

Safe Aggregation in Modular Sustainability Assessment: A Method for Interface-Aware Pareto Screening

Juan J. Segura

`juan.segura.f@unab.cl`

Universidad Andres Bello

Method Article

Keywords: life cycle assessment, industrial ecology, process aggregation, foreground modeling, technology screening, Pareto optimization, prospective LCA, functional equivalence

Posted Date: September 25th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11001612/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Safe Aggregation in Modular Sustainability Assessment: A Method for Interface-Aware Pareto Screening

Juan J. Segura

Article type: Methodology

Affiliation: Universidad Andres Bello, Facultad de Ingenieria, Republica 220, Santiago, Chile

Abstract

Modular sustainability assessment requires deciding which properties of intermediate flows must remain explicit when process models are reused, screened, or connected. We propose a method for determining when such information can be aggregated without altering downstream environmental conclusions. Conditional life-cycle inventories are mapped to partially ordered sufficient-burden relations on typed process interfaces. The central result is a factorization criterion: a coarse interface preserves composition exactly when every intended downstream continuation is insensitive to the distinctions removed. For a declared upstream technology library and downstream continuation family, this yields a reachable observational quotient that is the coarsest exact interface preserving all relevant distinctions. The resulting interface defines where local Pareto pruning is valid before full system composition. In a constructed five-technology benchmark, nominal-product screening eliminates the unique system-level Pareto optimum, whereas interface-aware screening preserves it while still removing a genuinely redundant alternative and an irrelevant interface variable. Controlled Monte Carlo stress tests over 160,000 randomly generated modular systems show that the risk of false nominal pruning increases with downstream state sensitivity, continuation diversity, and technology-library size; exact interface-aware screening preserves the complete system frontier in all simulated cases. The method provides an auditable basis for process aggregation, reusable foreground-model design, and early technology screening in industrial ecology.

Keywords: life cycle assessment; industrial ecology; process aggregation; foreground modeling; technology screening; Pareto optimization; prospective LCA; functional equivalence

1. Introduction

Industrial-ecology models are modular by necessity. Product systems, production chains, recycling networks, industrial symbioses, and prospective technology scenarios are assembled from process descriptions that exchange materials, energy, services, and wastes. Life cycle assessment (LCA) formalizes many of these connections through unit processes, technosphere exchanges, elementary flows, reference flows, functional units, and system boundaries (ISO 14040, 2006; ISO 14044, 2006; Hellweg and Milà i Canals, 2014). Prospective LCA further emphasizes the need to model foreground technologies and their backgrounds consistently when technologies, scales, and future contexts change (Arvidsson et al., 2018; Mendoza Beltran et al., 2020; Joyce and Björklund, 2022; Douziech et al., 2024).

The usefulness of modularity depends on abstraction. A reusable process model cannot expose every internal variable and is therefore represented by selected exchanges, parameters, and functional properties. The unresolved methodological question is not whether abstraction should occur, but **which distinctions can be removed without invalidating later system-level conclusions.**

Consider two foreground technologies that both provide one kilogram of the same nominal intermediate material. If one requires less energy, water, and material than the other, local technology screening may

classify it as environmentally preferable and discard the second alternative. Such pruning is valid if the processes provide genuinely interchangeable intermediate states. It need not be valid when the states differ in purity, pressure, temperature, geographic location, timing, residual-stream composition, or another property that changes the requirements of downstream processes.

This issue is related to, but distinct from, the final functional unit. Comparative LCA already recognizes that nominally similar materials may differ in properties and performance, and that the appropriate functional unit depends on the comparison goal and available information (Furberg et al., 2022). The present paper asks the analogous question at **intermediate substitution boundaries** inside modular models:

Under what conditions may intermediate process states be aggregated without changing the sustainability-relevant conclusions obtained after the process is connected to the remainder of the system?

The question is especially relevant for reusable foreground libraries, prospective LCA, automated model assembly, technology roadmapping, and early-stage Pareto screening. A process representation that is sufficient for one downstream use may be inadequate for another. Conversely, a highly detailed representation may retain variables no intended continuation actually uses. The modeling problem is therefore one of **minimal sufficient interface resolution**.

Several established traditions provide parts of the answer. LCA supplies unit-process, system-boundary, foreground/background, inventory, and functional-equivalence discipline. Parameterized and prospective LCA frameworks show the value of transparent, reusable, and shareable foreground and background representations (Joyce and Björklund, 2022; Douziech et al., 2024), while dynamic industrial-ecology frameworks demonstrate that the appropriate model structure depends on temporal and system context (Sigüenza et al., 2021). Absolute environmental sustainability assessment (AESAs) adds a separate question: whether modeled burdens lie within allocated carrying-capacity or other admissibility thresholds (Bjørn et al., 2019; Ryberg et al., 2021; Guinée et al., 2022; Heide and Gjerris, 2024; Paulillo et al., 2025).

Multiobjective optimization contributes Pareto ordering, but it does not by itself determine **when local pruning is valid before system composition**. Monotone co-design provides the closest mathematical predecessor: functionality, resources, implementation choices, feasibility relations, and Pareto antichains are composed across engineering subsystems (Censi, 2016). Applied category theory provides a general language for systems whose components retain explicit interfaces until interconnected (Fong and Spivak, 2019; Baez et al., 2022). The present work does not claim novelty for the underlying ordered-relational, Pareto, or open-system machinery. These are used as formal infrastructure for a specific industrial-ecology problem: **when does model aggregation commute with future composition?** Full algebraic definitions and proofs are collected in the appendices so that the main text can focus on the industrial-ecology method.

The paper makes three contributions.

First, it connects conditional foreground inventories to a compositional sufficient-burden semantics. Each process transformation is represented by the set of burden allowances under which the declared boundary transformation can be realized, rather than by a single scalar score.

Second, it establishes a factorization criterion for safe interface aggregation and derives a reachable observational quotient: the coarsest exact representation that preserves all distinctions observable by the declared downstream continuation family among states the upstream library can actually produce.

Third, it converts that criterion into an interface-aware Pareto screening procedure. Genuine dominance within a continuation-sufficient interface class is preserved by later composition, whereas dominance inferred only after premature nominal-product aggregation has no such guarantee.

Contextual environmental thresholds are treated as a secondary assessment layer rather than part of the aggregation criterion. This keeps descriptive process semantics separate from choices about carrying-capacity allocation, policy limits, or other admissibility budgets.

The method is evaluated in two complementary ways. A constructed five-technology benchmark makes the failure mechanism transparent. A Monte Carlo stress test then examines how often premature nominal-product pruning fails in generic modular systems as downstream state sensitivity, continuation diversity, technology-library size, and burden dimensionality vary. The simulations are controlled methodological stress tests, not estimates of error rates in empirical LCA studies.

Throughout the paper, **safe** means invariant with respect to the declared burden semantics, reachable upstream library, and intended continuation family. It does not mean environmentally benign, risk-free, or universally valid.

2. Industrial-ecology problem formulation

2.1 Process modules, interfaces, and behaviours

Let a process module be written

$$P: X \rightarrow Y,$$

where X and Y are typed input and output interfaces. An interface specifies which exchanges and state descriptors are exposed for connection to other modules. A hydrogen-stream interface, for example, might expose mass, pressure, temperature, and purity; an electricity interface might expose quantity, location, voltage class, and time; a material stream might expose quantity, composition, moisture, particle size, or grade.

An interface must be distinguished from a particular boundary state. We write

$$\text{Beh}(Y)$$

for the admissible behaviours of interface Y . A specific state is an element


$$y \in \text{Beh}(Y).$$

Thus a port declaration might say that pressure is exposed, whereas a boundary behaviour specifies its value.

Parallel interfaces combine by product, so that

$$\text{Beh}(X \otimes Y) \cong \text{Beh}(X) \times \text{Beh}(Y).$$

The analysis requires only that process modules can be connected serially and placed in parallel while retaining their declared interfaces. Appendix A gives the corresponding formal open-system construction. Structured or decorated cospans provide one richer realization of this syntax (Baez et al., 2022), but the applied method does not depend on a particular cospan representation.



Aggregation through π is exact only when all intended downstream semantics are insensitive to the removed distinctions.

Figure 1. Modular process representation and interface aggregation. An upstream module produces a detailed intermediate state. A coarse representation may replace it only when every intended downstream continuation is insensitive to the distinctions removed by the aggregation.

2.2 Translation to industrial-ecology terminology

Table 1 summarizes the relation between the formal objects used below and standard industrial-ecology interpretations.

Table 1. Notation and industrial-ecology interpretation.

Symbol	Formal meaning	Industrial-ecology interpretation
$P: X \rightarrow Y$	Open process morphism	Foreground unit process or reusable process module
X, Y, Z	Typed interfaces	Exposed technosphere/intermediate exchanges and relevant properties
$\text{Beh}(Y)$	Interface behaviour space	Admissible states of exposed exchanges
y	Boundary behaviour	One specific intermediate-flow state
$Q \circ P$	Serial composition	Connected process modules
B	Ordered burden space	Selected characterized burdens/resources
$\mathcal{I}_P(x, y)$	Conditional inventory set	Inventories capable of providing $x \mapsto y$
$\mathcal{S}(P)(x, y)$	Sufficient-budget relation	Burden allowances sufficient to realize the boundary transformation
$\pi: Y \rightarrow \bar{Y}$	Coarse-graining	Removal/aggregation of state information
$\pi_!$	Pushforward	Process semantics after interface aggregation
$\mathcal{P}$	Upstream library	Admissible foreground technologies
$Y_{\mathcal{P}}$	Reachable state set	Intermediate states realizable by the upstream library
$\mathcal{C}$	Continuation family	Intended downstream model library
$\sim_{\mathcal{C}}$	Observational equivalence	Downstream-indistinguishable reachable states

Symbol	Formal meaning	Industrial-ecology interpretation
$Y_P / \sim_C$	Exact observational interface	Coarsest reachable interface preserving the intended continuations
β	Contextual budget	Declared environmental/admissibility threshold

2.3 Three distinct aggregation operations

The term aggregation is used for several operations that should be distinguished.

State-variable aggregation removes or coarsens a descriptor, for example

$$(p, T, \chi) \mapsto \chi.$$

Spatial or temporal aggregation replaces fine resolution by a coarser one, such as hourly electricity by annual electricity or plant location by region.

Process aggregation composes several unit processes and hides their common internal exchanges, producing a larger reusable module.

The factorization result in Section 6 applies directly to state, spatial, and temporal coarse-graining. Process aggregation is valid when the composite module preserves the external semantics required by its intended continuations.

3. From conditional inventories to burden semantics

3.1 Foreground inventories, background resolution, and characterization

A central accounting requirement is to avoid treating arbitrary cradle-to-gate impact vectors as independently additive modules. A foreground process may consume electricity, fuel, reagents, transport, or materials whose upstream burdens are resolved in a background system. If those upstream burdens are already embedded in a process result and the supplying modules are then composed explicitly, double counting can occur.

We therefore distinguish three semantic stages:

$$\mathbf{Proc} \xrightarrow{\mathcal{I}} \mathbf{LCI} \xrightarrow{\mathcal{R}} \mathbf{ElementaryFlows} \xrightarrow{\mathcal{K}} B.$$

Here $\mathcal{I}$ generates conditional foreground inventories, $\mathcal{R}$ resolves the selected background technosphere requirements into elementary flows, and $\mathcal{K}$ characterizes those flows into the burden dimensions used in the assessment. This separation is consistent with the foreground/background distinction emphasized in prospective and modular LCA practice (Mendoza Beltran et al., 2020; Joyce and Björklund, 2022).

For process $P: X \rightarrow Y$ and boundary behaviours x and y , let

$$\mathcal{I}_P(x, y)$$

be the set of conditional inventories capable of realizing $x \mapsto y$. Multiple inventories may correspond to alternative operating points, supply choices, equipment configurations, or implementations.

Let

$$K = \mathcal{K} \circ \mathcal{R}.$$

The key bridge from conventional inventory modeling to the semantics used in this paper is

$$\mathcal{S}(P)(x, y) = \uparrow \{K(i) : i \in \mathcal{I}_P(x, y)\}.$$

Thus $\mathcal{S}(P)(x, y)$ is the upward-closed set of burden allowances sufficient for at least one inventory realizing the specified boundary transformation. If several implementations are feasible, their nondominated characterized burdens form a Pareto representation of the same sufficient-budget set.

The exact burden-level composition developed below assumes that the selected characterization semantics is compatible with the declared process decomposition. If characterization depends jointly on spatial, temporal, concentration, or other states, those dependencies must either remain in the interface or composition must be performed at the inventory level before characterization. This is not an exception to the interface principle; it is an instance of it.

3.2 Burden spaces and sufficient-budget sets

Let

$$(B, \oplus, 0, \preceq)$$

be an ordered burden space in which $b \preceq b'$ means that b is no more burdensome than b' in every modeled dimension. For standard additive indicators one may take

$$B = \mathbf{R}_{\geq 0}^k$$

with componentwise Pareto order and vector addition. The more general notation allows other monotone composition rules where needed; for example, a peak requirement may compose by max rather than by addition.

The semantics $\mathcal{S}(P)(x, y)$ introduced above is an **upward-closed sufficient-budget set**: if allowance b is sufficient and $b \preceq b'$, then the larger allowance b' is also sufficient. For two such sets U and V , define their joint burden requirement by

$$U \star V = \uparrow \{u \oplus v : u \in U, v \in V\}.$$

This operation is the only ordered-algebraic construction needed in the main text. Appendix A supplies the formal algebraic foundation and establishes the associated composition laws.

4. Compositional burden semantics

Let

$$\Phi: A \rightsquigarrow C$$

associate each pair (a, c) with the sufficient-budget set for realizing the transformation $a \mapsto c$. If $\Psi: C \rightsquigarrow D$ is a downstream relation, their serial composition is

$$(\Psi \circ \Phi)(a, d) = \bigcup_{c \in C} \Phi(a, c) \star \Psi(c, d).$$

Applied to process modules, this equation says that the model considers every compatible intermediate state, combines the corresponding upstream and downstream burden requirements, and retains the alternatives that

realize the same external transformation. Internal process states disappear from the external system description, but their burden consequences do not.

The complete algebraic construction, including identities, associativity, parallel composition, and the extension from primitive process modules to arbitrary networks, is given in Appendix A. These constructions are standard; the contribution of this paper begins with the consequences of **which interface distinctions are retained before this composition is performed**.

Exact burden-level composition assumes **compositional closure**: every variable through which connected modules can modify one another's feasible burden semantics is either exposed in their common interface or represented explicitly by the burden-composition rule. Hidden shared infrastructure, nonlinear joint effects, or omitted coproduct relations can violate this assumption. In such cases the interface must be enriched, composition moved to the inventory level before characterization, or an approximate/lax semantics used.

5. Safe technology substitution

5.1 Interface-aware semantic dominance

Consider two alternatives

$$P, P': X \rightarrow Y.$$

We say that P' semantically dominates P at the declared interface, written

$$P \leq_S P',$$

when

$$\mathcal{S}(P)(x, y) \subseteq \mathcal{S}(P')(x, y)$$

for every admissible x and y . The inclusion direction is important: a superior process has a **larger** set of sufficient budgets because any allowance sufficient for the inferior process is also sufficient for the superior one.

Theorem 1 (dominance preservation). If $P \leq_S P'$, then every compatible upstream or downstream context preserves that dominance. In particular, for any downstream continuation Q ,

$$Q \circ P \leq_S Q \circ P'.$$

The result follows from monotonicity of burden composition and is proved in Appendix B. Intuitively, if P' can realize every identical boundary behaviour under every allowance that is sufficient for P , connecting the same continuation cannot make P superior.

Theorem 1 gives the safe pruning rule used later:

A locally dominated alternative may be removed before full system composition when the comparison is made at a complete interface sufficient for the declared continuation family.

The theorem also identifies the source of apparent ranking reversals. **True interface-aware dominance cannot reverse under monotone composition.** A reversal requires that the local comparison has first forgotten a distinction relevant to the surrounding system.

6. Safe interface aggregation

6.1 Coarse-graining

Let Y denote a detailed intermediate behaviour space and let

$$\pi: Y \rightarrow \bar{Y}$$

be a surjective coarse-graining. Examples include forgetting pressure, grouping purity ranges, replacing hourly states by annual averages, or aggregating locations into regions.

For an upstream relation $\Phi: X \rightsquigarrow Y$, define the pushed-forward relation

$$(\pi_! \Phi)(x, \bar{y}) = \bigcup_{\pi(y)=\bar{y}} \Phi(x, y).$$

This is exactly the information retained when distinct detailed states are represented by the same coarse state.

6.2 Main factorization theorem

Let

$$\Psi: Y \rightsquigarrow Z$$

be a downstream burden relation.

Theorem 2 (universal semantic criterion for safe aggregation). The following are equivalent:

1. There exists a downstream relation $\bar{\Psi}: \bar{Y} \rightsquigarrow Z$ such that

$$\Psi(y, z) = \bar{\Psi}(\pi(y), z)$$

for every y, z ; that is, the downstream model is insensitive to distinctions removed by π .

2. For every upstream budget relation $\Phi: X \rightsquigarrow Y$, aggregating before composition gives exactly the same external semantics:

$$\Psi \circ \Phi = \bar{\Psi} \circ \pi_! \Phi.$$

Proof sketch. If the downstream model is constant on each fiber of π , grouping upstream states within those fibers cannot change the union over compatible intermediate states. Conversely, semantic point probes show that exact equality for every upstream relation forces any two states identified by π to have identical downstream semantics. The complete proof, including the scope of the converse, is given in Appendix B.

The converse is exact at the abstract semantic-relation level because its proof uses arbitrary point probes. A physical process library need not contain a realizable module corresponding to every such probe. The forward direction, however, immediately yields the engineering criterion; the distinction is formalized in Appendix B.

Corollary 2.1. If Ψ factors through π , then aggregation through π is safe for every physically realizable upstream process whose semantics lies in the image of $\mathcal{S}$.

The industrial-ecology interpretation is direct:

An intermediate descriptor may be omitted exactly when all downstream models admitted by the declared application produce identical sustainability-relevant consequences for states that differ only in that descriptor.

Thus pressure can be removed if no intended continuation is pressure-sensitive; location can be aggregated only if transport, energy supply, regulation, characterization, and other included semantics are insensitive to the removed spatial distinction; and temporal resolution can be reduced only if later models are insensitive to the timing that is lost.

7. Reachable observational interfaces

7.1 Upstream reachability

A detailed interface may contain admissible states that no process in the actual technology library can generate. Let $\mathcal{P}$ be the upstream process library and define

$$Y_{\mathcal{P}} = \{y \in Y : \exists P \in \mathcal{P}, x, \mathcal{S}(P)(x, y) \neq \emptyset\}.$$

States outside $Y_{\mathcal{P}}$ cannot appear in systems assembled from the declared upstream library and therefore need not determine the interface resolution of that library.

7.2 Observational equivalence

Let

$$\mathcal{C} = \{\Psi_1, \dots, \Psi_m\}$$

be the intended downstream continuation family. For $y, y' \in Y_{\mathcal{P}}$, define

$$y \sim_{\mathcal{C}} y'$$

when

$$\Psi_j(y, z) = \Psi_j(y', z)$$


for every j and every compatible z .

Theorem 3 (coarsest exact reachable interface). The quotient

$$Y_{\mathcal{P}} / \sim_{\mathcal{C}}$$

is the coarsest exact representation of the reachable states through which every continuation in $\mathcal{C}$ factors.

The theorem follows by identifying exactly those reachable states that have identical downstream semantics across the declared continuation family. Its universal factorization proof is given in Appendix B.



Every intended continuation factors through the quotient when removed distinctions are downstream-unobservable.

Figure 2. Reachable observational quotient. The quotient retains only states realizable by the upstream library and only distinctions observable by the intended downstream continuation family.

The exact interface is therefore determined jointly by two questions:

What can the upstream library produce?

and

What can the downstream library distinguish?

This two-sided dependence is important for reusable model libraries. The observational interface is relative to the declared pair $(\mathcal{P}, \mathcal{C})$, not an immutable property of the physical process. If a new continuation is admitted later, a previously safe aggregation may need refinement. Accordingly, reusable model metadata should record the source interface, reachable-state domain, continuation scope, aggregation map, and whether the interface guarantee is exact or approximate.

7.3 Finite-library construction

For a finite continuation library, define the downstream observational signature

$$\sigma_{\mathcal{C}}(y) = [\Psi_1(y, \cdot), \dots, \Psi_m(y, \cdot)].$$

Reachable states with identical signatures belong to the same exact observational class. An implementable finite workflow is therefore:

1. enumerate or sample reachable intermediate states;
2. evaluate every declared downstream model on each state;
3. form downstream observational signatures;
4. group states having identical signatures;
5. expose the resulting quotient as the reusable exact interface;
6. perform local dominance pruning only within identical observational classes.

7.4 Approximate interfaces

Numerical models rarely produce exact equality. Let $D_C(y, y')$ measure the maximum sustainability-relevant difference between downstream responses to two states. A partition Π of Y_P is called ε -safe when every block $C \in \Pi$ satisfies

$$\max_{y, y' \in C} D_C(y, y') \leq \varepsilon.$$

Because tolerance-based indistinguishability need not be transitive, an ε -safe partition is not a canonical quotient. The present simulations use exact observational classes; Appendix C gives the approximate-partition construction and discusses complete-link and optimization-based alternatives.

8. Pareto screening and contextual thresholds

8.1 Pareto representations

The primary semantic object is the upward-closed set $\mathcal{S}(P)(x, y)$. When minimal elements exist, define

$$\text{PF}_P(x, y) = \text{Min}\mathcal{S}(P)(x, y).$$

For finite or finitely Pareto-generated systems,

$$\mathcal{S}(P)(x, y) = \uparrow \text{PF}_P(x, y),$$

and the composite frontier can be obtained by composing candidate frontiers across compatible intermediate states and then removing dominated burden vectors:

$$\text{PF}_{Q \circ P}(x, z) = \text{Min} \left[\bigcup_y \left(\text{PF}_P(x, y) \oplus \text{PF}_Q(y, z) \right) \right].$$

Minimal elements need not exist for arbitrary continuous upward sets, so Pareto antichains are treated as a computational representation where appropriate, not as the universal semantics. The finite-generation assumptions and derivation of the composite-frontier formula are collected in Appendix C.

The methodological distinction is between two screening orders.

Nominal early screening: aggregate nominal products $\rightarrow$ calculate local indicators $\rightarrow$ Pareto-prune $\rightarrow$ compose survivors.

Interface-aware screening: determine the continuation-sufficient interface $\rightarrow$ compare alternatives within observationally equivalent classes $\rightarrow$ prune safely $\rightarrow$ compose $\rightarrow$ Pareto-reduce complete systems.

8.2 Contextual environmental thresholds

The burden relation describes what a process can achieve. It does not determine which burden levels are ecologically, legally, or institutionally acceptable.

When components are assigned local burden contracts, a sufficient condition for a system-level admissibility claim is that the local contracts are jointly compatible with the system threshold under the same burden-composition rule. For additive indicators, this reduces to requiring that allocated local budgets sum to no more than the system budget. The formal statement and proof are given in Appendix B.4.

The point is methodological: process semantics and threshold allocation are separate layers. Changing an allocation rule can change admissibility without changing the underlying process representation or the safe-aggregation criterion.

This separation mirrors a central issue in AESA: environmental pressures and the allocation of admissible carrying-capacity shares are distinct methodological layers (Bjørn et al., 2019; Guinée et al., 2022; Heide and Gjerris, 2024). The present framework leaves the allocation rule external to the process semantics.

9. Constructed modular benchmark

9.1 System definition

We first use a deliberately transparent benchmark to show the mechanism of unsafe pruning. Five alternative foreground processes $P_1, \dots, P_5$ produce the same nominal intermediate material M . The detailed output state is

$$y = (\chi, \ell, T),$$

where χ is a purity class, ℓ a location class, and T a temperature class. Four detailed states are considered:

$$y_1 = (L, A, H), \quad y_2 = (H, A, W), \quad y_3 = (L, B, H), \quad y_4 = (L, A, W).$$

Processes P_1 and P_5 both produce y_1 . Burdens are represented by three synthetic pseudo-indicators b_1, b_2, b_3 , all minimized. They are dimensionless and do not represent empirical LCIA results.

Table 2. Local burdens in the constructed benchmark.

Process	Detailed state	b_1	b_2	b_3
P_1	y_1	5.0	4.0	5.0
P_2	y_2	6.0	5.0	6.0
P_3	y_3	4.5	4.5	5.5
P_4	y_4	5.2	3.8	5.1
P_5	y_1	5.6	4.6	5.8

If all states are collapsed to the nominal label M , P_1 dominates both P_2 and P_5 . Nominal early Pareto screening therefore discards P_2 and P_5 .

The downstream process depends on purity and location but, by construction, not temperature.

Table 3. Downstream burden increments in the constructed benchmark.

Detailed state	b_1	b_2	b_3
$y_1 = (L, A, H)$	4.0	4.0	5.0
$y_2 = (H, A, W)$	1.0	1.0	1.0
$y_3 = (L, B, H)$	6.0	2.0	4.0

Detailed state	b_1	b_2	b_3
$y_4 = (L, A, W)$	4.0	4.0	5.0

Because the downstream semantics of y_1 and y_4 is identical, temperature is observationally irrelevant in this model. Purity and location are not. The exact quotient therefore retains (χ, ℓ) but removes T .

Within the quotient class containing P_1, P_4, P_5 , P_1 safely dominates P_5 , so P_5 can be discarded by Theorem 1. P_2 and P_3 cannot be pruned against P_1 because they occupy different downstream-observational classes.

9.2 Full system result

After downstream composition, the system burdens are shown in Table 4.

Table 4. Composite burdens in the constructed benchmark.

System	Upstream process	b_1	b_2	b_3
S_1	P_1	9.0	8.0	10.0
S_2	P_2	7.0	6.0	7.0
S_3	P_3	10.5	6.5	9.5
S_4	P_4	9.2	7.8	10.1

System S_2 , generated by the technology discarded under nominal local screening, strictly dominates every other complete system. Nominal screening therefore eliminates the unique global Pareto optimum. Interface-aware pruning removes one irrelevant descriptor and one genuinely dominated technology while retaining the correct system frontier.




Figure 3. Composite burdens in the constructed benchmark. Lower values are preferred in all three synthetic dimensions. System S_2 , generated by the technology discarded by nominal local screening, dominates the remaining complete systems.

The benchmark is a counterexample and is intentionally constructed. It establishes possibility, not prevalence. Section 10 therefore evaluates the screening procedures over randomized systems in which no reversal is imposed.

10. Generic computational stress test

10.1 Purpose and random-system generator

The Monte Carlo experiment is not intended to estimate the frequency of false technology screening in empirical LCA studies. The random generator has no empirical distributional interpretation. Its purpose is to evaluate the competing screening rules over a controlled ensemble while varying the strength and diversity of downstream state dependence.

Each reference system contains $n = 10$ upstream technology alternatives and $k = 3$ burden dimensions. Local burdens are independently sampled as

$$b_i^{\text{loc}} \sim U([0,1]^3).$$

Each technology produces one intermediate state. In the reference experiment six detailed states are organized into three downstream-observational classes:

$$Y = \{(c, v) : c \in \{1, 2, 3\}, v \in \{0, 1\}\}.$$

The class c affects downstream burden; nuisance descriptor v does not. Each technology is assigned uniformly to one of the detailed states. For observational class c , draw

$$u_c \sim U([0,1]^k)$$

and define the downstream burden increment

$$d_c = \lambda u_c,$$

where $\lambda \geq 0$ controls downstream interface sensitivity. When $\lambda = 0$, all intermediate states have identical downstream consequences and nominal aggregation is exact. Composite burdens are

$$b_i^{\text{sys}} = b_i^{\text{loc}} + d_{c_i}.$$

No ordering reversal is imposed.

Three sets are compared. The **nominal screening set** is the local Pareto frontier after all state distinctions are forgotten. The **interface-aware screening set** is the union of local Pareto frontiers computed separately within exact downstream-observational classes. The **reference frontier** is the Pareto frontier of all complete system burdens.

A false-pruning event occurs when an early-screening rule removes at least one technology belonging to the true complete-system Pareto frontier. We also record early reduction rate and true-frontier coverage.

All experiments were generated with a master random seed of 20260909 using independent deterministic child seeds for each parameter cell. Code, cell-specific seeds, aggregate outputs, and software versions are supplied in the Supporting Information.

10.2 Downstream sensitivity

For each

$$\lambda \in \{0, 0.1, 0.25, 0.5, 1, 2\},$$

10,000 independent systems were generated, for a total of 60,000 realizations.

Table 5. Effect of downstream interface sensitivity on early Pareto screening.

λ	Nominal false pruning	95% Wilson interval	Nominal reduction	Interface-aware reduction	Nominal coverage of true frontier
0.00	0.0000	0.0000-0.0004	49.59%	49.59%	100.00%
0.10	0.1909	0.1833-0.1987	49.29%	26.44%	95.92%
0.25	0.3678	0.3584-0.3773	49.60%	26.59%	91.14%
0.50	0.5245	0.5147-0.5343	49.42%	26.61%	85.69%
1.00	0.6525	0.6431-0.6618	49.28%	26.30%	79.17%
2.00	0.7387	0.7300-0.7472	49.50%	26.57%	74.15%

At $\lambda = 0$, downstream semantics is independent of state. The observational interface therefore collapses to the nominal product and both screening procedures are identical and exact. Once downstream consequences differ, nominal-product pruning begins to remove system-relevant alternatives. Even at $\lambda = 0.1$, at least one true system-frontier technology is lost in approximately 19.1% of systems. The probability exceeds one half at $\lambda = 0.5$ and reaches 73.9% at $\lambda = 2$.

Interface-aware pruning produced zero false eliminations in the exact finite-class experiment. This is not an empirical reliability estimate: it is the computational consequence of Theorem 1 because pruning occurs only within classes having identical downstream semantics. The meaningful empirical quantity in the simulation is therefore the amount of reduction that remains available under that guarantee: approximately 26% of technologies in the three-class setting, compared with approximately 49% under nominal screening.




Figure 4. Probability that nominal early Pareto screening removes at least one technology on the true system Pareto frontier as downstream state sensitivity λ increases. Error bars show 95% Wilson intervals. Interface-aware pruning is exact in the finite-class model because pruning is restricted to states with identical downstream semantics.

10.3 Continuation diversity

A second experiment varied the number of potential downstream-observational classes in $\{1,2,3,5\}$ for $\lambda = 0.5$ and $\lambda = 1$, with 5,000 realizations per cell (40,000 systems). Because only ten upstream technologies are sampled, not every potential class is necessarily reachable in each realization. The mean numbers of reachable classes were approximately 2.95 when three classes were available and 4.46 when five were available.

At $\lambda = 0.5$, nominal false-pruning probability increased from 0 with one class to 41.5% with two classes, 53.6% with three, and 59.3% with five. At $\lambda = 1$, the corresponding values were 0, 55.4%, 64.7%, and 71.5%. Meanwhile, safe early reduction declined as more distinctions became downstream-observable. Complete results are reported in Supporting Information Table S1.

This is the intended behavior. A model library with greater continuation diversity legitimately permits less aggressive local aggregation and pruning.

10.4 Technology-library size and burden dimensionality

A third experiment varied

$$n \in \{5, 10, 20, 50\}, \quad k \in \{2, 3, 5\},$$

at $\lambda = 0.5$ with three potential downstream classes. Five thousand realizations were evaluated for each of the 12 cells, adding 60,000 systems.

The failure of nominal-product screening is not peculiar to the reference $n = 10, k = 3$ setting. With five candidate technologies, the probability of deleting at least one true system-Pareto solution ranges from approximately 16% to 28%, depending on burden dimensionality. With 50 alternatives it ranges from approximately 81% to 97%. Interface-aware pruning again produced no false eliminations under the exact class semantics. Complete numerical results are given in Supporting Information Table S2.




Figure 5. Sensitivity of false nominal pruning to technology-library size n and burden dimensionality k at $\lambda = 0.5$ with three potential downstream-observational classes. Each cell is based on 5,000 independently generated systems.

Increasing k makes random local vectors more likely to be mutually incomparable, so local dominance and early reduction can decrease. Increasing n supplies more opportunities for both local dominance and downstream-dependent ranking changes. The combined result is a strong library-size effect: in large alternative sets, aggressive nominal pruning can be highly unreliable even when local Pareto screening appears attractive computationally.

Across the three experiments, aggregation risk is controlled by at least three factors:

- the magnitude of downstream state sensitivity;
- the number of downstream-observable distinctions; and

- the size and dimensionality of the upstream technology library.

The amount of safe simplification is therefore an endogenous property of the model ecosystem, not a fixed modeling preference.

11. Physical illustration: hydrogen delivery pressure

A simple physical example clarifies the factorization criterion without claiming superiority of any commercial electrolyzer technology.

Suppose two production routes provide one unit of hydrogen at different pressures,

$$y_L = (1, p_L), \quad y_H = (1, p_H),$$

and the final delivery requirement is $p_D = p_H > p_L$. Let the low-pressure production stage require energy E_0 , while achieving the high-pressure interface during production requires an incremental amount δ_H . If pressure is forgotten, the low-pressure production stage appears preferable because

$$E(P_L) = E_0 < E_0 + \delta_H = E(P_H).$$

Now let c_L denote the downstream energy required to transform the low-pressure stream to p_D . The complete pathways require

$$E_L^{\text{sys}} = E_0 + c_L, \quad E_H^{\text{sys}} = E_0 + \delta_H.$$

The high-pressure pathway has the lower complete-system energy whenever

$$\delta_H < c_L.$$

The purpose of the example is not to assert that any commercial technology lies on either side of this inequality. It is to show that a downstream compression model distinguishes y_L from y_H . Therefore a projection that retains hydrogen quantity but removes pressure does not satisfy the factorization criterion for that continuation.

If the declared continuation family were instead entirely insensitive to pressure, pressure could legitimately be removed. The required interface is therefore continuation-relative rather than an intrinsic label attached to the production technology.

12. Discussion

12.1 Relation to LCA and functional-equivalence reasoning

The framework is intended to complement, not replace, LCA. ISO-style functional units and system boundaries determine the common final function and the scope of an assessment. The present contribution addresses a different but closely related problem: the validity of **intermediate** model substitutions before the complete product system has been assembled.

Furberg et al. (2022) show that material comparisons can employ reference-flow, property, or performance functional units depending on purpose and data maturity, with different degrees of functional equivalence. The observational equivalence introduced here should not be interpreted as a redefinition of the LCA functional unit. It is more specific: it is an intermediate-interface analogue of functional-equivalence

reasoning. Two internal states are equivalent relative to continuation family $\mathcal{C}$ when no continuation in that family produces different modeled burden semantics from them.

This distinction clarifies why nominal product identity can be insufficient. “Both processes produce one kilogram of material” establishes reference-flow similarity. It does not establish downstream interchangeability if properties relevant to later operations differ.

12.2 Foreground/background modularity and prospective LCA

The foreground/background separation in Section 3 is essential for avoiding an incorrect interpretation of burden composition. Joyce and Björklund (2022) emphasize that prospective foreground models sit on a larger background industrial system, while Mendoza Beltran et al. (2020) and Douziech et al. (2024) show how changing backgrounds and parameterized foregrounds affect prospective assessments. Our relation $\mathcal{S}(P)(x, y)$ is therefore not an instruction to add independently calculated cradle-to-gate scores. It is a semantics derived from conditional inventories under a declared foreground/background accounting convention.

The factorization principle also applies to prospective spatial and temporal resolution. Annual electricity is an exact substitute for time-resolved electricity only when all downstream inventory, characterization, and decision semantics used by the study are insensitive to within-year timing. Regional location is exact only when finer spatial distinctions are irrelevant to every included continuation. This provides a common formal interpretation of state, temporal, and spatial aggregation.

12.3 Relation to compositional engineering design

The formal infrastructure is closely related to compositional engineering design. Censi’s monotone co-design provides functionality-resource relations and Pareto antichains that compose across engineering subsystems (Censi, 2016), while Fong and Spivak (2019) provide a broader open-system language for preserving interfaces until components are interconnected. The present paper uses these ideas as mathematical infrastructure rather than as the primary contribution.

The sustainability-specific contribution lies instead in combining conditional burden semantics with an explicit **information-preservation problem**: when may an intermediate environmental model be coarsened before future composition? The factorization theorem answers this exactly at the semantic level, and the reachable observational quotient turns the result into a model-design rule for a specified upstream/downstream library.

The distinction is also useful computationally. Pareto optimization normally begins once the decision space and system model have been constructed. Here the question is prior: which local Pareto eliminations are sound **before** all continuations have been composed? Theorem 1 shows that dominance on a sufficient interface is stable; the benchmark and simulations show why dominance created only after unsafe aggregation is not.

12.4 Contextual budgets and absolute sustainability

AESA compares modeled environmental pressures with allocated carrying-capacity or other threshold shares (Bjørn et al., 2019; Ryberg et al., 2021). The literature also makes clear that allocation principles and temporal choices remain consequential and value-laden (Guinée et al., 2022; Heide and Gjerris, 2024), with recent technical guidance continuing to develop the field (Paulillo et al., 2025).

The present semantics deliberately separates

$$\mathcal{S}(P)$$

from the admissibility threshold β . Changing an allocation rule can therefore change whether a modeled system satisfies its contextual contract without changing the process burden semantics itself. This separation prevents normative allocation from being hidden inside the physical process model or a single sustainability score.

12.5 Practical implications for reusable model libraries

The reachable quotient suggests a concrete metadata principle for modular industrial-ecology software. A reusable component should record not only what it produces, but also the continuation scope against which its interface has been simplified. At minimum, a certified interface description should identify:

- source/detail interface;
- reachable-state domain;
- intended continuation family or continuation specification;
- exact aggregation map or approximate partition;
- approximation tolerance and distance metric, if applicable;
- burden semantics against which equivalence was tested.

This makes model resolution auditable. A newly introduced continuation can trigger refinement rather than silently invalidating earlier aggregation.

The method is not an argument for maximal detail. At $\lambda = 0$ in the Monte Carlo experiment, all downstream states are identical and the method collapses exactly to nominal screening. It retains additional information only when the intended continuation family makes that information consequential.

12.6 Limitations

The framework is static and deterministic. Dynamic stocks, feedback, recycling loops, temporal environmental capacities, path dependence, and ecological recovery require behavioural or dynamical extensions. A future theory could replace static boundary states with trajectories and combine burden feasibility with viability constraints.

Second, the continuation family must be declared. The coarsest exact interface is relative to $(\mathcal{P}, \mathcal{C})$. It is not guaranteed to remain sufficient after the modeling universe expands.

Third, exact burden-level composition assumes compositional closure. Shared utilities, scale economies, coproduct allocation, or nonlinear characterization may create couplings not captured by the current interface. Such couplings must be exposed or resolved at an earlier semantic layer.

Fourth, the ε -safe construction requires a distance metric and tolerance. These are modeling choices, and complete-link clustering is only one conservative algorithm; it is not a canonical optimal quotient.

Fifth, the burden space is best interpreted here as environmental/resource semantics with optional hard constraints. Social sustainability involves procedural, institutional, distributive, and stakeholder dimensions that should not be reduced uncritically to additional scalar burden coordinates.

Finally, the Monte Carlo generator is intentionally generic. It demonstrates structural behavior under controlled variation but does not estimate the prevalence or magnitude of aggregation errors in published LCAs. Empirical testing on public modular LCA libraries would be a natural next step.

13. Conclusions

Modular sustainability models inevitably discard detail. The methodological problem is to determine whether the discarded information can alter conclusions after a module is connected to something else.

The central result is exact. A coarse interface

$$\pi: Y \rightarrow \bar{Y}$$

preserves composition with a downstream semantic relation precisely when that continuation factors through π . For a declared upstream library $\mathcal{P}$ and continuation family $\mathcal{C}$, the reachable observational interface

$$Y_{\mathcal{P}} / \sim_{\mathcal{C}}$$

is the coarsest exact representation that retains every downstream-observable distinction among states the upstream library can actually produce.

This result yields a safe technology-screening rule: local dominance can be used for pre-composition pruning only after alternatives are compared at a continuation-sufficient interface. In the constructed benchmark, nominal-product screening discards the unique system-level Pareto optimum. In the randomized stress tests, the risk of such false pruning rises sharply with downstream state sensitivity, continuation diversity, and technology-library size. Interface-aware pruning remains exact in the finite-class experiments while still achieving substantial early reduction whenever genuinely redundant alternatives exist.

The practical principle is therefore simple:

A reusable process model should expose no more detail than its intended downstream uses can observe, but no distinction should be discarded before its downstream irrelevance has been established.

This converts safe aggregation from an informal judgment about “enough detail” into an auditable property of a declared model ecosystem. It provides a basis for continuation-aware foreground libraries, safe early Pareto screening, and compositional sustainability assessment.

Acknowledgments

No external acknowledgments.

Funding

No external funding was received for this work.

Conflict of interest

The author declares no conflict of interest.

Data availability statement

No empirical dataset was analyzed in this study. The constructed benchmark data and all Monte Carlo datasets are synthetic. The complete simulation code, cell-specific random seeds, aggregate CSV outputs, benchmark data, and software-environment information are provided in the Supporting Information package so that every reported computational result can be reproduced.

Declaration of generative AI and AI-assisted technologies in the writing process

During preparation of this work, the author used generative AI tools for language editing and translation from Spanish. The author reviewed the final manuscript and takes full responsibility for the content.

Appendix A. Algebraic foundations

The main text uses only the composition operation required by the applied method. This appendix records the standard ordered-relational construction that guarantees associativity and compatibility with serial and parallel process composition.

A.1 Ordered burden monoids and upward sets

Let $(B, \oplus, 0, \leq)$ be a commutative partially ordered monoid with monotone multiplication:

$$b_1 \leq b_2 \Rightarrow b_1 \oplus c \leq b_2 \oplus c.$$

A subset $U \subseteq B$ is upward closed when $b \in U$ and $b \leq b'$ imply $b' \in U$. Let $\text{Up}(B)$ denote all upward-closed subsets and define

$$U \star V = \uparrow \{u \oplus v : u \in U, v \in V\}.$$

Proposition A.1. With inclusion as order, arbitrary union as join, $\star$ as multiplication, and $\uparrow 0$ as unit,

$$\mathcal{Q}_B = (\text{Up}(B), \subseteq, \cup, \star, \uparrow 0)$$

is a commutative unital quantale.

Proof. Associativity and commutativity follow from those of $\oplus$. Monotonicity of $\oplus$ ensures that upward closure is compatible with multiplication. The unit property follows because $u \oplus 0 = u$, and $U \star \bigcup_i V_i = \bigcup_i (U \star V_i)$ follows by distributing set formation and upward closure over arbitrary unions. $\square$

This is a standard construction (Rosenthal, 1990). Its role here is to represent alternatives by joins and simultaneous burden requirements by multiplication.

A.2 Budget relations

A budget relation $\Phi: A \rightsquigarrow C$ is a function

$$\Phi: A \times C \rightarrow \mathcal{Q}_B.$$

For $\Phi: A \rightsquigarrow C$ and $\Psi: C \rightsquigarrow D$, define

$$(\Psi \circ \Phi)(a, d) = \bigcup_{c \in C} \Phi(a, c) \star \Psi(c, d).$$

The identity relation is

$$I_A(a, a') = \begin{cases} \uparrow 0, & a = a', \\ \emptyset, & a \neq a'. \end{cases}$$

Proposition A.2. Sets and budget relations form a category $\mathbf{BRel}_B$. With Cartesian product on state spaces and $\star$ on burden values, $\mathbf{BRel}_B$ is symmetric monoidal.

Proof. Associativity is obtained by expanding either parenthesization of triple composition and using associativity and distributivity of $\star$ over unions. The stated identity is neutral by the unit law. For parallel composition define

$$(\Phi \boxtimes \Psi)((a, a'), (c, c')) = \Phi(a, c) \star \Psi(a', c').$$

The monoidal axioms and interchange law follow from commutativity, associativity, and distributivity in $\mathcal{Q}_B$.
 $\square$

A.3 Process syntax and semantic extension

Let $\mathbf{Proc}_\Sigma$ be the free symmetric monoidal category generated by a typed process signature Σ , and let Beh assign to each interface its admissible behaviour set, with

$$\text{Beh}(X \otimes Y) \cong \text{Beh}(X) \times \text{Beh}(Y).$$

For each primitive process generator $g: X \rightarrow Y$, assign

$$\hat{\mathcal{S}}(g): \text{Beh}(X) \rightsquigarrow \text{Beh}(Y).$$

Proposition A.3. These object and generator assignments extend uniquely to a symmetric monoidal functor

$$\mathcal{S}: \mathbf{Proc}_\Sigma \rightarrow \mathbf{BRel}_B.$$

Proof. This is the universal property of the free symmetric monoidal category, since $\mathbf{BRel}_B$ is symmetric monoidal by Proposition A.2. $\square$

The construction assumes compositional closure: every variable through which connected modules can modify one another's feasible burden semantics is either exposed at the interface or represented explicitly in the burden-composition operation. If this assumption fails, exact strong monoidality is not justified; the model must expose the missing coupling, compose at an earlier semantic level, or use a lax/approximate semantics.

Appendix B. Proofs of the interface and admissibility results

B.1 Proof of Theorem 1: dominance preservation

Suppose $P \preceq_\mathcal{S} P'$, so that

$$\mathcal{S}(P)(x, y) \subseteq \mathcal{S}(P')(x, y)$$

for every admissible x, y . For any downstream continuation Q and external states x, z ,

$$\mathcal{S}(Q \circ P)(x, z) = \bigcup_y \mathcal{S}(P)(x, y) \star \mathcal{S}(Q)(y, z).$$

Because $\star$ is monotone in each argument,

$$\mathcal{S}(P)(x, y) \star \mathcal{S}(Q)(y, z) \subseteq \mathcal{S}(P')(x, y) \star \mathcal{S}(Q)(y, z)$$

for every y . Taking the union preserves inclusion, giving

$$\mathcal{S}(Q \circ P)(x, z) \subseteq \mathcal{S}(Q \circ P')(x, z).$$

The upstream case is analogous. $\square$

B.2 Proof of Theorem 2: factorization criterion for safe aggregation

Assume first that there is $\bar{\Psi}: \bar{Y} \rightsquigarrow Z$ satisfying

$$\Psi(y, z) = \bar{\Psi}(\pi(y), z).$$

For any upstream relation $\Phi: X \rightsquigarrow Y$,

$$(\bar{\Psi} \circ \pi_! \Phi)(x, z) = \bigcup_{\bar{y}} (\pi_! \Phi)(x, \bar{y}) \star \bar{\Psi}(\bar{y}, z).$$

Using

$$(\pi_! \Phi)(x, \bar{y}) = \bigcup_{\pi(y)=\bar{y}} \Phi(x, y),$$

and distributivity over unions gives

$$(\bar{\Psi} \circ \pi_! \Phi)(x, z) = \bigcup_{\bar{y}} \bigcup_{\pi(y)=\bar{y}} \Phi(x, y) \star \bar{\Psi}(\pi(y), z).$$

Since $\Psi(y, z) = \bar{\Psi}(\pi(y), z)$, the right-hand side is

$$\bigcup_y \Phi(x, y) \star \Psi(y, z) = (\Psi \circ \Phi)(x, z).$$

Conversely, suppose the composite equality holds for every upstream budget relation. If $\pi(y) = \pi(y')$, take point-probe upstream relations from a singleton state space that select y and y' respectively at the quantale unit and no other output state. Their pushforwards through π coincide. Equality after composition therefore forces

$$\Psi(y, z) = \Psi(y', z)$$

for every z . Hence Ψ is constant on every fiber of π . Surjectivity of π then defines a unique $\bar{\Psi}$ on $\bar{Y}$ with

$$\Psi = \bar{\Psi} \circ \pi.$$

The necessity direction is exact at the semantic-relation level because arbitrary point probes need not correspond to physically realizable process modules. ▀

B.3 Proof of Theorem 3: coarsest exact reachable interface

By definition,

$$y \sim_{\mathcal{C}} y'$$

exactly when every declared continuation assigns identical downstream semantics to y and y' . Therefore each $\Psi_j \in \mathcal{C}$ is constant on the equivalence classes and factors uniquely through

$$q_{\mathcal{C}}: Y_{\mathcal{P}} \rightarrow Y_{\mathcal{P}} / \sim_{\mathcal{C}}.$$

Now let $\pi: Y_{\mathcal{P}} \rightarrow \bar{Y}$ be another surjective coarse-graining that is safe for every continuation in $\mathcal{C}$. If $\pi(y) = \pi(y')$, Theorem 2 implies that $\Psi_j(y, z) = \Psi_j(y', z)$ for every j, z . Thus $y \sim_{\mathcal{C}} y'$, so every fiber of π lies within one observational equivalence class. The quotient map is therefore constant on every fiber of π and factors uniquely as

$$q_{\mathcal{C}} = r \circ \pi$$

for some $r: \bar{Y} \rightarrow Y_{\mathcal{P}} / \sim_{\mathcal{C}}$. Hence the observational quotient is the coarsest exact safe representation. ▀

B.4 Compatible local admissibility contracts

For completeness, suppose two connected processes have compatible intermediate behaviour and local admissibility contracts. If each process satisfies its local contract and the two local burden allowances compose to no more than the declared system allowance, then the composite satisfies the system contract.

$$\beta_P \in \mathcal{S}(P)(x, y), \quad \beta_Q \in \mathcal{S}(Q)(y, z).$$

By the composition rule,

$$\beta_P \oplus \beta_Q \in \mathcal{S}(Q \circ P)(x, z).$$

If

$$\beta_P \oplus \beta_Q \leq \beta_{\text{sys}},$$

then upward closure gives

$$\beta_{\text{sys}} \in \mathcal{S}(Q \circ P)(x, z),$$

which is precisely

$$Q \circ P \models_{(x, z, \beta_{\text{sys}})} \text{Adm.}$$

□

Appendix C. Pareto and approximate-interface details

C.1 Pareto representation of sufficient-budget sets

The upward set $\mathcal{S}(P)(x, y)$ is the primary semantic object. When it is finitely Pareto-generated,

$$\mathcal{S}(P)(x, y) = \uparrow \text{PF}_P(x, y),$$

where

$$\text{PF}_P(x, y) = \text{Min} \mathcal{S}(P)(x, y).$$

For two composable processes, substituting the finite generators into the relational composition equation and then taking minimal elements gives

$$\text{PF}_{Q \circ P}(x, z) = \text{Min} \left[\bigcup_y \left(\text{PF}_P(x, y) \oplus \text{PF}_Q(y, z) \right) \right].$$

Thus the computational rule is to combine candidate frontier points across compatible intermediate states and remove dominated results. For arbitrary continuous upward sets, minimal elements need not exist, so closures or ε -frontiers may be required.

C.2 Approximate safe partitions

Let $D_{\mathcal{C}}(y, y')$ measure the largest declared sustainability-relevant difference between the downstream responses to two reachable states. A partition Π is ε -safe if every block $C \in \Pi$ satisfies

$$\max_{y, y' \in C} D_C(y, y') \leq \varepsilon.$$

The binary relation $D_C(y, y') \leq \varepsilon$ need not be transitive; consequently an ε -safe partition is not a canonical quotient. Complete-link agglomeration is one conservative construction because it permits a merge only when the resulting block diameter stays within ε . It guarantees the tolerance condition but need not give the smallest possible number of blocks.

A maximally aggregated admissible interface can instead be formulated as

$$\min_{\Pi} |\Pi|$$

subject to

$$\max_{y, y' \in C} D_C(y, y') \leq \varepsilon \quad \text{for every } C \in \Pi.$$

The exact finite-class experiments in this paper use $\varepsilon = 0$ and therefore do not depend on a clustering heuristic.

C.3 Nonadditive burden coordinates

The ordered-monoid representation is not restricted to vector addition. If a coordinate represents a peak load, its local values may compose by max. A noncompensable binary violation may use logical OR. The general admissibility condition

$$\oplus_i \beta_i \preceq \beta_{\text{sys}}$$

should therefore be read as compatibility of local contracts. The more specific term **budget allocation** is appropriate only for additive resource or carrying-capacity coordinates.

References

- Arvidsson, R., Tillman, A.-M., Sandén, B. A., Janssen, M., Nordelöf, A., Kushnir, D., & Molander, S. (2018). Environmental assessment of emerging technologies: Recommendations for prospective LCA. *Journal of Industrial Ecology*, 22, 1286-1294. <https://doi.org/10.1111/jiec.12690>
- Baez, J. C., Courser, K., & Vasilakopoulou, C. (2022). Structured versus decorated cospan. *Compositionality*, 4, 3. <https://doi.org/10.32408/compositionality-4-3>
- Bjørn, A., Richardson, K., & Hauschild, M. Z. (2019). A framework for development and communication of absolute environmental sustainability assessment methods. *Journal of Industrial Ecology*, 23, 838-854. <https://doi.org/10.1111/jiec.12820>
- Censi, A. (2016). Monotone co-design problems; or, everything is the same. In *2016 American Control Conference (ACC)* (pp. 1227-1234). IEEE. <https://doi.org/10.1109/ACC.2016.7525085>
- Douziech, M., Besseau, R., Jolivet, R., Shoai-Tehrani, B., Bourmaud, J.-Y., Busato, G., Gresset-Bourgeois, M., & Pérez-López, P. (2024). Life cycle assessment of prospective trajectories: A parametric approach for tailor-made inventories and its computational implementation. *Journal of Industrial Ecology*, 28, 25-40. <https://doi.org/10.1111/jiec.13432>
- Fong, B., & Spivak, D. I. (2019). *An Invitation to Applied Category Theory: Seven Sketches in Compositionality*. Cambridge University Press. <https://doi.org/10.1017/9781108668804>
- Furberg, A., Arvidsson, R., & Molander, S. (2022). A practice-based framework for defining functional units in comparative life cycle assessments of materials. *Journal of Industrial Ecology*, 26, 718-730. <https://doi.org/10.1111/jiec.13218>
- Guinée, J. B., de Koning, A., & Heijungs, R. (2022). Life cycle assessment-based Absolute Environmental Sustainability Assessment is also relative. *Journal of Industrial Ecology*, 26, 673-682. <https://doi.org/10.1111/jiec.13260>
- Heide, M., & Gjerris, M. (2024). Embedded but overlooked values: Ethical aspects of absolute environmental sustainability assessments. *Journal of Industrial Ecology*, 28, 386-396. <https://doi.org/10.1111/jiec.13472>
- Hellweg, S., & Milà i Canals, L. (2014). Emerging approaches, challenges and opportunities in life cycle assessment. *Science*, 344, 1109-1113. <https://doi.org/10.1126/science.1248361>
- International Organization for Standardization. (2006a). *ISO 14040:2006 Environmental management - Life cycle assessment - Principles and framework*. Geneva: ISO.
- International Organization for Standardization. (2006b). *ISO 14044:2006 Environmental management - Life cycle assessment - Requirements and guidelines*. Geneva: ISO.
- Joyce, P. J., & Björklund, A. (2022). Futura: A new tool for transparent and shareable scenario analysis in prospective life cycle assessment. *Journal of Industrial Ecology*, 26, 134-144. <https://doi.org/10.1111/jiec.13115>
- Mendoza Beltran, A., Cox, B., Mutel, C., van Vuuren, D. P., Font Vivanco, D., Deetman, S., Edelenbosch, O. Y., Guinée, J., & Tukker, A. (2020). When the background matters: Using scenarios from integrated

- assessment models in prospective life cycle assessment. *Journal of Industrial Ecology*, 24, 64-79. <https://doi.org/10.1111/jiec.12825>
- Paulillo, A., Sanyé-Mengual, E., De Laurentiis, V., Veà, E., Hauschild, M. Z., Sala, S., & Bjørn, A. (2025). Advancing absolute environmental sustainability assessment: Introducing the first technical guidance and a call for feedback. *Journal of Industrial Ecology*, 29, 1933-1935. <https://doi.org/10.1111/jiec.70113>
- Rosenthal, K. I. (1990). *Quantales and Their Applications*. Longman Scientific & Technical.
- Ryberg, M. W., Bjerre, T. K., & Nielsen, P. H. (2021). Absolute environmental sustainability assessment of a Danish utility company relative to the Planetary Boundaries. *Journal of Industrial Ecology*, 25, 765-777. <https://doi.org/10.1111/jiec.13075>
- Sigüenza, C. P., Steubing, B., Tukker, A., & Aguilar-Hernández, G. A. (2021). The environmental and material implications of circular transitions: A diffusion and product-life-cycle-based modeling framework. *Journal of Industrial Ecology*, 25, 563-579. <https://doi.org/10.1111/jiec.13072>

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [JIECodeandData.zip](#)