section register	Nickel new-scrap recovery rate of 0.844, reflecting integration of new scrap into stainless mill feed
[HARPER-12][12] [121]	Harper, E.M. et al. (2012), Environ. Sci. Technol. 46, 1079–1086	Cobalt new-scrap recovery rate of 0.047; cobalt in-use stock benchmark
[BRGM-19][122]	BRGM, silicon metal criticality assessment (2019, public deck)	Treatment of silicon new scrap as internal to the combined fabrication-and-manufacturing yield
[SHOJAEDDINI-25]	Shojaeddini, E. et al. (2025), Mineral Economics 38, 985–996 (U.S. Geological Survey), short-run global price elasticity of supply	Central per-material primary-supply elasticity: Copper: 0.10; Nickel: 0.31; Cobalt: 0.16 Lithium: 0.33; Silicon: 0.28
[CODELCO]	Codelco corporate historical record: Chile and Copper: An Intertwined History; and 41 Years Since the Nationalization of Copper	Copper 1971–1978 supply-shock window
[USGS-SIR-16]	USGS Scientific Investigations Report 2016-5075, copper metal and scrap use and trade patterns	Copper 2003–2008 supercycle scarcity reference

[117] Helbig, C., et al. (2020). *[Recorded only in the model's parameter provenance comments; no formal bibliographic entry exists in the repository and the reference has not been independently verified for this document. The per-stream values it supports are given in the parameter surface tables.]*
[118] Glöser-Chahoud, S. (2017). *Quantitative Analyse der Kreislaufwirtschaft metallischer Rohstoffe* [Doctoral dissertation]. Karlsruhe Institute of Technology.
[119] Reck, B. K., & Rotter, V. S. (2012). *[Recorded only in the model's parameter provenance comments and not independently verified for this document; see Reck et al. (2008) in the Data section register for the traced nickel-cycle reference.]*
[120] Reck, B. K., Müller, D. B., Rostkowski, K., & Graedel, T. E. (2008). Anthropogenic nickel cycle: Insights into use, trade, and recycling. *Environmental Science & Technology, 42*(9), 3394–3400. doi:10.1021/es072108l
[121] Harper, E. M., Kavlak, G., & Graedel, T. E. (2012). Tracking the metal of the goblins: Cobalt's cycle of use. *Environmental Science & Technology, 46*(2), 1079–1086. doi:10.1021/es201874e
[122] Bureau de Recherches Géologiques et Minières. (2019). *Silicon metal criticality assessment*. BRGM.

Key	Source	Settings supported
[USGS-OFR-08]	USGS Open-File Report 2008-1356, factors that influence metal prices	Copper supercycle and nickel 2006–2009 boom-and-crash windows
[IMF-PC-89]	International Monetary Fund, Primary Commodities: Market Developments and Outlook (1989)	Nickel 1987–1989 boom on record-low exchange inventories
[USGS-FS-16]	USGS Fact Sheet 2016-3072, resource nationalism in Indonesia	Nickel 2014–2016 ore export ban window
[NIOSH-79]	U.S. Bureau of Mines / NIOSH, Impact of the Recent Cobalt Supply Situation (1979)	Cobalt 1978–1979 Shaba/Kolwezi supply crisis
[CAMPBELL-19]	Campbell, G.A. (2019), The cobalt market revisited, Michigan Technological University	Cobalt Shaba/Kolwezi crisis magnitude and market context
[USGS-MYB-Li]	USGS, Minerals Yearbook — Lithium (2018 and 2019 editions)	Lithium 2016–2018 boom and the 2019 oversupply correction
[IEA-POLY-20]	IEA, share of photovoltaics in polysilicon demand and polysilicon price, 1975–2010	Silicon 2007–2011 polysilicon shortage and rebalance window

Scenarios

Documents the scenario architecture: the three orthogonal axes along which futures are varied, the levers and ambition rungs that define each circularity setting, the cumulative ladder that makes results attributable to individual interventions, the resilience shock inventory, and the executed run surface.

The three axes

A scenario in this model is a coordinated setting of three independent things: how much material the economy demands and in what composition, how aggressively circular-economy interventions are pursued, and whether the system is subjected to a disruption. These are varied along three axes that are orthogonal by construction — each resolves at a different level of the configuration, so any combination of the three is expressible, and none constrains another. Table 48 sets them out.

Table 48. *The three scenario axes, the question each answers, and where each resolves in the configuration.*

Axis	Question	Mechanism	Resolves at
1 — Demand	How much material is required, and in what stream composition?	A single selector that swaps the total-demand trajectory and the stream-composition trajectory together	Run level: which data trajectories load
2 — Circularity	How aggressively are circular-economy interventions pursued?	Three lever families at three ambition rungs, applied across five materials	Scenario level: lever blocks and value surfaces
3 — Shocks	How does the system behave under disruption, and how fast does it recover?	Named shock events applying a multiplier to one channel over a defined window	Per scenario: shock blocks on a host cell

The two demand and circularity axes are named differently in three places a reader may be holding at once — the digital report, the data explorer, and this document — and the differences are not cosmetic: the report's ambition labels and this document's rung labels do not share a word in common. Table 50 gives the crosswalk once, so it is not repeated at each mention below.

Table 49. *Terminology crosswalk for the demand and circularity ambition axes: report language, explorer controls, and the identifiers used in this document and its configuration.*

Axis	Digital report	Data explorer	This document / configuration
Demand — low	Low decarbonisation	STEPS	STEPS (demand_sel: steps) — IEA Stated Policies Scenario
Demand — medium	Medium decarbonisation	APS	APS (demand_sel: aps) — IEA Announced Pledges Scenario

Axis	Digital report	Data explorer	This document / configuration
Demand — high	High decarbonisation	NZE	NZE (demand_sel: nze) — IEA Net Zero Emissions by 2050 Scenario
Circularity — low	Low ambition	Not a selectable control — the explorer's rung selector offers only Medium and High	Current-practice rung (“L”): each lever neutral, implicit baseline
Circularity — medium	Medium ambition	Medium (M)	Moderate rung (rung_sel: M)
Circularity — high	High ambition	High (H)	Frontier rung (rung_sel: H)

The active baseline configuration sits at the low end of both axes at once — stated-policies demand and the current-practice circularity rung, referred to elsewhere in this document as the STEPS-L world. It is the only cell of the nine-cell grid that is also the calibrated historical continuation; the other eight are counterfactual by construction.

Axis 1 — Demand

Three demand trajectories are available, aligned to the International Energy Agency's published scenario framework: a stated-policies trajectory reflecting measures actually legislated, an announced-pledges trajectory reflecting commitments made but not yet enacted, and a net-zero trajectory reflecting a pathway consistent with limiting warming to 1.5 °C. In the model's configuration, these are named after the IEA scenarios they follow, and each maps to a matched pair of trajectories in the underlying data: a total-demand series and a stream-composition series.

Why the swap is joint. Stream-level demand in the runtime is the product of a total and a composition, so these two could in principle be varied independently. They are not, and the reason is central to what the model measures. A transition scenario is not simply more metal: it redistributes metal toward battery and clean-energy streams whose lifetimes, collection rates and recycling routes differ sharply from the legacy streams they displace. Holding composition fixed while raising the total would keep the end-of-life mix constant and so understate exactly the recycling and end-of-life divergence this study exists to quantify. One selector therefore swaps both factors coherently. Historical values are identical across all three scenarios for both factors; the trajectories diverge only from the reporting window onward, which prevents alternative-history paths from contaminating the calibrated historical stock. Table 50 gives the 2050 totals and the composition surface each scenario selects.

Table 50. *Total demand at 2050 by scenario, and the composition surface each selects. The active baseline configuration runs the stated-policies column.*

Material	Unit	Stated policies	Announced pledges	Net zero	Ordering
Copper	Mt	37.45	39.64	41.28	Increasing with ambition
Nickel	kt	6,337	7,230	7,612	Increasing with ambition
Cobalt	kt	390	453	503	Increasing with ambition

Material	Unit	Stated policies	Announced pledges	Net zero	Ordering
Lithium	kt	1,170	1,540	1,834	Increasing with ambition
Silicon	kt	16,263	15,700	14,935	Decreasing with ambition — see below
Composition surface selected	—	Low	Central	High	Battery and clean-energy stream shares rise with ambition

Silicon runs the other way, and this matters for how results are read. Four of the five materials show the expected ordering — more transition ambition means more material. Silicon does the opposite, and the reason is structural rather than an artefact. Silicon does not appear in the IEA's critical-minerals demand tables at all, so its scenario totals are built by decomposing demand into end uses and scaling each by its own driver. Roughly 42% of silicon demand is ferrosilicon, which tracks crude steel production, and a further large share is silicon metal entering aluminium alloys, which tracks aluminium production. Both steel and aluminium output *fall* under more ambitious transition pathways — crude steel at 2050 runs from 2,424 Mt under stated policies down to 1,925 Mt under net zero, a 26% spread in the opposite direction to solar. The steel and aluminium effect dominates the solar effect, so silicon demand decreases as ambition rises. Semiconductor and silicon are held scenario-invariant, being driven by general economic activity rather than by energy technology. The reporting consequence is stated here because the stated policies, low-circularity corner is silicon's *high*-extraction case, while it is the *low*-extraction case for the other four materials. Results must be read per material. Avoided extraction tonnes must never be summed across the five metals — they are physically incommensurable quantities moving in different directions.

Two further caveats attach to the demand axis. First, the sourced scenario design ends at 2050; the model's horizon in the active configuration ends there too, but the underlying data files extend to 2100 on a fitted continuation that is a modelling tail rather than a published result, and which produces an ordering inversion for cobalt beyond 2050. Any extension of the study horizon past 2050 must make an explicit choice about that tail rather than inheriting it. Second, the demand vintages are mixed: the totals for four materials come from the 2024 edition of the IEA critical-minerals outlook while the silicon decomposition and the stream-composition anchors cite the 2025 edition. This is accepted and documented, because re-sourcing the totals would move the baseline.

Axis 2 — Circularity ambition

Circularity ambition is expressed through three lever families, each corresponding to a distinct strategy for reducing primary material requirements, and each set at one of three ambition rungs. The levers introduce no new mechanics: they select values within fields the model already reads. Table 51 maps each lever to what it actually changes.

Table 51. *The three circularity levers, the mechanism each drives, and the deployed values at the moderate and frontier rungs.*

Lever	Strategy	What it changes	Moderate rung	Frontier rung
NARROW	Use less material per unit of service	A material-intensity multiplier reducing base demand before any recovery mechanism operates	0.88–0.95 at 2050, varying by demand column	0.82–0.88 at 2050, varying by demand column
SLOW	Keep products in service longer	A product-recovery gate: repair and maintenance lifetime stretch, direct reuse, and refurbishment/remanufacturing, each with its own diversion rate, screening pass rate and second-life service ratio	Moderate product-recovery surface	Frontier product-recovery surface
CLOSE	Recover more material at end of life	Imports the collection trajectory and the recycling-to-downcycling split from the target routing surface, then raises recycling process yield	Central collection and split; yield uplift +0.05	High collection and split; yield uplift +0.10

What the rungs mean. The lowest rung represents current practice: the low routing surface, with all other levers neutral — lifetimes unmultiplied, no transition policy, no yield uplift. It is not a no-circularity world. Today's economy already collects, remanufactures and recycles these metals, and the low rung encodes those observed volumes; the model's product-recovery surface at this rung carries non-zero diversion rates derived from the deployed routing data. The moderate rung represents legislated and pledged instruments — recycled-content floors, ecodesign requirements, critical-raw-materials legislation. The frontier rung represents literature upper bounds, subject to the feasibility caps discussed below.

Why the NARROW lever differs by demand column. A single material-intensity multiplier applied uniformly across all three demand columns would double-count. Each IEA demand pathway already embeds its own assumptions about material efficiency, and those assumptions are strongest in the most ambitious pathway. Applying the same additional reduction to all three would credit the model twice for efficiency the net-zero trajectory already contains. The lever therefore uses separate per-column paths, modelling only the increment beyond what each scenario already embeds — so the reduction is largest for the stated-policies column, which embeds least, and smallest for netzero, which embeds most. Table 52 gives the deployed values.

Table 52. *Material-intensity multiplier at 2050 by demand column and ambition rung. Lower values mean more material saved; the column ordering reflects how much efficiency each demand pathway already embeds.*

Demand column	Moderate rung	Frontier rung	Reading
Stated policies (STEPS)	0.880	0.820	Embeds least efficiency, so carries the largest additional reduction
Announced pledges (APS)	0.900	0.840	Intermediate
Net zero (NZE)	0.950	0.880	Embeds most efficiency, so carries the smallest additional reduction

The SLOW lever and its structure. The product-recovery lever is the most detailed of the three. It is authored as three named surfaces, each specifying, for each material stream where product recovery

is physically meaningful, three sub-mechanisms: repair and maintenance, represented as a uniform stretch of the first-life lifetime distribution; direct reuse, in which retiring units are diverted to a second life before any material recovery; and refurbishment or remanufacturing, in which units are diverted, remediated and then returned to service. Each recovery pathway carries a diversion rate, a return rate representing the share of diverted units passing condition screening, and a second-life service ratio representing how much service a recovered unit delivers relative to a new one. A price-response amplifier scales diversion rates with the price signal at the two upper rungs, bounded by explicit clamps. Table 53 gives the deployed values for a battery stream and an industrial stream as worked examples; the complete surface spans twelve material-stream blocks across the three rungs.

Table 53. *Product-recovery parameters for two representative streams across the three ambition rungs. Diversion rates are shares of first-life retirements entering the recovery gate, shown at 2050.*

Stream	Parameter	Current practice	Moderate	Frontier
Cobalt battery sulfate	Lifetime stretch factor	1.00	1.10	1.20
Cobalt battery sulfate	Reuse diversion rate	0.040	0.080	0.150
Cobalt battery sulfate	Reuse return rate	0.55	0.58	0.62
Cobalt battery sulfate	Refurb/reman diversion rate	0.020	0.050	0.250
Cobalt battery sulfate	Refurb/reman return rate	0.60	0.63	0.68
Cobalt battery sulfate	Second-life service ratio	0.65	0.65	0.65
Cobalt battery sulfate	Price-response amplifier	0.0 (off)	0.075	0.15
Copper alloy grade	Lifetime stretch factor	1.00	1.05	1.10
Copper alloy grade	Reuse diversion rate	0.002	0.005	0.010
Copper alloy grade	Refurb/reman diversion rate	0.048	0.075	0.140
Copper alloy grade	Second-life service ratio	0.70	0.70	0.70
Copper alloy grade	Price-response amplifier	0.0 (off)	0.05	0.10

Two features of that table carry interpretive weight. The second-life service ratio is held constant across rungs within a stream but varies substantially between materials — 0.70 for copper, 0.65 for nickel and cobalt, 0.55 for silicon, and 0.50 for lithium — because it encodes a physical property of what a recovered unit can still do, not an ambition level. A second-life lithium battery in stationary storage delivers roughly half the service of a new one because its usable capacity has degraded. The price-response amplifier is off entirely at the lowest rung, which preserves the property that the baseline is invariant to the price signal on this channel, and its per-material values reflect how responsive each recovery market plausibly is: cobalt battery take-back is the most price-elastic stream in the set, lithium the least because regulatory mandate rather than price is expected to drive its recovery.

The evidence behind these values is of mixed strength. Table 54 sets out the basis for each parameter family together with an explicit confidence label, so that a reader can see which values rest on published measurement and which are structural assumptions chosen to be plausible.

Table 54. *Evidence basis and confidence for each product-recovery parameter family.*

Parameter family	Deployed values	Evidence basis	Confidence
Second-life service ratio	Copper 0.70, nickel 0.65, cobalt 0.65, silicon 0.55, and lithium 0.50.	Measured performance of recovered units. Copper from a collection-weighted assessment across motor rewind, heat exchangers, cable and electronics [JEROME-25];[123] nickel and cobalt from a review of battery, stainless and aerospace recovery [NERI-24];[124] lithium additionally reflects capacity derating in stationary second-life duty [NERI-24], [IEA-GEO-21];[125] silicon from photovoltaic early-decommission analysis [IEA-PVPS-12][126]	High — anchored on measured recovered-unit performance
Return rate (screening pass rate)	Reuse 0.55–0.62; refurbishment/remanufacturing 0.60–0.68	State-of-health screening thresholds and remanufacturability assessment: electric-vehicle batteries retire at roughly 80% state of health, with 55–60% of components found remanufacturable [NERI-24]	Medium-high — a single review is the primary anchor
Diversion rate, current-practice rung	0.001–0.048 by stream	Derived arithmetically from the deployed routing data as the product of the low-surface collection rate and remanufacturing rate, so it inherits that artefact's provenance rather than adding an independent claim	Medium — as good as the underlying routing data
Diversion rate, upper rungs	Up to 0.150 reuse and 0.250 refurbishment for battery streams	Structural bounds set to represent industry-scale take-back programmes and frontier ambition. Not fitted to observed recovery volumes, and no published collection-to-remanufacturing elasticity was located to anchor them	Medium — plausible bounds, not measured rates
Lifetime stretch factor	1.00 to 1.20 by stream and rung	Bracketed against single-product lifetime-extension evidence: appliances roughly ×1.5, vehicles effective lifetime more than doubled, buildings where over 90% of frontier material savings are attributable to lifetime extension [SANDEZ-24],[127] [MATECON-18],[128] [IEA-ME-19].[129] Deployed values sit well below those figures because the lever is a stock-weighted fleet multiplier, not a per-product factor — only a fraction of units in a stream are repaired or reused	Medium — conservative against the bracket, but the evidence concerns appliances, buildings and vehicles rather than these five metals, and rebound effects are documented and not modelled [BAKKER-21][130]
Price-response amplifier	0.0 under current practice; copper 0.05/0.10, nickel 0.04/0.08, cobalt 0.075/0.15, lithium 0.03/0.06, and silicon 0.0 at the upper rungs.	Structural analogue to the remanufacturing price response in the endogenous circular-loop surface. Per-material ordering reflects judged market responsiveness: cobalt battery take-back most price-elastic, lithium least because regulatory mandate rather than price is expected to drive recovery	Medium — structural bounds, explicitly not fitted coefficients

One caveat applies across the whole circularity axis rather than to any single parameter. Much of the frontier evidence in the circular-economy literature concerns bulk materials — steel, aluminium,

[123] Jerome, A., et al. (2025). Copper product reuse across motors, heat exchangers, cable and electronics. *[Journal details as recorded in the model's scenario evidence register; not independently verified for this document.]*
[124] Neri, A., et al. (2024). Remanufacturability and state-of-health screening of recovered units. *[Journal details as recorded in the model's scenario evidence register; not independently verified for this document.]*
[125] International Energy Agency. (2021). *Global EV outlook 2021*. IEA Publications.
[126] International Energy Agency, Photovoltaic Power Systems Programme. (2024). *End-of-life management of photovoltaic panels* (IEA-PVPS Task 12). IEA-PVPS.
[127] Sandez, S., et al. (2024). A review of product lifetime extension. *Sustainable Production and Consumption, 51*, 445. Incorporating per-product figures from Iraldo (2017), Baxter (2019) and Jerome (2023).
[128] Material Economics. (2018). *The circular economy: A powerful force for climate mitigation*. Material Economics Sverige AB.
[129] International Energy Agency. (2019). *Material efficiency in clean energy transitions*. IEA Publications.
[130] Bakker, C. A., Mugge, R., Boks, C., & Oguchi, M. (2021). Understanding and managing product lifetimes in support of a circular economy. *Journal of Cleaner Production, 279*, 123764. doi:10.1016/j.jclepro.2020.123764

cement — rather than the five metals modelled here. Transferring it involves a modelling assumption about whether recovery behaviour observed for structural materials carries over to critical raw materials in dispersed applications. That modelling assumption is stated rather than tested and represents a material source of uncertainty in the circularity axis.

The wedge ladder: Making results attributable

Knowing that a frontier-circularity cell uses less primary material than a current-practice cell is of limited use if the reduction cannot be attributed to individual interventions. Each reported cell is therefore constructed as the endpoint of a cumulative ladder rather than as a single configured state. Starting from the current-practice baseline for a given demand column, the levers are added one at a time: first the material-intensity reduction, then the product-recovery gate, then the end-of-life bucket. Each rung is a real, executed run, so the difference between consecutive rungs is the contribution of the lever that was added.

Conservation holds by construction: because the rungs are cumulative, the individual lever contributions sum exactly to the difference between the final cell and the baseline. This is verified as a blocking gate for cumulative primary supply and cumulative secondary supply across all six moderate and frontier cells — a build fails rather than reporting an unattributable result.

Two limits on how the decomposition may be read. First, the attribution is order-conditional. Because the endogenous circular loops respond to the state the previous lever left behind, adding the end-of-life lever to a system that has already had its demand reduced produces a different increment than adding it first. The ladder order is fixed and documented, so the decomposition is consistent across cells and comparable between them, but the individual increments are contributions given that order rather than context-free properties of each lever. Second, decomposition is valid only within a fixed demand column. Comparing a rung of the net-zero ladder against a rung of the stated-policies ladder confounds the circularity effect with the demand effect and is not a wedge.

Axis 3 — Shocks and resilience

The third axis tests behaviour under disruption. Four shock channels are defined, each applying a multiplier to one model quantity over a window beginning in 2025. Each is run in two durations: a sustained form lasting ten years, and a pulse form lasting three, so that the difference between a prolonged constraint and a sharp one can be separated. The sustained form leaves a sixteen-year recovery window within the study horizon and the pulse form twenty-three.

Three of the four channels are grounded per material rather than applied uniformly, which is the most consequential feature of the shock design. A uniform supply cut would imply that each material's supply chain is equally vulnerable, an assumption that is not substantiated for the materials assessed in this study. The supply-cut multipliers are instead read from published supply-concentration analysis measuring what fraction of demand could still be met if the single largest supplier were removed. Table 55 gives the deployed values and Table 56 the evidence behind each.

Table 55. *Shock channels and deployed multipliers. Sustained shocks run ten years from 2025; pulse shocks run three. Values below one are reductions while above one are increases.*

Channel	Reading	Copper	Nickel	Cobalt	Lithium	Silicon	Basis
Supply cut	Refined-output disruption	0.96	0.53	0.65	0.67	0.70	Supply concentration: fraction of demand met with the largest supplier removed
Demand surge	Accelerated transition pull	1.06	1.16	1.23	1.39	1.00	Spread between demand pathways per material
Capacity loss	Refining bottleneck	0.81	0.82	0.69	0.75	0.62	Refining-stage concentration by material
Collection collapse	End-of-life system failure	0.75	0.75	0.75	0.75	0.75	Uniform; not re-grounded per material

The per-material spread is large and carries the analytical content of this axis. Copper's supply cut is barely a shock at 0.96, because copper production is the least concentrated of the five — that mildness is itself a finding rather than a modelling artefact. Nickel's at 0.53 is the most severe, reflecting the concentration of refined nickel supply in a single country. Silicon is the one material with no demand surge — its multiplier is held at 1.00, consistent with its demand falling rather than rising under transition ambition, so a transition-pull surge would be wrongly signed for it.

Table 56. *Derivation and evidence for each shock channel's magnitudes.*

Channel	How the magnitude is derived	Evidence	Status
Supply cut	Read directly from published supply-concentration analysis: the multiplier is the share of demand still met when the single largest supplier is removed, at 2035 under stated policies	[IEA-GCMO-25],[131] supply-concentration chart. Copper 96%, nickel 53% (reflecting Indonesian concentration), cobalt 65%, lithium 67%	Grounded for four materials. Silicon is not covered by the concentration analysis and retains the design default of 0.70 — a stated residual gap, not a sourced value
Demand surge	The ratio between the net-zero and announced-pledges demand trajectories at 2030, taken from the model's own scenario data — a credible one-step acceleration rather than an arbitrary uplift	Derived from the held demand trajectories [IEA-GCMO-24].[132] Ordering independently corroborated by observed growth rates: lithium near 30% year-on-year against nickel and cobalt at 6–8% [IEA-GCMO-25]	Grounded. Silicon correctly receives no surge
Capacity loss	Computed as one minus the dominant refiner's production share multiplied by a single governed outage depth of 0.40, applied uniformly across materials	Dominant-refiner shares cross-checked between two independent compilations and found to agree: copper 47%, nickel 45%, lithium 62% (chemical), cobalt 77%, silicon 95% for high-purity polysilicon [IEA-GCMO-25]. Mechanism	Structure grounded and cross-checked; the outage depth is a governed assumption. This is silicon's best-grounded shock, refining being its most concentrated stage

[131] International Energy Agency. (2025). *Global critical minerals outlook 2025*. IEA Publications. Base-case total-supply projections.
[132] International Energy Agency. (2024). *Global critical minerals outlook 2024*. IEA Publications. Accompanying Critical Minerals Data Explorer workbook.

Channel	How the magnitude is derived	Evidence	Status
		evidenced by named episodes including a four-month cobalt export suspension and pandemic-era smelter halts	
Collection collapse	Retained as a uniform stress rather than re-derived per material	Bracketed by two independent anchors: EU waste-electronics collection fell from 48.6% (2019) to 37.5% (2023), a ratio of 0.77 [EEA-WEEE];[133] and the low-to-central ratio in published battery-metal recycled-share sensitivity is near 0.70 [IEA-RCM-24].[134] Global formal collection is projected to decline further to 2030 [GEM-24][135]	Design-level stress at the severe end of an observed decline. Uniformity should be read as a default, not as evidence that end-of-life systems are equally fragile everywhere. Battery collection is legislated upward, so a collapse is tail risk rather than a central projection

Two channels warrant a caution about combination. The supply cut and the capacity loss both reduce refined availability and are correlated, but they represent distinct events — one removes a supplier net of that supplier's own demand, the other derates throughput at the dominant refiner. They are never applied within the same scenario, and their effects must not be added together to construct a combined disruption.

What the resilience test asks. The headline question is not whether a shock can impact, but whether higher circularity ambition confers faster recovery. Each shocked run is paired with an unshocked run on the same host cell, and the comparison of interest is whether the frontier-circularity host recovers within the horizon and does so faster than the current-practice host. This is posed as a finding to be reported rather than a result assumed in advance.

An important constraint attaches. The recovery window is short relative to the shock magnitude, so a substantial share of recovery-time measurements are expected to be censored — meaning recovery has not completed by the end of the horizon and the true recovery time is unknown, only bounded below. Censored values are reported explicitly as censored. They must not be read as evidence of recovery, and if all values for a given comparison are censored, the correct response is to revisit the shock window rather than to reinterpret the output.

The executed run surface

The scenario architecture describes what is expressible. What is actually executed and reported is a curated manifest of sixty-nine runs, all at the central primary-supply-elasticity setting documented in the Calibration section. Table 57 gives its composition.

Table 57. *Composition of the executed run manifest.*

[133] European Environment Agency & Eurostat. (2019–2023). *Waste electrical and electronic equipment collection rates* [Data set].
[134] International Energy Agency. (2024). Recycling and secondary supply. In *Global critical minerals outlook 2024*. IEA Publications.
[135] International Telecommunication Union & UNITAR. (2024). *The global e-waste monitor 2024*. International Telecommunication Union & UNITAR.

Block	Content	Rows
Grid	Three current-practice baselines plus six moderate and frontier final cells	9
Ladder	The two intermediate rungs for each of the six moderate and frontier cells	12
Sustained shocks	Four channels on each of six hosts (three demand columns × current-practice and frontier rungs)	24
Pulse shocks	The same four channels and six hosts in the three-year pulse form	24
Total	—	69

A separate family of diagnostic scenarios exists outside this manifest and is not reported as a result. These are ablations that isolate individual mechanisms by switching them off — endogenous loops inert, strategic reserve disabled, historical warm-start removed, warm-start present but without event windows — together with two endpoint sensitivities holding the primary-supply elasticity at zero and at one. They exist to answer the question of *what this mechanism contributes*, and are run manually when that question arises rather than as part of the reported surface.

Feasibility bounds and an open finding

The frontier rung is bounded by literature-derived feasibility caps on how much secondary material can realistically supply demand by 2050. These are output bounds checked after a run rather than lever inputs — the levers are set on the input side, and the caps ask whether the resulting realised secondary share is credible. Table 58 gives them together with the evidence behind each.

Table 58. *Frontier-rung feasibility caps on the 2050 realised secondary share, with the binding physical constraint and supporting evidence for each material.*

Material	Cap	Stress upper	Binding constraint	Evidence
Copper	40%	50%	Collection, sorting, and secondary smelter capacity	Secondary supply excluding manufacturing scrap approaching 40% by 2050 under announced pledges, with end-of-life scrap reaching roughly 70% of demand [IEA-RCM-24]; independent long-run cycle modelling at 37.5% [WATARI-21];[136] end-of-life recycling rate assessed at 43–53% [UNEP-RR][137]
Nickel	40%	42%	Battery fleet lag; the stainless loop is already mature	Battery secondary share of 30% under net zero [IEA-RCM-24]; long-run cycle modelling at 42.0% [WATARI-21]; recycled content assessed at 29–41% and end-of-life recycling at 57–63% [UNEP-RR]
Cobalt	40%	40%	End-of-life pool from electric vehicles and storage; chemistry	Battery secondary share of 40% under net zero and approximately 30% under announced pledges [IEA-RCM-24]. Cap and stress upper

[136] Watari, T., Nansai, K., & Nakajima, K. (2021). Major metals demand, supply, and environmental impacts to 2100: A critical review. *Resources, Conservation and Recycling, 164*, 105107. doi:10.1016/j.resconrec.2020.105107

[137] Graedel, T. E., et al. *Recycling rates of metals: A status report*. United Nations Environment Programme, International Resource Panel. [As compiled in the project's circularity calibration evidence register; the underlying report was not re-verified directly for this document.]

Material	Cap	Stress upper	Binding constraint	Evidence
			substitution away from cobalt	coincide, so there is no headroom above the central bound
Lithium	30%	35%	Young fleet, so little material is available for recovery before 2040	Battery secondary share of 30% under net zero, with a high-collection sensitivity of 25–35% [IEA-RCM-24]; historical recycling assessed at below 1% [UNEP-RR]
Silicon	20%	20% unless non-photovoltaic recovery is sourced	Photovoltaic waste availability and the difficulty of high-purity recovery	Photovoltaic waste reaching 80–89% of new module mass by 2050 [IRENA-16];[138] input-silicon recovery of 66–98% at module level [IEA-PVPS-12]; high-purity silicon not widely recycled at scale [IEA-RCM-24]. The stress upper equals the cap because no non-photovoltaic recovery route has been sourced

The metric these caps are measured on is not the model's headline recycling indicator. This distinction is an important variable and is worth stating precisely. The caps are expressed on an end-of-life-only basis — post-consumer material exclusively. The model's recycled-input-rate indicator and its secondary-supply series both *include* manufacturing scrap, which accounts for up to roughly 63% of copper and 76% of nickel secondary supply. Comparing the model's recycled-input rate directly against these caps therefore compares two different quantities and will show breaches that are artefacts. The cap-consistent quantity subtracts manufacturing scrap from secondary supply before taking the share; on that basis, copper and nickel sit well under their caps.

An open feasibility finding. On the correct metric, and under the exact-import construction described above, cobalt and lithium exceed their stress upper bounds in the frontier cells. This is not a double-counting failure — the no-double-count verification passes — and it is not an artefact of the wrong metric. It is one of the consequences of importing the high routing surface exactly: the collection and recycling-split trajectories that surface specifies, applied to these two materials, imply a realised secondary share above what the feasibility review considered defensible. Resolving it would require either reducing the imported surface, which would break the exact-frontier construction the grid cells depend on, or revising the caps. The project documentation identifies this as an unresolved configuration decision for the modelling team, subject to the applicable technical review and publication sign-off process. At the time of writing, the decision remains open. . Frontier-rung cobalt and lithium results should therefore be read as an upper bound whose feasibility is contested by the project's own review, not as a validated projection.

Sources for the scenario evidence base

The sources below support the scenario design. As with the preceding sections, they are listed separately from the exogenous-data and calibration registers; where a source appears in more than one, the key is the same. Table 59 lists them.

Table 59. *Source register for the scenario design evidence base.*

[138] International Renewable Energy Agency & IEA-PVPS. (2016). *End-of-life management: Solar photovoltaic panels*. IRENA.

Key	Source	Settings supported
[IEA-GCMO-24]	IEA, Global Critical Minerals Outlook 2024, and its Critical Minerals Data Explorer workbook	Total-demand milestones for copper, nickel, cobalt and lithium under all three scenarios
[IEA-GCMO-25]	IEA, Global Critical Minerals Outlook 2025	Stream-composition anchors; supply-concentration data underlying the per-material supply-cut multipliers
[IEA-WEO-24][139]	IEA, World Energy Outlook 2024, Annex A.5	Crude steel and aluminium output by scenario, driving the silicon end-use decomposition
[IEA-RCM-24]	IEA, recycling and secondary-supply analysis within the critical-minerals outlook series	Frontier-rung realised secondary-share caps for all five materials
[WATARI-21]	Watari, T. et al. (2021), on global metal cycles and long-run secondary supply potential	Copper and nickel realised-share cap corroboration
[HERTWICH-20][140]	Hertwich, E. et al. (2020), UNEP International Resource Panel, Resource Efficiency and Climate Change	Frontier-rung ambition bounds for the circularity levers
[MATECON-18]	Material Economics (2018), The Circular Economy: A Powerful Force for Climate Mitigation	Frontier-rung circular-economy ambition bounds
[IRENA-16]	IRENA (2016), End-of-Life Management: Solar Photovoltaic Panels	Photovoltaic waste availability underlying the silicon feasibility cap
[IEA-PVPS-12]	IEA PVPS Task 12 (2024 update), PV End-of-Life Management	Silicon recovery rates and second-life service ratio for the solar stream
[NERI-24]	Neri, A. et al. (2024), on remanufacturability and state-of-health screening of recovered units	Return rates and second-life service ratios for battery, stainless and aerospace streams
[JEROME-25]	Jerome, A. et al. (2025), on copper product reuse across motors, heat exchangers, cable and electronics	Copper second-life service ratio of 0.70
[IEA-GEO-21]	IEA, Global EV Outlook 2021	Lithium second-life service ratio, reflecting capacity derating in stationary duty
[IEA-CET-21][141]	IEA, Clean Energy Transitions programme analysis (2021)	Embedded material-efficiency overlap by scenario, underlying the per-column material-intensity paths
[IEA-ME-19]	IEA, Material Efficiency in Clean Energy Transitions (2019)	Frontier ceiling bounding the material-intensity lever; buildings lifetime-extension attribution underlying the repair and maintenance parameters
[UNEP-RR]	UNEP International Resource Panel, metal recycling-rates assessment, as compiled in the project's circularity calibration evidence register; the underlying report was not re-verified directly for this document	End-of-life recycling and recycled-content ranges corroborating the copper, nickel and lithium feasibility caps

[139] International Energy Agency. (2024). *World energy outlook 2024*. IEA Publications.
[140] International Resource Panel. (2020). *Resource efficiency and climate change: Material efficiency strategies for a low-carbon future*. United Nations Environment Programme, International Resource Panel.
[141] International Energy Agency. (2021). *Clean energy transitions programme: Annual report*. IEA Publications.

Key	Source	Settings supported
[SANDEZ-24]	Sandez, S. et al. (2024), Sustainable Production and Consumption 51:445, review of product lifetime extension; incorporating per-product figures from Iraldo (2017), Baxter (2019) and Jerome (2023)	Single-product lifetime-extension bracket against which the repair and maintenance stretch factors are set
[BAKKER-21]	Bakker, C. et al. (2021), Journal of Cleaner Production 279:123764, on rebound in product-lifetime extension	Documented rebound caveat on the lifetime stretch parameters; rebound is not represented in the model
[EEA-WEEE]	European Environment Agency and Eurostat, waste electrical and electronic equipment collection rates, 2019–2023	Observed collection decline of 48.6% to 37.5% bracketing the collection-collapse shock magnitude
[GEM-24]	Global E-waste Monitor (2024)	Projected decline in global formal collection to 2030, supporting the collection-collapse shock as a plausible stress

Assumptions and limitations

Sets out what the model deliberately simplifies, where its evidence is weakest, which of its current states are temporary, and how its outputs should and should not be read.

How this section is organised

A model of this scope embeds a large number of choices, and they are not all of the same kind. Some are deliberate simplifications made because it was anticipated the added resolution would not change the conclusions and would cost transparency. Some are limitations of the available evidence. Some are temporary states of the current configuration with an identified route out. And some are not limitations of the model at all, but ways its outputs can be misread. Grouping them together as a single undifferentiated list of caveats would obscure the difference between a considered design decision and a known weakness, so they are separated here into four categories, treated in turn.

Two general cautions apply across all four. First, each element below describes the configuration documented in this report. The model's framework is more general than the configuration it is currently run in — several capabilities described as inactive here are implemented and executable, and a limitation that follows from the active configuration is not necessarily a limitation of the model. Where that distinction matters, it is stated explicitly. Second, the items below are the ones that materially affect how results should be read; they are not an exhaustive inventory of each simplification in the code.

Deliberate structural simplifications

The following are choices, not oversights. In each case, a more resolved representation was possible and was not adopted, either because the data to support it does not exist at global scale or because it was anticipated the added complexity would not change what the model is used to conclude. What matters for the reader is knowing which questions each simplification puts out of reach. Table 60 pairs each simplification with the questions it forecloses.

Table 60. *Structural simplifications in the deployed model, and the questions each puts out of reach.*

Area	What is simplified	What this means for interpretation
Productive capacity	Capacity is a single aggregate envelope index per material rather than explicit stage-specific capital stocks with their own construction lead times and retirement schedules	The model can say that capacity is binding and how quickly it adjusts, but not which specific stage — mine, smelter, refinery — is the constraint, nor what investment would relieve it
Collection systems	Collection is a single aggregated rate per stream, bounded and lagged in its response, but not resolved by infrastructure type, geography or collection channel	Collection ambition can be varied and its consequences traced, but the model cannot distinguish kerbside from deposit-return from industrial take-back, and cannot cost a collection intervention
Secondary inventory	Recovered material is buffered in a refinery stockpile with a release-rate	Stockpile behaviour is a smoothing device, not a representation of how traders actually hold and release scrap in response to price

Area	What is simplified	What this means for interpretation
	control rather than an endogenous market-clearing inventory policy	
Strategic reserve	Strategic inventory is an aggregate stock per material with no commodity-quality or form detail, and its fill and release depend strongly on configured threshold settings	Reserve results are sensitive to policy thresholds that are modelling choices rather than observed policy; they indicate direction and rough magnitude, not the behaviour of any actual stockpile programme
Demand response	Demand responds to scarcity through a price-elasticity loop, with no representation of material substitution or technology switching	The model cannot represent a material being designed out in response to sustained scarcity — arguably the most important real-world response to a critical-materials constraint, and one entirely absent here
Behavioural heterogeneity	Differences between materials in the behavioural layer are supplied by configuration rather than learned from data within the run	Per-material behavioural differences are asserted from evidence and calibration, not emergent; they will not adapt if the underlying market structure changes over the horizon
Unmet demand	When supply for a specific material stream is short, unmet service remains attached to that stream — demand does not shift to another stream or product	Stream-specific shortfalls should be read as indicative of pressure on that stream, not as a forecast that the shortfall will physically materialise; in reality some substitution across grades would occur
Product structure	A single placeholder product category is active; the product and component detail structure exists in the framework but is switched off	Results cannot be attributed to end-use sectors. This is a configuration state rather than a framework limit, but it holds for each result reported here
Source and route detail	Deposit type and processing route both collapse to a single category carrying unit share and unit grade	The model cannot currently distinguish sulfide from laterite nickel, or brine from hard-rock lithium, and cannot represent ore-grade decline through the deposit axis. The upstream chain is structurally live but carries no differentiating information

The demand-response and source-route rows deserve emphasis over the others. Substitution is the principal mechanism by which economies have historically escaped material constraints, and its absence means the model's scarcity results should be read as describing pressure under a no-substitution assumption rather than as projections of what could actually happen. The source-and-route collapse means that a headline result about primary supply is a result about aggregate refined availability, not about geology.

Evidence and data limitations

Where the previous subsection concerns what the model represents, this one concerns how well the numbers behind that representation are known. Table 61 gives the limitations that most affect confidence in results, with the direction of the risk in each case.

Table 61. *Evidence limitations affecting confidence in results.*

Area	Limitation	Effect on results
Upstream yields	Primary ore-input anchoring is highly sensitive to the quality of stage-yield data, which is thinner and less consistently reported than production statistics	Errors in stage yields propagate directly into modelled primary availability; the ore-to-metal conversion is among the least certain quantities in the model
Silicon forward supply	Silicon has no published mine-supply trajectory in the source used for the other four materials. Its forward path retains an inherited support series with an assumed growth rate beyond 2040	Any conclusion about silicon supply adequacy is contingent on that assumed pathway rather than on a sourced projection, and should be presented with that dependence stated
Refining-boundary representation	Availability at the refining boundary is a sum of domestic availability and net trade terms, which can miss unmodelled channels and material-quality effects	Refined availability is an approximation of a physically more complex boundary; small discrepancies against reported refined production are expected
Circular-loop elasticities	Limited historical price variation in secondary-metal markets under-determines the price and core-availability elasticities in the endogenous loops. They are set against literature budgets and structural priors rather than freely estimated	Endogenous-loop responses should be read as ranges and compared against the exogenous-anchor baseline, never as point estimates
Quality representation	Material quality in the downcycling cascade is a single scalar index per material — a reduced form of a metallurgical state that is genuinely stream- and alloy-specific	Captures aggregate direction and order of magnitude of quality degradation, not within-material heterogeneity. Tramp-element behaviour in specific alloy systems is out of reach
Remanufacturing representation	Remanufacturing is driven by mass flow as a proxy for what are really unit-level and condition-level decisions, with core suitability derived from cohort retirement age	Does not distinguish product condition, certification status or sector-specific restorability; interpret remanufacturing results as a range
Circularity frontier evidence	Much of the frontier circular-economy literature used to bound the ambition rungs concerns bulk materials rather than these five metals	The frontier rung embeds a transferability modelling assumption that is stated but not tested — the single largest uncertainty in the circularity axis
Cobalt behavioural response	Cobalt's behavioural parameters were selected from sensitivity analysis under a deliberately conservative stance rather than from historical calibration of the kind applied to copper and nickel	Cobalt deficit results carry a different evidentiary status from the other materials and should be compared against observed substitution, expansion and recycling evidence before being relied on
Lithium mid-century scope	The lithium demand series is civilian apparent consumption, with the United States lithium-6 stockpile programme excluded as a separate reservoir, but the stream-composition shares are transcribed from a benchmark end-use split that includes that defence use — the 1960 specialty share of 0.54 is set by a demand component the model does not carry	Mid-century lithium is over-routed into the specialty stream, whose retirement profile is in turn anchored on a sub-mix dominated by material outside the modelled system. Lithium stream-level results across roughly 1955-1985 should not be relied on; the tonnages are small beside the battery era, and the inconsistency is unresolved
Calibration observation coverage	The observed series available for evaluating the model cover roughly a quarter of the full artefact grid, and for three of the five	Historical agreement is genuinely testable for copper and nickel; for silicon, cobalt

Area	Limitation	Effect on results
	materials the in-use stock comparator is derived from the model's own earlier output rather than sourced independently	and lithium the stock-level comparison is not independent evidence

Temporary states of the current configuration

The items in Table 62 differ from those above in that each has an identified route out. They are states of the configuration described in this report, not properties of the model, and each would be resolved by supplying data or making a governance decision rather than by changing the model's structure.

Table 62. *Temporary states of the deployed configuration, with the route to resolving each.*

Item	Current state	Route out
Source and route allocation	Deposit share, ore grade, processing-route share and product share are each populated with a single canonical category at unit value, so the upstream allocation chain is structurally live but carries no differentiating information	Supply sourced deposit-type and processing-route distributions per material, then re-run validation and the benchmark gates
Stock comparator for three materials	The observed in-use stock series used to check silicon, cobalt and lithium is derived from an earlier baseline run of the model, present so the warm-start gate can evaluate all five materials consistently	Replace those rows with curated external historical stock estimates and re-run the warm-start report; until then, treat those three gate results as non-independent
Solver convergence in two extreme cells	The inner coupling loop reaches its 0.02 tolerance within the 30-iteration cap for each reported cell except two extreme corners — net-zero demand at the frontier circularity rung, under a sustained supply-side shock, in the supply-capped regime. There, a strategic-reserve release relay leaves a residual of roughly 0.05, some two and a half times the tolerance. The affected cells are shock scenarios for cobalt and lithium; the residual is a property of the relay's discrete switching rather than of the smoothing factor, and is not removable by tuning the solver settings.	Replacing the discrete fill and release relay with a continuous or rate-limited response would make the loop contractive in these corners. Until that change is made, indicator values for the affected shock cells should be read as approximate, and the coverage and release paths in particular should not be interpreted at fine resolution. A cobalt convergence re-check is also outstanding against the current remanufacturing-on baseline.
Scenario packaging	Scenario classes for the grid, ladder, shock and pulse blocks ship and execute through dedicated run entrypoints, but the workflow-preset layer still exposes only a baseline-only option	A governance decision on whether scenario packs should be exposed as selectable workflow presets, or remain accessible through run entry points and the batch manifest as at present
Feasibility of the frontier rung for two materials	Under the exact frontier-import construction, cobalt and lithium exceed the stress upper bounds of their realised secondary-share feasibility caps	A decision between relaxing the exact-import construction and revising the caps; until taken, frontier-rung results for these two materials are an upper bound whose feasibility is contested

Item	Current state	Route out
Data provenance framework	A source hierarchy, promotion rules for provisional data, and evidence thresholds for material and process assumptions are identified as needed but not yet formalised	Establish and apply the framework, then record promoted provenance for each artefact

How results should and should not be read

The final category is not a limitation of the model but a set of ways its outputs invite misreading. Each has caused a real error at some point in the model's development, which is why they are recorded here rather than left to the reader's assumption. Table 63 sets them out.

Table 63. *Interpretation cautions.*

Caution	Detail
Terminal inventories are stock, not loss	End-of-horizon refinery stockpile and strategic inventory balances are material held in inventory. They are not losses, and must not be netted against supply or counted as unmet demand when closing a mass balance at the final year
Never sum avoided extraction across metals	The five materials move in different directions under transition ambition — silicon's demand falls as the others rise — and their tonnages are physically incommensurable. Aggregating avoided tonnes across metals produces a number with no physical meaning. Report per material
The recycling indicator is not the feasibility metric	The model's recycled-input rate includes manufacturing scrap, which reaches roughly 63% of copper and 76% of nickel secondary supply. The feasibility caps are defined on post-consumer material only. Comparing the two directly will show breaches that are artefacts of the mismatch
Censored recovery times are not recovery	Where a shocked run has not recovered by the end of the horizon, the recovery time is censored — bounded below, but unknown. Censored values must be reported as censored. If all values in a comparison are censored, the shock window needs revisiting rather than the output reinterpreting
Wedge decomposition is order-conditional and column-bound	Lever contributions are contributions given the fixed ladder order, because the endogenous loops respond to the state each preceding lever leaves. Decomposition is valid only within a single demand column; comparing rungs diagonally across columns confounds the circularity effect with the demand effect
Shock channels are not additive	The supply-cut and capacity-loss channels both reduce refined availability and are correlated, but represent distinct events. They are never applied together, and their effects must not be summed to construct a combined disruption
Resilience metrics depend on a configured threshold	Service-level resilience scalars are computed against a configured service threshold. Changing that threshold changes the metrics; they are comparative measures across scenarios, not absolute statements about adequacy
Scenario results are conditional on the demand pathway	Each circularity result is conditional on its demand column. A frontier-circularity outcome under one demand pathway is not evidence about what circularity achieves under another, which is why the full grid is run rather than a diagonal

Caution	Detail
Trade is not represented in these results	The origin-destination trade layer is implemented but disabled in the configuration reported here. All results are single-region global aggregates, and none of them describes inter-regional dependency, supplier concentration or trade exposure

A final point applies to the horizon itself. The configuration reported here runs from 1950 to 2050, which is shorter on both ends than the model's full authored data range. The shortened start compresses the historical run-up relative to the longer canonical horizon, and where time-window sensitivity matters, a comparison against the longer baseline is the appropriate check. The 2050 end is where the sourced scenario design stops; the underlying data files extend further on a fitted continuation that is a modelling artefact rather than a published projection, and any extension of the reporting horizon would need to address that explicitly rather than inherit it.

Verification: What would catch an error

The above concerns the model's design and evidence. A separate question is whether an error introduced into the model — an inappropriate merge, a mis-specified parameter, an accidental regression — would actually be caught before it reached a reported result. This matters because the model has a documented history of both outcomes: a cobalt stream-composition regression was caught by exactly this kind of check and now carries dedicated safeguards against recurrence, while the controls then in place did not identify a cobalt lifetime divergence or several stale claims in the reference documentation. Table 64 sets out the automated checks that exist, distinguishing between checks that prevent a defective artefact or run from completing and those that flag a finding for review without blocking completion.

Table 64. *Automated verification checks in the deployed model, and the consequence when one fails.*

Check	What it protects	When it runs	Consequence of failure
Mass-balance residual	Physical conservation within a single model run, checked against a configurable tolerance	Live, inside each run	The run raises an error and stops. A result cannot be reported from a run that failed this check
Demand-scenario ordering	The requirement that the three demand pathways be internally consistent in direction for four materials, with silicon's reversed ordering checked against its own headline-consistency bound instead	When the exogenous demand data is generated, not per model run	Data generation raises an error and stops; a violating demand file cannot be produced or shipped
Routing and stream-composition benchmarks	The six per-material recycling-rate reconciliation gates and the twelve stream-composition gates described in the Data section, including the three sentinel gates added specifically to prevent the cobalt composition regression from recurring	When the exogenous routing and composition data is generated	Data generation raises an error and stops; a violating file cannot ship. This is why the composition and routing artifacts can be asserted as internally consistent rather than merely plausible

Check	What it protects	When it runs	Consequence of failure
Wedge conservation	The requirement that summed lever contributions equal the total difference between a scenario's final cell and its baseline, for each moderate and frontier grid cell	Run as a verification script across the six cells after they are executed	A documented pass is required before wedge-decomposed results are reported; the check does not run automatically inside each individual model execution
No double-counting between exogenous levers and endogenous loops	The design requirement that circularity ambition is expressed only through inputs to the endogenous loops, never by separately writing their outputs	Run as a verification script and confirmed alongside the routing-import construction change described in Scenarios	A documented pass is required before the affected construction is treated as valid; not a live per-run gate
Frontier-rung feasibility caps	Whether a scenario's realised secondary-material share is credible against literature-derived upper bounds	Checked after a run completes, on the corrected ex-direct-scrap metric	Produces a flagged finding, not a block. This is why the cobalt and lithium frontier-rung breach recorded above is an open, reported issue rather than a result the model refused to produce
EU recycled-content floor plausibility	Whether the moderate-rung grid cells land in a plausible range against an announced regulatory floor	Checked after a run completes	Produces a flagged finding, not a block, and is explicitly not read as a compliance determination
Endpoint identity	That the two extremes of the primary-supply-elasticity continuum reproduce the legacy supply-capped and demand-responsive constructions exactly	A one-time regression verification performed when the elasticity continuum was introduced, not a check repeated on each run	Recorded as a passed verification at a point in time; it is not live and would not catch a future regression unless re-run

The pattern across Table 64 is informative. What blocks automatically is concentrated at the data-generation and single-run level — mass balance, demand ordering, and the composition and routing benchmarks — because those are checked with an unambiguous pass condition that can be evaluated the moment an artifact or run completes. What surfaces only as a flagged finding is concentrated at the cross-scenario and feasibility level — wedge conservation, no-double-count, the frontier caps, the regulatory-floor plausibility check — because those require modelling assumptions about what counts as an acceptable deviation, and a hard block would remove exactly the governed-decision latitude that the frontier-rung cap breach was deliberately left open under. Readers should draw the corresponding conclusion: a reported result is guaranteed to be internally consistent and mass-conserving, but a flagged feasibility or double-count finding is not automatically re-checked on each subsequent run and depends on the check having actually been re-run since the configuration last changed.

Glossary

Defines the terms used in this document that carry a specific technical meaning, grouped by the part of the model they belong to.

This glossary covers the vocabulary of this document. It is not a complete inventory of the model's internal terminology: terms that appear only in the codebase and not in this document are deliberately omitted, as defining concepts that readers will not encounter may hinder rather than aid understanding. The entries are grouped by the part of the model they belong to, in Tables 65 to 69.

Relationship to the report's own glossary. The CGR Critical Raw Materials digital report carries a separate, shorter glossary aimed at a general readership. Where a term appears in both, the definitions here are consistent with it and add technical precision rather than replacing it. Two mappings are worth stating explicitly, because the same words are used differently in the two documents and the difference matters exactly where supply falls short of demand. What the report calls *real demand* is the quantity this document also calls real demand — demand after scarcity and price response. What the report calls *realised demand* — material actually delivered — this document calls *delivered service*, because the model measures it at the point of entry into use rather than at the point of sale. A reader moving between the two documents should treat those two terms as the bridge.

Demand and service

These terms describe successive stages between what the economy wants and what it gets. They are easy to conflate, and the distinctions between them carry most of the model's analytical content, so they are given first.

Table 65. *Demand and service terms.*

Term	Definition
Desired demand	The demand trajectory of the model is given as input, before any price or scarcity response. Taken from the selected published transition pathway; sometimes called exogenous demand.
Real demand	Demand after the behavioural layer has applied its scarcity and price response. This is the target the supply system attempts to meet in a given year, and it is lower than desired demand whenever scarcity has suppressed consumption.
Suppressed demand	The gap between desired and real demand — consumption that did not occur because scarcity raised prices and buyers deferred or reduced purchases. It is a behavioural response, not a physical shortfall.
Service demand	The refined-material requirement entering fabrication, before fabrication and manufacturing yield losses. This is the boundary at which the model measures whether supply is adequate.
Delivered service	Material actually entering in-use stock, after fabrication and manufacturing losses and after any supply cap has bound. The report's glossary calls this quantity realised demand.
Unmet service	Called unmet demand in the report and in the indicator explorers. The shortfall between service demand and delivered service — requirements the system could

Term	Definition
	not physically satisfy from any source. Distinct from suppressed demand, which never became a requirement.
Service level	Delivered service as a fraction of service demand. A value of 1 means demand was fully met. The model's primary measure of supply adequacy.
Service deficit	The complement of service level: the unmet fraction of demand in a given year.
Functional service	The service-level family recomputed on a functional basis, crediting second-life units at their reduced service-delivery ratio rather than as full-capability units.

Material flow and mass balance

Table 66. *Material flow, stock and end-of-life terms.*

Term	Definition
Material stream	A labelled refined-equivalent flow within a material, differentiated by application grade — for example battery-grade against stainless-feed nickel. Streams activate at the refining stage and are carried on each flow and stock downstream of it. Twenty are defined across the five materials.
Refined-metal equivalent	The common mass basis on which all flows are tracked: the contained refined metal, so that quantities remain comparable across ore, intermediate and product forms.
Cohort	All material entering use in a single year, tracked as a group so that it retires according to its own lifetime distribution rather than at an average rate.
In-use stock	The mass of material currently in service, accumulated from past inflows and depleted by cohort retirement.
Burn-in	A synthetic spin-up period ahead of the study horizon in which the model accumulates a physically plausible starting in-use stock rather than beginning from zero. Length is set per material against a benchmark stock estimate.
Ore-to-metal ratio	The fraction of extracted ore that becomes refined material, combining ore grade with the cumulative efficiency of mining, beneficiation, smelting and refining. A ratio of 1% means 100 tonnes of ore per tonne of refined metal.
Primary available to refining	The refined-equivalent material reaches the refining stage from primary sources, after upstream allocation and yield losses.
Two-gate routing	The end-of-life routing structure: a first gate determines what fraction of retiring material is collected, and a second distributes collected material across recycling, remanufacturing and downcycling, with those three shares summing to one by construction.
Collection rate	The share of retiring material entering the collection system — the first routing gate. Supplied per stream and per year as an exogenous trajectory.
Functional recycling	Recovery in which material returns to the same material's usable supply pool. Material that is collected and processed but ends up in a lower-value circuit is downcycling, not functional recycling — a restrictive convention adopted throughout.

Term	Definition
Downcycling	Collected material processed into a lower-quality output that does not return to the originating material's supply pool.
Secondary supply generated vs used	Two quantities: the report and the explorer's report separately. Generated is all secondary material the system produces in a year — functionally recycled old scrap plus recovered new scrap. Used is the portion actually drawn into fabrication that year. They differ because the refinery stockpile buffers between them, so a year of high recovery need not be a year of high secondary use.
Old scrap	Material recovered from products that have reached the end of their service life.
New scrap	Material recovered from fabrication and manufacturing losses, never having entered use. Included in the model's secondary supply, which matters when comparing against end-of-life-only benchmarks.
Sorting rejects	Collected material lost during sorting and pre-processing, before any recovery route is reached.
System losses	The consolidated sink for material leaving the system without recovery — uncollected retirements, sorting rejects, and process losses.
Refinery stockpile	A buffer stock holding recovered secondary material between recovery and refining demand, governed by a release rate. A smoothing device rather than a market-clearing inventory.
Strategic inventory	A reserve stock per material that fills and releases according to policy intents generated by the behavioural layer. Tracked as an aggregate without quality or form detail.
Mass-balance residual	The maximum absolute discrepancy between inflows and outflows across the system in a run. Checked against a tolerance that halts the run if exceeded.

Behavioural layer and coupling

Table 67. *Behavioural response and coupling terms.*

Term	Definition
Slice	A single material and region pair. Most of the model's calculation and diagnostics happen at this level.
Loose iterative coupling	The scheme by which the behavioural and material-flow layers are solved alternately over the full horizon and iterated until their results stop changing materially — as distinct from solving both together year by year.
Inner loop	The repeated behavioural-then-material-flow solves within one slice.
Convergence metric	A single value summarising how much the exchanged signals changed between successive iterations; iteration stops when it falls below the convergence tolerance.
Capacity envelope	An aggregate index representing productive capacity per material, adjusting with a lag toward a target set by demand pressure. Not a representation of specific plants or capital stock.
Flow utilisation	Throughput as a fraction of the capacity envelope — how hard the system is running against its own ceiling.

Term	Definition
Bottleneck pressure	A measure of constraint severity derived from utilisation, which amplifies the scarcity signal reaching the price channel.
Scarcity multiplier	The amplified scarcity signal that drives price formation, combining the coupling-supplied stress signal with bottleneck amplification.
Price ratio	Modelled price relative to its baseline level; the quantity to which demand responds.
Collection effort	An accumulated state representing investment in collection capability, which builds under circular-supply pressure and decays otherwise.
Collection multiplier	The factor by which collection effort scales the exogenous collection trajectory, bounded by configured floors and ceilings.
Service stress signal	The feedback signal sent from the material-flow layer to the behavioural layer representing unmet service relative to demand.
Circular supply stress signal	The feedback signal represents dependence on primary rather than secondary sources. Interpreted as pressure to build collection capability, not as physical scarcity.
Stress multiplier	The combined stress term through which feedback signals enter the behavioural layer's price and capacity response.
Warm-start	Activating the behavioural layer ahead of the reporting window so that capacity, backlog and collection states are populated before results begin.
Tether strength	During warm-start, the weight blends observed historical demand against the model's own computed demand. One means fully anchored to history, zero means fully model-driven — the setting always used in the reporting window.
Dynamics start year	The per-material year at which behavioural dynamics and the background noise process activate. Kept separate from the gate that activates price response.

Scenarios and circularity levers

Table 68. *Scenario design and circularity terms.*

Term	Definition
Demand scenario	One of three trajectories aligned to published transition pathways, selecting a matched pair of total-demand and stream-composition series together.
Circularity rung	One of three ambition levels — current practice, moderate, frontier — at which the circularity levers are set.
NARROW	The lever reducing material required per unit of service, applied as a material-intensity multiplier on base demand before any recovery mechanism operates.
SLOW	The lever keeping products in service longer, through repair and maintenance lifetime extension, direct reuse, and refurbishment or remanufacturing.
CLOSE	The lever recovering more material at end of life, by improving the collection trajectory and recycling split and raising recycling process yield.

Term	Definition
Diversion rate	The share of retiring units routed into a product-recovery pathway rather than to material recovery.
Return rate	The share of diverted units that pass condition screening and actually re-enter service.
Second-life service ratio	The service a recovered unit delivers relative to a new one, reflecting degraded capability. A physical property of the recovered unit, not an ambition setting.
Wedge ladder	The construction in which each reported scenario cell is reached by adding levers cumulatively from its own baseline, so that the difference between consecutive rungs is the contribution of the lever added.
Wedge stages	The rungs of that ladder: the baseline, then successive additions of the narrowing, slowing and closing levers, terminating at the reported cell.
Shock event	A named disruption applying a multiplier to one model quantity over a defined window — supply cut, demand surge, capacity loss or collection collapse.
Pulse shock	The short-duration form of a shock event, used to separate the effect of a brief disruption from a sustained one.
Host cell	The unshocked scenario cell on which a shock is imposed, and against which the shocked run is compared.
Realised secondary share	The fraction of supply met from secondary sources in a given year — the quantity the frontier feasibility caps bound.
Ex-direct-scrap basis	The convention of excluding manufacturing scrap when measuring secondary share, so that only post-consumer material counts. The basis on which the feasibility caps are defined, and not the basis of the model's headline recycled-input indicator.

Indicators and resilience

Table 69. *Indicator and resilience terms.*

Term	Definition
Primary supply	Newly mined and refined material entering the system in a given year.
Secondary supply	Material recovered and returned to supply, comprising both end-of-life recycling and manufacturing scrap. The inclusion of manufacturing scrap is material to how this figure compares against external benchmarks.
End-of-life recycling rate	The share of material reaching the end of life that is functionally recycled — returned to the same material's usable supply pool.
Recycled input rate	Also written recycling input rate, the form the report uses; both denote the same indicator. The share of total supply met from secondary material. Includes manufacturing scrap, and so is not the quantity the frontier feasibility caps bound.
Old scrap ratio	The share of recycling input originating from end-of-life products rather than manufacturing scrap. A higher ratio means the recycling system is fed by used material rather than factory waste.

Term	Definition
Resilience triangle area	The cumulative service shortfall over a disruption and its recovery — the integrated gap between service level and its threshold.
Years below service threshold	The count of years in which service level falls below the configured adequacy threshold.
Censored recovery time	A recovery-time measurement in which recovery had not completed by the end of the horizon. The true value is unknown and bounded below; a censored result is not evidence of recovery.
Threshold service level	The configured service level below which a year counts as inadequate for resilience measurement. Resilience scalars are comparative across scenarios and move with this threshold.

Haarlemmerweg 331
1051 LH Amsterdam
info@circle-economy.com

