

The Joule Standard: A Unified Economics of Work, Trust and Intelligence

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

23 July 2026

Abstract

Every economically organized process uses energy, but energy does not provide a natural unit for its consequences. A motor produces mechanical work. An inference system produces candidate task service whose usefulness depends on correctness, adoption, and causal realization. A proof-of-work system expends energy in a protocol that may change settlement risk under a named threat model. A national account records production flows, not all useful outcomes and not the stocks those outcomes protect. This paper synthesizes a ten-paper research program into a common measurement architecture without collapsing these objects into one score. The Joule Standard is a typed, vector-valued account. Its common denominator is a declared, non-overlapping energy inventory. Its numerators retain their native semantics and units. A boundary manifest fixes the functional unit, physical system, time, location, counterfactual, energy stage, allocation rules, outcome functional, attribution rule, uncertainty model, and exclusions. A transduction graph then distinguishes physical efficiency, exergy or useful work, verified service, causal economic realization, avoided loss, settlement assurance, and quality-adjusted machine intelligence. Exact interface matching permits stage yields to compose. Semantic labels do not.

Four formal results govern aggregation. First, addition is undefined for quantities that differ in semantic type, unit, stock-flow status, boundary, or horizon. Second, non-dominating outcome vectors admit opposing rankings under admissible positive weights, so energy normalization cannot reveal a universal scalar order. Third, a strictly positive, monotone scalarization can select a Pareto-efficient allocation, but unsupported frontier points need not be recoverable without convexity. Fourth, adding primary energy, conversion output, and final energy as independent inputs repeats a physical lineage. Related accounting controls prevent AI output already recorded in value added and protected asset stocks from being added to gross domestic product.

The paper supplies a crosswalk for the equations and notation of Papers 1 through 9, a boundary map, incompatibility and double-counting audits, a falsification and

identification agenda, and precise nonclaims. A deterministic standard-library Python implementation preserves exact decimal arithmetic, rejects binary floats at public numeric boundaries, computes typed transduction paths and Pareto frontiers, and tests adversarial substitutions. The proposed standard is therefore a discipline for asking what a joule becomes under declared conditions. It is not an energy theory of value, a conversion of trust or intelligence into thermodynamic units, or a universal league table.

Keywords: energy productivity; exergy; useful work; machine intelligence; causal inference; proof of work; settlement assurance; Pareto frontier; national accounts; composite indicators.

Evidence status. Definitions, type rules, accounting identities, and algebraic propositions are **PROVED** under their stated assumptions. Results produced by the reference implementation and its synthetic fixtures are **COMPUTATIONAL**. Empirical estimates and causal interpretations are **CONDITIONAL**. Claims whose identification conditions have not been met are **OPEN**. The paper does not report a sector ranking.

Contents

1	The synthesis problem	5
1.1	Research question	5
1.2	Contributions	6
1.3	Precise nonclaims	6
2	Prior foundations and intellectual boundary	7
2.1	Energy productivity, exergy, and useful work	7
2.2	Information, computation, and physical limits	7
2.3	Settlement, assurance, and adversarial systems	7
2.4	Causal effects and accounting boundaries	8
2.5	Composite indicators and social choice	8
3	Measurement objects	8
3.1	The boundary manifest	8
3.2	Estimand-family tags	9
3.3	Typed quantities	10
3.4	Energy stages and physical lineage	10
3.5	Physical efficiency and economically realized value	11
4	Typed transduction	12
4.1	Interface composition	12
4.2	Evidence-labeled edges	13
4.3	The assurance branch	14

5	Quality-adjusted intelligence and causal realization	15
5.1	Task service rather than token volume	15
5.2	Quality is not realized value	15
5.3	Identification ladder	16
6	The common vector	16
6.1	Definition	16
6.2	Dominance and the Pareto frontier	17
6.3	Conditional scalarization	17
6.4	No natural common scalar	18
7	Marginal allocation and energy markets	18
7.1	The next joule	18
7.2	Vector allocation	19
7.3	Flexible workloads and grid settlement	19
8	Efficiency and rebound	20
9	National accounting and double-counting control	21
9.1	Preserve the production core	21
9.2	Stocks, flows, and assurance	22
9.3	The audit	22
9.4	Twenty-case cross-program screen	23
10	Cross-paper equation and notation map	24
10.1	Imported conflicts and explicit repairs	25
10.2	Symbol discipline	25
11	Boundary map	27
11.1	Boundary-reversal counterexample	27
12	Incompatibility and non-aggregation	27
12.1	Why avoided loss does not solve every incompatibility	27
12.2	Why exergy does not solve every incompatibility	30
12.3	Why price does not solve every incompatibility	30
13	A synthetic integrated fixture	30
13.1	Independent national closures	31
14	Reference implementation	31

14.1 Executable objects	31
14.2 Refusal behavior	32
14.3 Adversarial test contract	32
15 Identification and falsification agenda	33
15.1 Boundary falsification	33
15.2 Transduction falsification	33
15.3 Quality-adjusted intelligence	33
15.4 Settlement assurance	34
15.5 Causal realization	34
15.6 Marginal allocation and rebound	34
15.7 National-account audit	34
16 Claim ledger	35
17 Governance and reporting protocol	36
17.1 Minimum publication table	37
18 Limitations	37
19 Conclusion	38
A Formal reporting schema	38
B Worked audit cases	39
B.1 Token substitution	39
B.2 Protected exposure added to production	39
B.3 Lifecycle plus operation	39
B.4 Near-zero causal denominator	39
B.5 Weight-sensitive frontier	39
C Reproducibility map	40
D Checklist for an empirical Joule Standard claim	41

1 The synthesis problem

The phrase “value per joule” invites a scalar. The research program developed in Papers 1 through 9 reaches a different conclusion. A common physical resource can organize comparisons, budgets, and marginal allocation without making the resulting services commensurate. Joules can be counted across a declared energy boundary. Mechanical work, verified task service, reliability, expected avoided loss, settlement risk, value added, welfare, and asset exposure cannot simply be counted together.

This distinction matters because energy normalization often hides a change of question. A device-level efficiency experiment asks how much physical output follows from a measured input. A deployment study asks whether a service was useful outside the laboratory. A causal evaluation asks what changed relative to a counterfactual. A settlement study asks how a protocol changes the distribution of attack costs or losses under an adversary model. A national account asks which production flows occurred within an accounting period. The denominator may be expressed in joules in every case, but neither the numerator nor the inferential claim is the same.

The Joule Standard proposed here has three layers.

1. A *typed physical ledger* records energy stages and prevents the same physical lineage from entering a total more than once.
2. A *transduction record* follows energy through physical work, computation, service, adoption, causal realization, or assurance. Every edge states its evidence status.
3. A *decision view* retains a vector of native outcomes, removes dominated allocations, and permits scalarization only after normalization, orientation, and weights are published.

The common theory is therefore a theory of disciplined interfaces. It says when a ratio is interpretable, when two ratios are directly comparable, when stage yields compose, when an addition is legal, and which judgments enter an allocation decision. It does not say that all economically valuable order is a single physical substance.

1.1 Research question

The synthesis asks:

Can work, machine intelligence, settlement assurance, and economic realization share one energy-normalized accounting architecture while retaining their different meanings, units, causal requirements, and institutional boundaries?

The answer is yes for a typed vector account and no for an undeclared scalar account. This answer follows from the foundations in Paper 1, the typed transduction graph in Paper 2, the task-service measure in Paper 3, the scenario-conditioned assurance account in Paper 4, the causal inference chain in Paper 5, the marginal allocation model in Paper 6, the rebound identity in Paper 7, the grid-market model in Paper 8, and the satellite-account structure in Paper 9 [1, 2, 3, 4, 5, 6, 7, 8, 9].

1.2 Contributions

The paper makes six synthesis contributions.

1. It defines a type system for energy flows, physical services, task services, money flows, money stocks, probabilities, rates, and indices.
2. It specifies a boundary manifest that is common across the nine application papers and separates physical scope from causal attribution.
3. It proves non-addition, non-aggregation, composition, and Pareto-selection results and states the assumptions that limit each result.
4. It publishes an equation and notation crosswalk so that a symbol used in one application cannot silently acquire another meaning in the synthesis.
5. It joins physical-lineage and national-account controls in one double-counting audit.
6. It supplies executable reference code and adversarial tests for the generic rules rather than hard-coding any paper-specific answer.

1.3 Precise nonclaims

The Joule Standard makes the following twenty nonclaims explicit.

1. Joules are not a currency, utility unit, welfare unit, or conserved quantity of value.
2. Energy alone does not determine price, intelligence, legitimacy, trust, or social order.
3. Exergy does not extend to semantic value without a task, observer, institution, and value functional.
4. Tokens are not useless operational measures. They remain useful under a fixed tokenizer and service contract, but they are not stable value units.
5. QAIJ is neither general intelligence nor realized economic value.
6. Aggregate proof-of-work energy does not identify security, attack cost, censorship resistance, or losses prevented.
7. Waiting for confirmations does not necessarily cause proportional network energy use.
8. Gross protected exposure is not settlement output or value added.
9. Observed revenue per joule does not identify causal productivity.
10. A causal ratio with a weak or zero energy effect need not have a stable finite interpretation.
11. A high average value per joule does not identify the best recipient of the next joule.
12. The marginal-allocation rule does not automatically apply under market power, non-convexity, omitted constraints, or unknown response curves.
13. Efficient AI does not necessarily lower or raise total energy.
14. Flexible computation is not automatically beneficial to the grid.
15. Direct-current power flow or reserve headroom does not prove voltage, frequency, or dynamic stability.
16. Settlement closure does not prove price formation, truthful bidding, or social optimality.
17. A value-per-joule satellite account neither replaces nor enlarges GDP.
18. Primary, final, operational, and lifecycle energy are not additive stages.
19. Equal weights do not create a politically neutral composite.
20. Passing code tests does not upgrade empirical or causal claims to mathematical proof.

2 Prior foundations and intellectual boundary

2.1 Energy productivity, exergy, and useful work

Energy productivity and energy intensity are established indicators. Patterson distinguishes thermodynamic, physical, economic-thermodynamic, and economic measures, showing that the label “efficiency” does not fix a numerator or boundary [10]. Ayres and Warr place useful work in historical production analysis, while Warr and coauthors construct long-run useful-work series [11, 12]. Exergy analysis asks how much useful work a system could obtain as it approaches equilibrium with a reference environment. It is indispensable for locating physical irreversibility. It does not determine whether the resulting service is wanted, correct, adopted, or welfare improving.

That boundary is important. Let E_b^e denote an energy input, X^{ex} the exergy carried by that input relative to a reference environment, and U^{work} useful physical work. Then physical measures such as U^{work}/E_b^e or $U^{\text{work}}/X^{\text{ex}}$ compare commensurate physical quantities. A monetary realization $\Delta V/E_b^e$ requires prices or another value functional, a baseline, and an attribution design. Replacing U^{work} by ΔV is not a unit conversion. It changes the object being measured.

2.2 Information, computation, and physical limits

Shannon’s information measure concerns uncertainty in messages, not semantic usefulness [13]. Landauer identifies a thermodynamic cost for logically irreversible operations under stated physical conditions [14]. These results motivate a careful connection between energy and computation, but neither provides an economic valuation of an inference. A token may carry different semantic content under different tokenizers. A correct answer may be useless because it arrives late, cannot be verified, or is not adopted. A useful answer may have large economic effects in one context and none in another.

Machine-learning evaluation has accordingly moved beyond raw throughput. Benchmark suites report accuracy, robustness, calibration, latency, and system performance under declared tasks and hardware [15, 16]. Paper 3 uses those ideas to define a quality-adjusted task service. Paper 5 then adds adoption and causal realization. The synthesis preserves this separation: quality adjustment is still not a welfare functional.

2.3 Settlement, assurance, and adversarial systems

Nakamoto’s proof-of-work construction joins a chain-selection rule to an economic cost of rewriting history [17]. Formal analyses give security properties under explicit network and adversary assumptions [18]. Economic analyses show that attack incentives depend on rewards, rents, market structure, outside options, and the value at risk [19]. Energy expenditure is one input to a security mechanism. It is not a complete measure of censorship resistance, finality, decentralization, or expected loss.

Paper 4 therefore treats assurance as a vector and permits avoided loss per joule only

for a matched intervention under a named scenario. This synthesis retains that constraint. The scalar ALPJ is a conditional monetary decision statistic, not a physical law and not a synonym for trust.

2.4 Causal effects and accounting boundaries

An incremental economic numerator is a counterfactual contrast. Potential outcomes distinguish such effects from observed differences [20, 21, 22]. When the denominator is itself an estimated intervention effect, a ratio of effects can become weakly identified near zero. Paper 5 makes this issue explicit and recommends joint uncertainty sets or Fieller-type logic rather than mechanically dividing two noisy point estimates.

At economy scale, the SNA records production flows and the SEEA connects physical environmental flows to economic accounts [23, 24, 25]. Gross supply-use tables may show a primary carrier, its conversion output, and final use. Those appearances are useful for balancing accounts, but their sum is not a total energy input. Paper 9 contributes an energy-stage bridge and keeps AI outcomes and settlement assurance in memorandum panels. The present paper turns that architecture into a general typed audit.

2.5 Composite indicators and social choice

Composite indicators require normalization, weights, aggregation, and sensitivity analysis [26]. Equal weights are still weights. A data-derived weighting method encodes a statistical criterion rather than discovering a natural social objective. More broadly, heterogeneous welfare dimensions cannot be ordered without assumptions about standing, tradeoffs, and aggregation [27, 28]. Energy normalization does not remove those assumptions. It changes the resource basis of the comparison.

3 Measurement objects

3.1 The boundary manifest

Paper 1 defines a measurement tuple

$$\mathcal{M} = (b, c_0, W, \tau_H, e^{\text{conv}}, a^{\text{attr}}, \Pi^{\text{unc}}), \quad \text{VPJ}_{\mathcal{M}}(s) = \frac{\mathbb{E}_{\Pi^{\text{unc}}}[\Delta W(s \mid c_0)]}{E_b^{e^{\text{conv}}}(s; \tau_H)}.$$

The synthesis expands this tuple into an operational manifest. For a study r , define

$$\mathfrak{M}_r = (f^{\text{unit}}, p^{\text{system}}, \tau_H, \ell^{\text{geo}}, c_0, \text{stage}^E, a^{\text{shared}}, i^{\text{failure}}, m^{\text{embodied}}, W, a^{\text{attr}}, \Pi^{\text{unc}}, z^{\text{exclude}}),$$

The entries are, in order, the functional unit, physical system, temporal interval, location, counterfactual, energy stage, shared-infrastructure rule, idle and failure rule, embodied-energy scope, outcome functional, attribution rule, uncertainty model, and exclusions.

Every denominator is accompanied by the record

$$D_E = (\text{stage, physical boundary, geography, interval,} \\ \text{allocation, idle/failure treatment, embodied scope, lineages}).$$

An equation that uses a short symbol E inherits this record from its manifest. The short symbol is typography, not permission to omit the fields.

The expansion does not change the theory in Paper 1. It makes recurrent implementation choices explicit. A direct comparison between two reported ratios passes the boundary gate only if the relevant fields match or a documented bridge maps one field to the other. A common string such as “facility energy” is not a bridge. The analyst must show that metering, allocation, time, location, and inclusion rules align.

Definition 3.1 (Direct comparability). Two measurements r and r' are directly comparable when their functional units, physical systems, time intervals, locations, counterfactuals, energy stages, infrastructure allocations, idle and failure rules, embodied scopes, outcome functionals, attribution rules, and uncertainty models match, and their exclusions are identical. A bridged comparison is a different object and must report the bridge transformation and its uncertainty.

Proposition 3.2 (Typed comparison class). *Direct comparability is an equivalence relation within a fixed functional unit and manifest. A comparison across classes is valid only through a documented order-preserving bridge for the reported claim.*

Proof. Equality of the required manifest fields is reflexive, symmetric, and transitive. A bridge supplies a new mapped object; without it, at least one field differs and the objects lie in different classes. Requiring the bridge to preserve the order relevant to the claim prevents a unit transformation from being mistaken for an arbitrary reranking. $\square$

3.2 Estimand-family tags

Identical printed units can conceal different mathematical objects. Every energy-normalized result in the synthesis therefore carries one of six estimand-family tags:

Level intensity. $I = X/E_b^e$, where $E_b^e > 0$ is a matched inventory. Paper 1’s basic VPJ is in this family when its expected numerator is divided by a fixed positive inventory.

Ratio of expectations. $I_{\text{RE}} = \mathbb{E}[X]/\mathbb{E}[E_b^e]$. Paper 3’s QAIJ is in this family.

Mean of unit ratios. $I_{\text{MR}} = \mathbb{E}[X/E_b^e]$. This is generally not I_{RE} .

Causal ratio of effects. $\theta^{\text{causal}} = \Delta_V^{\text{ATE}}/\Delta_E^{\text{ATE},b}$, where both effects describe the same intervention. The denominator can be positive or negative and the ratio does not exist at zero.

Marginal derivative. $m_{i,n,t} = \partial B_i/\partial e_{i,n,t}$, evaluated at an operating point under constraints. This is not an average intensity.

Descriptive national intensity. A matched period production or outcome flow divided by a matched period energy flow. This is not causal absent a separate design.

Proposition 3.3 (Estimand-family non-equivalence). *A positive level intensity, ratio of expectations, mean of unit ratios, causal ratio of effects, and marginal derivative are generally unequal. They can have different signs and existence conditions.*

Proof. For two units with outcomes $(0, 8)$ and energies $(2, 8)$, the ratio of means is $8/10 = 0.8$, while the mean of ratios is $(0/2 + 8/8)/2 = 0.5$. For a value function with $V(10) = 100$ and $V(11) = 101$, average value at 10 J is 10 per J while the local finite slope is 1 per J. An intervention can save energy, giving $\Delta_E^{\text{ATE},b} < 0$, even though every gross inventory is positive. At $\Delta_E^{\text{ATE},b} = 0$, the causal ratio is undefined. These constructions establish non-equivalence, possible sign differences, and different domains. $\square$

3.3 Typed quantities

An amount alone is not enough for an accounting operation. The synthesis type of a quantity ξ is

$$\text{type}(\xi) = (\sigma, u, k, b, \tau),$$

where σ is its semantic name, u its unit, k its accounting kind, b its boundary identifier, and τ its horizon identifier.

The accounting kinds used here are energy flow, physical service, task service, money flow, money stock, probability, rate, and index. “USD” does not erase the distinction between a flow of value added during a quarter and a protected asset stock at quarter end. “Count” does not erase the distinction between generated tokens and verified completed tasks. “Joule” does not erase the distinction between primary supply, conversion output, final energy, facility operation, and a lifecycle inventory.

Definition 3.4 (Typed addition). For quantities $\xi_1, \dots, \xi_n$, the sum $\sum_i \xi_i$ is defined in the Joule Standard only if

$$\text{type}(\xi_1) = \dots = \text{type}(\xi_n).$$

The equality includes semantic type, unit, accounting kind, boundary, and horizon.

Theorem 3.5 (Non-addition of heterogeneous outcomes). *If two quantities differ in at least one component of $\text{type}(\xi)$, their addition is undefined in the typed algebra. No multiplication or division by a common energy quantity makes the addition defined.*

Proof. Typed addition is defined only on a common carrier. Suppose $\text{type}(\xi_1) \neq \text{type}(\xi_2)$. Division by the same positive energy E_b^e appends a common denominator unit and boundary but preserves the unequal numerator types. Thus $\text{type}(\xi_1/E_b^e) \neq \text{type}(\xi_2/E_b^e)$. The premise for addition still fails. A mapping into a common scalar space can be supplied, but that mapping is an additional value functional, not a consequence of energy normalization. $\square$

Remark 3.6. The theorem is intentionally syntactic. It prevents an analyst from claiming that an undeclared sum already has meaning. It does not prohibit a declared model that maps different inputs to a common welfare, cost, or index space. Such a model must publish its assumptions and is evaluated separately.

3.4 Energy stages and physical lineage

For an energy component j , record

$$\mathbf{e}_j = (\text{id}_j, \chi_j^{\text{lineage}}, \text{stage}_j^E, E_j),$$

where χ_j^{lineage} is a physical-lineage identifier and stage_j^E is an energy stage. The standard distinguishes device operational, system operational, facility operational, uniquely allocated embodied increment, lifecycle, primary, conversion output, final, and marginal-grid energy. A total is legal only over a non-overlapping set of lineages at a coherent stage.

Consider 100 J of primary fuel entering a converter that produces 35 J of electricity, of which 32 J reaches final use. A physical supply-use table may display all three entries. The energy input is not 167 J. The entries are linked stages of one lineage. A bridge may report conversion loss 65 J and network loss 3 J, but the stages are not independent resources.

Theorem 3.7 (Transformation-lineage double counting). *Let a transformation chain have primary input $E_{\text{primary,supply}}$, conversion output O_{energy} , and final use $E_{\text{final,use}}$, with $0 \leq E_{\text{final,use}} \leq O_{\text{energy}} \leq E_{\text{primary,supply}}$. If the three entries carry one physical lineage, then their sum repeats that lineage. A non-overlapping input total chooses one stage or uses a bridge identity, such as*

$$E_{\text{primary,supply}} = E_{\text{final,use}} + (E_{\text{primary,supply}} - O_{\text{energy}}) + (O_{\text{energy}} - E_{\text{final,use}}).$$

Proof. The amount O_{energy} is produced from a subset of the energy represented by $E_{\text{primary,supply}}$, and $E_{\text{final,use}}$ is delivered from a subset of O_{energy} . They are not disjoint inputs. The right side of theorem 3.7 partitions the lineage into final use and two loss terms. Adding the three stage levels instead contains delivered energy three times, conversion loss once, and network loss twice. $\square$

Proposition 3.8 (Denominator non-substitution). *Device, system-wall, facility-operational, embodied-increment, lifecycle, primary, final, and marginal-grid energy are distinct denominator types. Equality of their joule units is insufficient for substitution. A valid bridge must identify the physical lineage, additions, losses, allocation, geography, and interval without overlap.*

Proof. Each denominator differs in at least one stage or boundary component of sections 3.1 and 3.3. Typed substitution therefore fails. A bridge may map between them, but theorem 3.7 requires the mapped physical components to form a non-overlapping partition. $\square$

3.5 Physical efficiency and economically realized value

The accepted program contains two local AI chains that use O , U , and T differently. The synthesis resolves that collision with the canonical optional-node graph

$$E_b^e \longrightarrow C_{\text{named}} \longrightarrow O_{\text{emit}} \longrightarrow U_{\mathcal{P}} \longrightarrow A^{\text{adopt}} \longrightarrow \Delta \mathbf{o}_{c_0} \longrightarrow W(\Delta \mathbf{o}_{c_0}).$$

Here O_{emit} is immediate output and $U_{\mathcal{P}}$ is a verified useful outcome under protocol $\mathcal{P}$. For a token-emitting system, the representation statistic

$$T_{\kappa}^{\text{tok}} = |\kappa(O_{\text{emit}})|$$

is attached to the immediate-output node. It is not substituted for that node and does not become the verified outcome. An adoption node can be bypassed only by a named bridge that explains why produced service equals institutionally used service in the application.

Not every application uses every node. A physical-work branch is

$$E_b^e \longrightarrow X^{\text{ex}} \longrightarrow U^{\text{work}} \longrightarrow \text{delivered physical service} \longrightarrow \Delta \mathbf{o}_{c_0}.$$

A proof-of-work branch proceeds from hashing energy through hashrate, productive share, protocol rules, a named risk measure, economic deterrence, and a loss model. A grid branch proceeds from power integrated over an interval to feasible dispatch, reserve, and cash settlement. The point is not to force one causal graph on every system. It is to prevent a jump from an upstream observable to a downstream claim without a declared edge.

4 Typed transduction

4.1 Interface composition

Paper 2 represents an economy as a typed transduction graph. For a path

$$X_0^{\text{int}} \xrightarrow{\mathcal{T}_1^{\text{stage}}} X_1^{\text{int}} \xrightarrow{\mathcal{T}_2^{\text{stage}}} \dots \xrightarrow{\mathcal{T}_n^{\text{stage}}} X_n^{\text{int}},$$

define a stage yield $y_k^{\text{stage}} = X_k^{\text{int}} / X_{k-1}^{\text{int}}$ only after the output of $\mathcal{T}_{k-1}^{\text{stage}}$ matches the input of $\mathcal{T}_k^{\text{stage}}$ in amount and type. Then

$$\frac{X_n^{\text{int}}}{X_0^{\text{int}}} = \prod_{k=1}^n y_k^{\text{stage}}.$$

Proposition 4.1 (Exact typed composition). *If every adjacent interface has identical semantic type, unit, accounting kind, boundary, horizon, and amount, then section 4.1 holds exactly for a deterministic path with positive intermediate inputs.*

Proof. The product telescopes:

$$\prod_{k=1}^n \frac{X_k^{\text{int}}}{X_{k-1}^{\text{int}}} = \frac{X_1^{\text{int}}}{X_0^{\text{int}}} \frac{X_2^{\text{int}}}{X_1^{\text{int}}} \dots \frac{X_n^{\text{int}}}{X_{n-1}^{\text{int}}} = \frac{X_n^{\text{int}}}{X_0^{\text{int}}}.$$

Exact interface equality makes each cancellation a cancellation of the same typed quantity. Without that equality, the written cancellation is only a symbolic resemblance. $\square$

Proposition 4.2 (Optional-node transduction). *Any selected path through section 3.5 composes exactly only when its adjacent typed amounts match and no branch charges a shared input more than once. The token statistic in section 3.5 can enter a token-specific stage, but it cannot replace O_{emit} . An omitted adoption node requires a declared type-preserving bridge.*

Proof. Exact path composition follows from proposition 4.1. A duplicated branch violates lineage uniqueness. Token count and emitted content have different semantic types, so substitution violates interface equality. The same argument applies to a silent relabeling of produced service as adopted service. $\square$

Proposition 4.3 (Stochastic nonfactorization). *Pathwise stage yields telescope, but in general*

$$\mathbb{E}\left[\prod_k Y_k\right] \neq \prod_k \mathbb{E}[Y_k].$$

Proof. Let $Y_1 = Y_2 = 0$ with probability $1/2$ and $Y_1 = Y_2 = 2$ with probability $1/2$. Then $\mathbb{E}[Y_1 Y_2] = 2$, while $\mathbb{E}[Y_1]\mathbb{E}[Y_2] = 1$. The pathwise product remains exact in each state. The failure arises from dependence, not from cancellation. $\square$

Correlation, selection, missingness, and endogenous adoption therefore matter. The reference implementation composes exact observed or modeled stages. It does not infer stochastic independence.

4.2 Evidence-labeled edges

Each transduction edge carries one of four labels.

PROVED A definition or algebraic consequence follows from stated assumptions.

COMPUTATIONAL An exact fixture or tested algorithm produces the result.

CONDITIONAL The interpretation depends on an empirical design, scenario, or behavioral assumption.

OPEN The relevant edge is proposed but not identified by the available evidence.

An energy meter can support a computational statement about facility electricity. A benchmark can support a conditional statement about verified task service on its sampled task distribution. Neither alone identifies economic value. The service-to-adoption and adoption-to-outcome edges need their own evidence. This prevents evidence from flowing farther down a graph than the design supports.

For terminal claims, the synthesis publishes the evidence-status vector

$$\mathbf{s}_{\text{ev}} = (\text{algebra}, \text{code}, \text{measurement}, \text{identification}, \text{transport}, \text{normative aggregation}).$$

For example, a tested avoided-loss calculation may have

$$\mathbf{s}_{\text{ev}} = (\text{PROVED}, \text{COMPUTATIONAL}, \text{CONDITIONAL}, \\ \text{OPEN}, \text{OPEN}, \text{CONDITIONAL}).$$

A proved algebra coordinate never upgrades an open identification coordinate.

For a composed transduction path, the reference implementation uses the conservative stage-status order

$$\text{PROVED} \prec \text{COMPUTATIONAL} \prec \text{CONDITIONAL} \\ \prec \text{OPEN} \prec \text{OBSTRUCTED}.$$

The path status is the join, which is the most limiting status present. Thus a COMPUTATIONAL meter edge followed by an OPEN service-to-value edge yields an OPEN terminal path. The multi-coordinate vector in section 4.2 remains visible because stage status cannot substitute for measurement, identification, transport, or normative status.

4.3 The assurance branch

Settlement assurance is not naturally one scalar node. For protocol $\mathcal{P}^{\text{settle}}$ and scenario ω , define an assurance record

$$\mathcal{A}(\mathcal{P}^{\text{settle}}, \omega) = (\alpha_A^{\text{hash}}, Pr^{\text{reorg}}, Loss_P, Cost_A^{\text{net}}, M_A, \mathbf{k}),$$

where α_A^{hash} is effective adversarial share, Pr^{reorg} a named reorganization risk, $Loss_P$ loss conditional on success, $Cost_A^{\text{net}}$ modeled net private attack cost, M_A attack margin, and $\mathbf{k}$ non-energy constraints. Each component keeps its own unit.

Paper 4 locally declares both an energy-account interval and a post-confirmation attack window. The synthesis names them τ_E^{hash} and τ_A^{attack} . Hashing energy is

$$E_H(\tau_E^{\text{hash}}) = H_H \epsilon_H^{\text{J/TH}} \tau_E^{\text{hash}}.$$

Risk is evaluated over τ_A^{attack} . The restriction $\tau_E^{\text{hash}} = \tau_A^{\text{attack}}$ may be imposed for a particular fixture, but it is never implicit.

Paper 4’s $Loss_P$ is loss if the modeled event succeeds. Paper 9’s local “gross protected exposure” is instead an exposure X_P^{gross} . The synthesis relates them only through a declared loss severity $s_P^{\text{loss}} \in [0, 1]$:

$$Loss_P = X_P^{\text{gross}} s_P^{\text{loss}}.$$

Treating the two as equal assumes 100 percent severity.

A monetary avoided-loss ratio is permitted only for matched risks and a named incremental intervention:

$$\text{ALPJ}(\mathcal{P}^{\text{settle}}; \omega) = \frac{Loss_P \{Pr_0^{\text{reorg}}(\omega) - Pr_1^{\text{reorg}}(\omega)\}}{\Delta E_P}, \quad \Delta E_P > 0.$$

The risk probabilities are matched under the baseline and intervention.

The expression is algebraically well defined. Its interpretation is **CONDITIONAL**. Equal energy does not imply equal risk, attack cost, decentralization, censorship resistance, or loss prevention. The protocol, network topology, hardware market, reward schedule, adversary, time horizon, and asset exposure all enter the scenario.

Proposition 4.4 (Assurance insufficiency). *Equal proof-of-work energy can coexist with different hashrate, scenario risk, attack margin, and protected-loss outcomes. The ratio in section 4.3 is causal only when its loss and energy increments arise from the same intervention.*

Proof. Two miners can each use 100 J while one uses hardware at 1 J/TH and the other at 2 J/TH, yielding 100 TH and 50 TH. Their effective shares and reorganization risks can therefore differ at equal energy. Propagation, access, rewards, and attacker benefit can further change risk and margin. The causal condition follows from proposition 5.1: unmatched numerator and denominator contrasts do not define one intervention effect ratio. $\square$

5 Quality-adjusted intelligence and causal realization

5.1 Task service rather than token volume

Paper 3 defines a quality-adjusted service contribution for task i :

$$s_i^{\text{svc}} = \tilde{d}_i^{\text{task}} u_i^{\text{svc}} q_i^{\text{AI}} c_i^{\text{correct}} \alpha_i^{\text{svc}} \ell_i^{\text{lat}},$$

where $\tilde{d}_i^{\text{task}}$ is a normalized difficulty weight, u_i^{svc} usefulness, q_i^{AI} quality, c_i^{correct} correctness, α_i^{svc} reliability, and ℓ_i^{lat} latency or timeliness. The factors and their scales must be declared. For evaluation protocol $\mathcal{P}^{\text{eval}}$,

$$\text{QAIJ}_{\mathcal{P}^{\text{eval}}} = \frac{\mathbb{E}[\sum_{i \in \mathcal{P}^{\text{eval}}} s_i^{\text{svc}}]}{\mathbb{E}[E_b^e]}.$$

The multiplicative form is a preregistered grading aggregation. It is not a physical conservation law and not the interface-cancellation identity in section 4.1. Raw task difficulty does not enter the dimensionless product without the declared normalization that produces $\tilde{d}_i^{\text{task}}$.

This measure is not “intelligence” in the unrestricted psychological or philosophical sense. It is quality-adjusted service on a named task distribution, evaluated with declared scoring and energy boundaries. A tokenizer change can alter token count while leaving user-visible content unchanged. Verbosity can increase tokens without improving correctness or usefulness. Task mix can reverse system rankings. The task distribution is therefore part of the estimand.

5.2 Quality is not realized value

A high QAIJ can coexist with low realized value. The system may not be adopted, may be used for a low-stakes task, or may displace an equally good baseline. Paper 5’s local inference chain is translated into the canonical graph in section 3.5. Its token count is T_κ^{tok} , its verified outcome is $U_{\mathcal{P}}$, and its adoption node is A^{adopt} . None of those nodes is silently renamed in the synthesis.

Paper 5 locally calls a component sum E_i^{wall} while permitting a separate facility-overhead component. The synthesis does not inherit that label. It uses

$$\begin{aligned} E_i^{\text{system-wall}} &= E_i^{\text{device}} + E_i^{\text{host}} + E_i^{\text{memory}} + E_i^{\text{network,inside}}, \\ E_i^{\text{facility-operational}} &= E_i^{\text{system-wall}} + E_i^{\text{cooling}} + E_i^{\text{power-conversion}} + E_i^{\text{shared,facility}}, \end{aligned}$$

with component lineages checked for overlap. A general component sum is named $E_i^{\text{operational},b}$, where b states which of these boundaries applies. Facility overhead is never hidden inside a quantity still labeled system-wall energy.

Let Δ_V^{ATE} be an intervention effect on realized value and $\Delta_E^{\text{ATE},b}$ its effect on operational energy at boundary b . The causal value-per-energy estimand is

$$\theta^{\text{causal}} = \frac{\Delta_V^{\text{ATE}}}{\Delta_E^{\text{ATE},b}}, \quad \Delta_E^{\text{ATE},b} \neq 0.$$

Both effects must refer to the same intervention, boundary, counterfactual, horizon, and population. Dividing an observational revenue difference by a metered energy effect does not meet that condition.

The executable identification key is

(intervention, boundary, counterfactual, horizon, population).

Population is not inferred from the other four fields.

Proposition 5.1 (Matched causal ratio). *The ratio in section 5.2 is a causal effect ratio only if numerator and denominator identify effects of the same intervention against the same counterfactual over the same boundary, population, and horizon. If $\Delta_E^{\text{ATE},b} = 0$, the finite ratio is undefined. If $\Delta_E^{\text{ATE},b}$ is weakly separated from zero, ordinary delta-method intervals can be misleading.*

Proof. The first claim follows from the definition of a ratio of intervention effects. Unmatched contrasts do not share an intervention estimand. Division by zero is undefined. Near zero, the ratio map is highly nonlinear, so uncertainty in the denominator can create unbounded or disjoint confidence sets. Joint inference or Fieller-type constructions are then required. $\square$

5.3 Identification ladder

The synthesis uses an identification ladder rather than one evidence label for an entire system.

1. **Metering:** identify operational energy for the declared physical boundary.
2. **Service evaluation:** estimate correctness, usefulness, reliability, latency, and task difficulty on a declared task distribution.
3. **Adoption:** observe whether the service enters the relevant workflow and whether failed attempts or retries are retained.
4. **Outcome effect:** identify the counterfactual change in a physical, monetary, or welfare outcome.
5. **Attribution:** allocate joint effects and shared burdens without adding the same outcome twice.
6. **External validity:** state the population, time, and institutions to which the result may travel.

Failure at a lower rung does not invalidate observations above it, but it limits their interpretation. For example, a well-measured energy denominator and benchmark accuracy remain useful engineering evidence even when economic realization is OPEN.

6 The common vector

6.1 Definition

For a feasible system or allocation s under manifest $\mathcal{M}$, define

$$\mathbf{J}_{\mathcal{M}}(s) = \left(E_b^e(s); \frac{U^{\text{work}}(s)}{E_b^e(s)}, \frac{S_{\text{task}}(s)}{E_b^e(s)}, \frac{\Delta V(s)}{E_b^e(s)}, \mathcal{A}(\mathcal{P}^{\text{settle}}, \omega), \frac{\Delta G_{c_0}^{\text{grid}}(s)}{E_b^e(s)} \right).$$

The displayed coordinates are illustrative, not mandatory. They represent useful work, quality-adjusted task service, causal monetary realization, scenario-conditioned assurance, and a declared environmental or grid effect. Every coordinate names its native numerator unit. The semicolon after energy is deliberate: energy is also a resource to minimize, not merely a denominator that disappears after ratios are formed.

Definition 6.1 (Joule vector). A Joule vector is an energy amount and an ordered set of outcome axes. Each axis has a semantic name, native unit, orientation, boundary, horizon, and evidence status. Two vectors are directly Pareto comparable only when their axis signatures and manifests match.

The vector representation can report “MJ of useful work per J”, “verified task-equivalent per J”, and “2026 USD of causally attributed avoided loss per J” in adjacent columns. It does not add the columns. If assurance is expressed as a probability or attack margin, that coordinate remains a probability or a scenario-specific economic quantity. If the causal value effect is not identified, the vector may report the upstream service coordinate while marking the value coordinate OPEN.

6.2 Dominance and the Pareto frontier

Let $\mathcal{F}$ be a feasible set of allocations. Energy is weakly preferred lower. For every beneficial outcome axis j , z_j is weakly preferred higher; cost or risk axes are oriented lower. An allocation s dominates r when it is weakly better on energy and every outcome and strictly better on at least one coordinate. The Pareto frontier is

$$\mathcal{P} = \{s \in \mathcal{F} : \nexists r \in \mathcal{F} \text{ such that } r \succ s\}.$$

Dominance is a partial order. It can reject a system that uses more energy and produces no better outcome. It cannot rank one system that provides more verified task service against another that provides more useful work or assurance.

6.3 Conditional scalarization

For a decision maker d , let $n_j(o_j) \in [0, 1]$ be a published monotone normalization rule and let $\omega_{dj}^{\text{index}} \geq 0$, $\sum_j \omega_{dj}^{\text{index}} = 1$, be published weights. Include energy as an oriented axis, for example $n_E(E_b^e) = (E_{\max} - E_b^e)/(E_{\max} - E_{\min})$. A conditional view is

$$S_d(s) = \sum_j \omega_{dj}^{\text{index}} n_j(o_j(s)).$$

The subscript d matters. The score belongs to a decision rule, not to physics.

Theorem 6.2 (Positive monotone scalarization selects a Pareto point). *Let $\mathcal{F}$ be finite and nonempty. Suppose every coordinate is oriented so that higher normalized values are better, every n_j is strictly increasing in that orientation, and every $\omega_j^{\text{index}} > 0$. Any maximizer of section 6.3 is Pareto efficient.*

Proof. Suppose a maximizer s^* were dominated by r . Every normalized coordinate of r would be at least that of s^* , and at least one would be strictly greater. Strictly positive weights would imply $S(r) > S(s^*)$, contradicting maximality. $\square$

Limitation 6.3. Zero weights weaken the result because a score can ignore the axis on which a candidate is dominated. A deterministic tie rule can then select a dominated point. The implementation permits zero weights as a declared analytic view, but the theorem applies only when every compared objective has positive weight.

Limitation 6.4. The converse requires more structure. With a convex attainable outcome set, supporting hyperplane results can recover supported frontier points using nonnegative linear weights. In a discrete or nonconvex set, some Pareto points are unsupported and cannot maximize a weighted sum. Reporting only weighted-sum optima can therefore hide feasible tradeoffs.

6.4 No natural common scalar

Theorem 6.5 (Conditional rank reversal). *Let systems a and b have matched manifests and a common positive energy denominator. Suppose their normalized outcome vectors do not dominate one another. If a is strictly better on one axis and b is strictly better on another, there exist admissible strictly positive weight vectors ω and ω' such that*

$$S_\omega(a) > S_\omega(b) \quad \text{and} \quad S_{\omega'}(b) > S_{\omega'}(a).$$

Proof. Place weight $1 - \varepsilon$ on an axis where a is better and distribute the remaining positive mass across other axes. For sufficiently small $\varepsilon > 0$, the strict advantage on that axis determines the sign. Repeat with weight concentrated on an axis where b is better. Because the number of axes is finite and normalized differences are bounded, both weight vectors can remain strictly positive. $\square$

The theorem is constructive. It does not say that all weights are equally reasonable. Evidence, rights, budgets, law, or institutional mandates can constrain admissible weights. It says that the ordering does not follow from joules alone. A claim of one “Joule Standard score” must therefore identify the decision maker, normalization, weights, and robustness region.

7 Marginal allocation and energy markets

7.1 The next joule

Paper 6 studies the marginal value of energy in an automated economy. Let $e_{i,n,t}$ be MWh assigned to service i , node n , and time t . Let $B_i^{\text{work}}(\mathbf{e}_i)$ be a declared benefit function and let grid and operational constraints define $\mathcal{F}$. A scalar social-planner problem can be written

$$\max_{\mathbf{e} \in \mathcal{F}} \sum_i B_i^{\text{work}}(\mathbf{e}_i) - \sum_{i,n,t} k_{n,t}^{\text{grid}} e_{i,n,t} - \sum_i \mathcal{R}_i^{\text{risk}}(\mathbf{e}_i),$$

where $\mathcal{R}_i^{\text{risk}}$ is a declared risk or external-cost functional. Under convexity and regularity, an interior allocation satisfies a marginal condition

$$\frac{\partial B_i^{\text{work}}}{\partial e_{i,n,t}} = k_{n,t}^{\text{grid}} + \lambda_{i,n,t}^{\text{KKT}} + \frac{\partial \mathcal{R}_i^{\text{risk}}}{\partial e_{i,n,t}},$$

where $\lambda_{i,n,t}^{\text{KKT}}$ collects the opportunity costs of binding service and feasibility constraints. A declared nodal market price, if used, is $p_{n,t}^{\text{nodal}}$, not this multiplier.

This condition does not produce a universal sector ranking. The marginal benefit function depends on the service, state, counterfactual, and decision objective. The next joule may be most valuable for a delayed manufacturing operation at one node, a latency-sensitive inference at another, or a settlement workload under an elevated threat scenario. Scarcity prices and shadow costs are local and time dependent.

Proposition 7.1 (Average-marginal separation). *A ranking by average value per joule does not identify the ranking of the next joule.*

Proof. At $E = 10$, let system a have value $V_a(10) = 100$ and local increment $V_a(11) - V_a(10) = 1$. Let system b have $V_b(10) = 60$ and local increment $V_b(11) - V_b(10) = 5$. Average intensities rank a above b , since $10 > 6$, while local slopes rank b above a , since $5 > 1$. Diminishing returns, node-time costs, nonconvexity, and opportunity cost can widen this difference. $\square$

7.2 Vector allocation

When the benefit components are not commensurate, solve a vector optimization problem first:

$$\max_{\mathbf{e} \in \mathcal{F}} (U^{\text{work}}(\mathbf{e}), S_{\text{task}}(\mathbf{e}), -Pr^{\text{reorg}}(\mathbf{e}), \Delta \mathbf{o}_{c_0}(\mathbf{e}), -\text{CO}_2 \mathbf{e}(\mathbf{e})), \quad \min E_b^e(\mathbf{e}).$$

The output is a Pareto frontier. A decision maker may then apply section 6.3, an epsilon-constraint method, lexicographic priorities, or legal constraints. The choice must remain visible. In particular, monetizing every effect is one optional model. It is not required by the Joule Standard.

7.3 Flexible workloads and grid settlement

Paper 8 represents AI and proof-of-work workloads as heterogeneous feasible sets in a nodal dispatch model. Both can be price-responsive, but their constraints differ. AI jobs may carry deadlines, data locality, checkpoint costs, latency targets, or quality degradation. Proof-of-work loads may have startup behavior, pool constraints, hardware efficiencies, and protocol revenue exposure.

Paper 8 uses v_w locally for both a downward ramp limit and an audited true value. The synthesis imports neither collision. It writes $r_j^\downarrow$ for a load's downward ramp limit and ν_j^{true} for audited true private value. Submitted bid value is b_j^{submit} , which remains distinct from ν_j^{true} and from realized downstream social value.

Reserve capability is a power quantity. A reserve offer is stated here in currency per MW-hour of reserved capability, equivalently currency per MWh-reserved for a one-hour interval. For an interval of δ_t^{hrs} hours, reserve payment multiplies the MW quantity by δ_t^{hrs} before applying a currency-per-MWh-reserved price. This resolves Paper 8's local currency-per-MW versus currency-per-MWh label inconsistency and prevents power from being treated as energy without integration.

For a direct-current dispatch approximation, node balance and line limits can be expressed as

$$\begin{aligned} p_{n,t}^{\text{gen}} - D_{n,t}^{\text{fixed}} - \sum_j p_{j,n,t}^{\text{load}} &= \sum_m b_{nm}^{\text{sus}}(\phi_{n,t} - \phi_{m,t}), \\ -\bar{F}_{nm} &\leq b_{nm}^{\text{sus}}(\phi_{n,t} - \phi_{m,t}) \leq \bar{F}_{nm}. \end{aligned}$$

A complete market settlement also needs revenue adequacy and an explicit reserve or uplift rule. The canonical system closure is

$$\text{Payment}_D - \text{Credit}_G = K^{\text{network}},$$

where Payment_D , Credit_G , and residual network or constraint rents K^{network} balance under the modeled settlement.

The synthesis uses the grid model to locate marginal energy cost and constraints. It does not treat a dispatch solution as proof of alternating current feasibility, voltage security, transient stability, or deliverable reserves. Those require additional models.

Proposition 7.2 (Market-accounting separation). *Lossless dispatch balance and nodal cash closure do not imply truthful bids, supporting prices, competitive equilibrium, dynamic stability, emissions benefit, or realized service value.*

Proof. Balance follows from feasible injections and withdrawals. Cash closure follows from applying a complete declared price vector to that balanced dispatch. Both identities can hold when bids differ from true private value, when the price vector is externally supplied, when omitted dynamic constraints fail, or when the dispatched service has no downstream realized value. Each stronger claim therefore needs an additional premise or model. $\square$

8 Efficiency and rebound

Let $Q^{\text{cog}}(\eta_Q^{\text{svc/J}})$ be total cognitive service demanded when efficiency is $\eta_Q^{\text{svc/J}}$ service units per joule. Direct facility-operational energy use is

$$E_{\text{AI}}^{\text{facility}}(\eta_Q^{\text{svc/J}}) = \frac{Q^{\text{cog}}(\eta_Q^{\text{svc/J}})}{\eta_Q^{\text{svc/J}}}.$$

Differentiating gives the identity from Paper 7:

$$\frac{d \ln E_{\text{AI}}^{\text{facility}}}{d \ln \eta_Q^{\text{svc/J}}} = \frac{d \ln Q^{\text{cog}}}{d \ln \eta_Q^{\text{svc/J}}} - 1 = \epsilon_{Q^{\text{cog}}, \eta_Q} - 1.$$

Energy falls when the service elasticity with respect to efficiency is below one, remains constant at one, and rises above one. For finite changes,

$$\frac{E_{\text{AI},1}^{\text{facility}}}{E_{\text{AI},0}^{\text{facility}}} = \frac{Q_1^{\text{cog}}/Q_0^{\text{cog}}}{\eta_{Q,1}^{\text{svc/J}}/\eta_{Q,0}^{\text{svc/J}}}.$$

These are accounting identities, not empirical estimates of rebound. The causal effect of an efficiency improvement on service demand can include lower prices, new use cases, quality changes, complementary capital, organizational adaptation, and economy-wide income effects. System boundaries again matter. A model-level reduction in joules per task may coincide with higher facility-level electricity, and facility electricity may differ from economy-wide primary energy. The Joule Standard requires reporting the efficiency numerator, price path, demand response, energy stage, time horizon, and displaced baseline before classifying rebound.

Proposition 8.1 (Rebound scope). *Under fixed service and direct-energy definitions, section 8 holds. It does not determine economy-wide energy without the response of non-AI energy.*

Proof. The direct identity follows by differentiating $\ln E_{\text{AI}}^{\text{facility}} = \ln Q^{\text{cog}} - \ln \eta_Q^{\text{svc/J}}$. For non-overlapping total energy $E_{\text{total}} = E_{\text{AI}}^{\text{facility}} + E_{\text{other}}$, the derivative also contains the response of E_{other} . The direct service elasticity does not identify that term. $\square$

Rebound also changes allocation logic. A lower energy cost per verified task shifts the feasible frontier outward. It does not ensure movement toward lower total energy. Quantity responses and system constraints determine the new allocation.

9 National accounting and double-counting control

9.1 Preserve the production core

The national-accounting core remains

$$\text{GDP} = \sum_j \text{GVA}_j + \text{taxes on products} - \text{subsidies on products}.$$

AI services, data centers, electricity generation, equipment manufacture, mining activity, and payment services can already enter gross output, intermediate consumption, and value added. A satellite account may expose their energy use and outcomes, but it must not add those same production flows to GDP a second time.

Let G be an AI-related flow already included in an industry’s GVA. Let G^{memo} display the same source flow in a thematic memorandum. Then

$$\text{GDP} + G^{\text{memo}}$$

double counts G . The memorandum can disaggregate or annotate GDP. It is not an additional production flow.

9.2 Stocks, flows, and assurance

Let K_t be an asset position protected or transferred using a settlement system during period t . It is a stock or exposure measure, not a production flow. Adding K_t to GDP violates both unit semantics and temporal type. Likewise, a modeled avoided expected loss $K_t\{R_0 - R_1\}$ is not automatically value added. It is a scenario-conditioned benefit estimate. Depending on the application, some associated services may already be purchased and recorded in production.

Proposition 9.1 (Production and memorandum non-addition). *An account entry cannot be added to GDP when it is either (i) the same source flow already recorded in GVA, (ii) a stock or exposure, (iii) a probability or physical service, or (iv) a causal benefit estimate not defined as an SNA production flow.*

Proof. Case (i) repeats a source flow. Cases (ii) and (iii) fail the typed-addition condition in theorem 3.5. Case (iv) has a different production boundary and meaning unless an SNA-consistent transaction or imputation is separately established. A memorandum display does not establish that mapping. $\square$

Proposition 9.2 (National non-addition). *Conversion-linked energy stages, production already recorded in GVA, nested produced-adopted-realized AI outcomes, and stock exposures cannot be added as independent production. Typed memorandum panels can be attached without changing GDP.*

Proof. Theorem 3.7 covers linked physical stages, while proposition 9.1 covers repeated production and stock-flow incompatibility. Produced, adopted, and realized quantities are nested stages of one outcome lineage, so their sum repeats earlier stages. Memorandum attachment changes display, not the production identity in section 9.1. $\square$

9.3 The audit

Every additive account entry should carry:

- a unique entry identifier;
- a source-flow identifier;
- a typed quantity;
- a role, either integrated production, satellite memorandum, or physical memorandum; and
- a classification and bridge to the source account.

The deterministic audit emits a repeated-source finding whenever a source-flow identifier appears more than once. It emits a stronger GVA-satellite overlap finding when the same source appears in integrated production and a satellite memorandum. It also rejects an integrated total that mixes semantic, unit, stock-flow, boundary, or horizon types. This is a useful minimum control, not a complete national-accounts compiler. Source systems must still resolve revaluations, transfers, consolidation, residence, imports, taxes, subsidies, and classification changes.

9.4 Twenty-case cross-program screen

The lineage and type rules apply to the following audit cases. A conforming implementation either rejects the aggregation or supplies the named bridge.

1. Do not add primary, conversion-output, final, facility, and lifecycle energy as independent inputs.
2. Do not add manufacturing operation to embodied allocation and then add both again to the economy-wide operational total.
3. Do not add component estimates for host, network, facility, idle, retry, or rework when a covering meter already contains them.
4. Do not charge one shared inference's full energy to several adopted uses and then sum the branches.
5. Do not add tokens, verified tasks, adopted tasks, and realized outcomes as independent outputs. They are typed stages or nested quantities.
6. Do not multiply quality factors that encode the same failure event.
7. Do not remove failed attempts from energy and then separately claim a reliability adjustment.
8. Do not add AI-enabled output to GVA when the source production flow is already recorded.
9. Do not add protected exposure, expected loss, avoided loss, settlement service, and insurance or risk-management production without an overlap bridge.
10. Do not sum overlapping protected exposures or repeated positions across assurance windows.
11. Do not add load payments, generator credits, congestion rent, reserve credits, and reserve charges without closing cash settlement and identifying transfers.
12. Do not add explicit congestion, loss, or energy components when the locational price already embeds them.
13. Do not charge reserve once as headroom and again as the same reserve requirement.
14. Do not sell one curtailment, battery, backup generator, or interconnection capability into overlapping reserve products without coupled constraints.
15. Do not count a demand-response baseline payment as grid surplus without verified reduction and a production-cost counterfactual.
16. Do not count direct AI energy again inside indirect digital or economy-wide energy.
17. Do not call displaced or shifted energy eliminated when it reappears at another site or time.
18. Do not add private attack cost, electricity-producer revenue, and social resource cost without an incidence account.
19. Do not sum ratios across units. Aggregate matched numerators and denominators, then form the declared ratio.
20. Do not add a normalized composite index to its underlying physical or monetary components.

Proposition 9.3 (Lineage-complete aggregation). *Suppose every physical flow, monetary production item, shared service input, and protected exposure has a stable lineage identifier; an aggregation admits each identifier at most once; and distinct identifiers have been verified*

non-overlapping. Then the aggregation contains no duplicate component.

Proof. Assume a duplicate component remains. It must either have the same lineage identifier twice, contradicting the admission rule, or have two identifiers, contradicting the verified non-overlap premise. The conclusion is conditional on the source records satisfying those premises. $\square$

10 Cross-paper equation and notation map

The following crosswalk is normative for this synthesis. It translates each prior paper’s local equation into the canonical namespace, records its estimand family, and names the dependency that prevents overreach. Local paper notation appears only in the conflict table that follows.

Paper	Equation	Role	Dependencies and limits
1	$VPJ_{\mathcal{M}} = \mathbb{E}_{\Pi^{\text{unc}}} [\Delta W] / E_b^e$	Measurement contract and level-intensity family.	Prior dependency: none. $\mathcal{M}$ fixes boundary, counterfactual, value functional, horizon, energy convention, attribution, and uncertainty. It does not create universal weights.
2	$X_n^{\text{int}} / X_0^{\text{int}} = \prod_k X_k^{\text{out}} / X_k^{\text{in}}$	Typed pathwise level-ratio composition.	Prior dependency: P1. Adjacent amounts and semantic types must match. Pathwise composition does not imply a product of expectations or a semantic valuation.
3	$QAIJ_{\mathcal{P}^{\text{eval}}} = \mathbb{E}[\sum_i \tilde{d}_i^{\text{task}} u_i^{\text{svc}} q_i^{\text{AI}} c_i^{\text{correct}} \alpha_i^{\text{svc}} \ell_i^{\text{lat}}] / \mathbb{E}[E_b^e]$	Quality-adjusted task service, ratio-of-expectations family.	Prior dependencies: P1, P2. Task distribution, scoring, reliability, latency, and energy boundary are declared. Task service is not tokens and not realized economic value.
4	$ALPJ = \text{Loss}_P(P r_0^{\text{reorg}} - P r_1^{\text{reorg}}) / \Delta E_P$	Conditional avoided loss, causal ratio-of-effects family when identified.	Prior dependencies: P1, P2. Intervention, threat scenario, loss severity, risk models, and boundary must match. It is not a general measure of trust.
5	$\theta^{\text{causal}} = \Delta_V^{\text{ATE}} / \Delta_E^{\text{ATE},b}$	Causal realized-value ratio-of-effects family.	Prior dependencies: P1, P2, P3. Numerator and denominator identify the same intervention contrast. Near-zero $\Delta_E^{\text{ATE},b}$ requires ratio-aware inference. Tokens remain intermediate.
6	$\frac{\partial B_i^{\text{work}}}{\partial e_{i,n,t}} = k_{n,t}^{\text{grid}} + \lambda_{i,n,t}^{\text{KKT}} + \frac{\partial \mathcal{R}_i^{\text{risk}}}{\partial e_{i,n,t}}$	Marginal-derivative family at node and time.	Prior dependencies: P1, P5. Requires a declared scalar objective, convexity, differentiability, and regularity. It supplies no context-free sector order.
7	$\frac{d \ln E_{\text{AI}}^{\text{facility}}}{d \ln \eta_Q^{\text{svc}/J}} = \epsilon_{Q^{\text{cog}}, \eta_Q} - 1$	Direct rebound elasticity identity, not a level intensity.	Prior dependencies: P3, P5. The demand elasticity is empirical and boundary dependent. Direct, indirect, and economy-wide rebound are distinct.

Paper	Equation	Role	Dependencies and limits
8	$Payment_D - Credit_G = K_{network}$	Market-settlement identity, not an energy-normalized ratio.	Prior dependencies: P4, P5, P6, P7. Dispatch and settlement need explicit network, reserve, and revenue-adequacy rules. A DC solution does not prove AC or dynamic feasibility.
9	$GDP = \sum_j GVA_j + Tax_p - Subsidy_p$	Production identity; attached intensities are descriptive-national family.	Prior dependencies: P1 through P6. Energy, task service, assurance, avoided loss, and protected stocks remain memoranda unless an SNA-consistent flow is identified.

10.1 Imported conflicts and explicit repairs

The independent cross-program audit identified ten actual or status-granularity conflicts. The synthesis disposition is explicit so that no local notation is silently inherited.

ID	Conflict	Paper 10 disposition
C1	O, T, U denote different inference nodes across the corpus.	Use section 3.5: O_{emit} is immediate output, U_P is verified outcome, and T_{κ}^{tok} is an attached representation statistic.
C2	Exergy is X in the registry and B in Paper 2.	Use X^{ex} exclusively for exergy in the synthesis.
C3	Paper 4 declares energy interval τ and attack window T , then writes energy over T .	Use $E_H(\tau_E^{hash})$ and evaluate risk over τ_A^{attack} . Equality is an explicit scenario restriction.
C4	Paper 5's local E^{wall} sum can include facility overhead.	Use the separate system-wall and facility-operational totals in section 5.2, or $E^{operational,b}$ with boundary b .
C5	Paper 8 uses v_w for downward ramp and audited true value.	Use $r_j^{\downarrow}$ and ν_j^{true} , respectively.
C6	Paper 8 labels reserve offers once per MWh and once per MW.	Use currency per MW-hour, convert MW reserve through δ_t^{hrs} , and report the resulting MWh-reserved.
C7	Paper 2's duplicate-branch prose says 20/200 and 20/100 are equal.	They are 0.1 and 0.2, differing by a factor of two. The example is retained only as a warning against duplicating a shared input.
C8	The shared register gives one computational label to Paper 3 results that include formal existence theorems.	Tokenizer and task-mix existence results are PROVED; fixture outputs are COMPUTATIONAL; empirical magnitude and transport remain CONDITIONAL or OPEN.
C9	$\lambda_{n,t}$ is a KKT shadow price in Paper 6 and an externally declared complete price vector in Paper 8.	Use λ_k^{KKT} for multipliers and $p_{n,t}^{nodal}$ for declared market prices. Settlement closure does not turn the latter into a supporting equilibrium price.
C10	Paper 4's loss-on-success L_P is called gross protected exposure in a Paper 9 fixture.	Use gross exposure X_P^{gross} and loss-on-success $Loss_P = X_P^{gross} s_P^{loss}$. Equality requires $s_P^{loss} = 1$.

10.2 Symbol discipline

Several letters recur across economics and engineering. The crosswalk avoids a false unification by qualifying them.

- E_b^e is an energy quantity under boundary b and convention e . It is not expected value, for which the paper uses $\mathbb{E}$.

- U^{work} is useful physical work. $W(\cdot)$ is a declared economic or welfare functional. They are not interchangeable.
- S_{task} is verified task service. S_d is a decision-maker-specific scalarization. Neither denotes GDP.
- $\mathcal{A}(\mathcal{P}^{\text{settle}}, \omega)$ is an assurance vector. A^{adopt} is adopted use. Neither is an empirical treatment assignment, for which the namespace uses Z^{assign} .
- $Pr_0^{\text{reorg}}, Pr_1^{\text{reorg}}$ in section 4.3 are matched scenario risks. $\mathcal{R}_i^{\text{risk}}(x)$ in section 7.1 is a general risk or external-cost functional.
- λ_k^{KKT} is a constraint multiplier and $p_{n,t}^{\text{nodal}}$ a declared nodal price. Neither is an intrinsic value of energy at all locations or dates.

Local symbol	Collision class	Canonical synthesis namespace
O	Immediate output, verified outcome, converted-energy output.	$O_{\text{emit}}, U_{\text{verify}}, O_{\text{energy}}$.
U	Verified outcome, usefulness, ramp-up, product use.	$U_{\mathcal{P}}, u_i^{\text{svc}}, r_i^{\uparrow}, U_{\text{se}ra}$.
T	Tokens, attack window, time-index set.	$T_{\kappa}^{\text{tok}}, \tau_A^{\text{attack}}, \mathcal{T}$.
C	Computation, attack cost, generator credit, conversion input.	$C^{\text{comp}}, \text{Cost}_A, \text{Credit}_G, E_{\text{conv}, \text{in}}$.
A	Adoption, availability, incidence matrix.	$A^{\text{adopt}}, \alpha^{\text{avail}}, \mathbf{B}^{\text{inc}}$.
E	Inventory, random denominator, treatment effect, dispatch quantity, facility amount, national flow.	Always attach stage, boundary, interval, and estimand family.
x	Outcome vector, interface amount, MWh allocation, MW schedule.	$\mathbf{o}, X_k, e_{i,n,t}, p_{j,t}^{\text{load}}$.
η	Exergy efficiency, hashing intensity, tokens/J, service/J.	$\eta_{\text{ex}}, \epsilon_H^{\text{J/TH}}, \eta_{\kappa}^{\text{tok/J}}, \eta_Q^{\text{svc/J}}$.
q	AI quality, adversarial share, workload energy, fallback power.	$q_i^{\text{AI}}, \alpha_A^{\text{hash}}, e_{i,t}^{\text{load}}, p_j^{\text{fallback}}$.
z	Outcome vector, success threshold, confirmations, assignment, restart, normalized indicator.	$\mathbf{j}, h_i^{\text{pass}}, d^{\text{conf}}, Z_i^{\text{assign}}, y_{j,t}^{\text{start}}, \tilde{x}_k$.
c	Counterfactual, correctness, nodal cost, marginal cost, generation cost.	$c_0, c_i^{\text{correct}}, k_{n,t}^{\text{grid}}, mc^{\text{svc}}, mc_g^{\text{gen}}$.
a	Attribution, service reliability, grid availability.	$a^{\text{attr}}, \alpha^{\text{svc}}, \alpha^{\text{grid}}$.
π	Uncertainty, price pass-through, demand-response payment.	$\Pi^{\text{unc}}, \epsilon_{p,c}^{\text{pass}}, \varpi^{\text{DR}}$.
τ, T, Δ	Horizon, energy interval, attack window, dispatch duration, and treatment contrast.	$\tau_H, \tau_E, \tau_A, \delta_t^{\text{hrs}}$, with Δ reserved for declared contrasts.
λ	Value scaling, Poisson mean, KKT multiplier, nodal price.	$\alpha_W, \mu_P, \lambda_k^{\text{KKT}}, p_{n,t}^{\text{nodal}}$.
ρ	Shared allocation, risk weight, pass-through, reserve price.	$a^{\text{shared}}, \gamma_i^{\text{risk}}, \epsilon^{\text{pass}}, p_t^{\text{reserve}}$.
R	Reorganization risk, reward, risk functional, rebound, reserve.	$Pr^{\text{reorg}}, \text{Reward}_A, \mathcal{R}_i^{\text{risk}}, \mathcal{B}^{\text{rebound}}, \text{Req}_t^{\text{reserve}}$.
L	Graph edges, latency, protected loss, conversion loss.	$\mathcal{E}_G, \ell_i^{\text{lat}}, \text{Loss}_P, E_{\text{conv}, \text{loss}}$.
P	Evaluation protocol, payment, primary supply.	$\mathcal{P}^{\text{eval}}, \text{Payment}_D, E_{\text{primary}, \text{supply}}$.
Q	Cognitive service, required energy, outcome counts.	$Q^{\text{cog}}, E_{i,k}^{\text{req}}, Q_k^{\text{outcome}}$.
N	Chain network, net value, active units, grid nodes.	$\mathcal{N}^{\text{chain}}, V_i^{\text{net}}, n^{\text{active}}, \mathcal{N}^{\text{grid}}$.
w	Valuation weights, task weights, workload index, composite weights.	$\omega^{\text{value}}, \omega_i^{\text{task}}, j \in \mathcal{J}^{\text{load}}, \omega_k^{\text{index}}$.
κ	Tokenizer and benefit curvature.	$\kappa^{\text{tok}}, \gamma^{\text{curve}}$.
B	Exergy, attacker benefit, workload benefit, susceptance.	$X^{\text{ex}}, \text{Benefit}_A, B_i^{\text{work}}, b_{\ell}^{\text{sus}}$.
X	Exergy, outcome space, feasible set, other energy disposition.	$X^{\text{ex}}, \mathcal{O}, \mathcal{F}, D^{\text{other}}$.

Local symbol	Collision class	Canonical synthesis namespace
v_w	Downward ramp and audited true value in Paper 8.	$r_j^\downarrow, \nu_j^{\text{true}}$.
θ	Semantic type, causal ratio, voltage angle.	$\text{type}(x), \theta^{\text{causal}}, \phi_{n,t}$.
d	Unit dimension, task difficulty, deadline or damage, fixed demand, line destination.	$\text{dim}, d_i^{\text{task}}, h_{i,k}^{\text{deadline}}, \text{damage}_{n,t}, D_{n,t}^{\text{fixed}}, \text{head}(\ell)$.

The software function `canonical_equation_crosswalk()` fixes Papers 1 through 9 in order, assigns each an evidence label, and validates uniqueness and completeness. That executable list is a regression guard against losing a paper or silently changing its role.

11 Boundary map

Three additional boundaries cross-cut table 4. The *outcome boundary* identifies the service, population, distribution, and external effects counted in the numerator. The *causal boundary* fixes the intervention and counterfactual. The *institutional boundary* fixes ownership, residence, legal obligations, protocol governance, and market settlement. A complete comparison aligns all four.

11.1 Boundary-reversal counterexample

Let systems a and b produce 100 and 120 matched service units. Their device energy is 10 J and 20 J, so the device-boundary intensities are 10 and 6 service units per J and a ranks first. Suppose the matched facility boundary adds 40 J of non-overlapping cooling and shared infrastructure to a and none to b . Facility intensities are then 2 and 6, and b ranks first. Neither ranking is an error. They answer different boundary questions. Comparing one device ratio directly with the other facility ratio would fail the manifest gate.

12 Incompatibility and non-aggregation

12.1 Why avoided loss does not solve every incompatibility

Analysts sometimes monetize every outcome as avoided loss or willingness to pay. That can support a decision model when exposure, counterfactual risk, standing, prices, distribution, and uncertainty are defensible. It does not produce a natural scalar. A service may protect rights or institutional features that are deliberately not traded off at a market price. The same loss event may already enter insurance payments, business interruption, asset revaluation, and welfare estimates. Monetization therefore creates additional overlap and identification duties.

Boundary	Included energy	Suitable question	Frequent invalid inference
Device operational	Metered accelerator or device energy during named work.	Hardware or kernel efficiency under controlled load.	Calling the result facility, marginal-grid, or lifecycle energy.
System operational	Device plus declared host, memory, storage, and network components.	End-to-end system comparison under matched service.	Ignoring idle, retries, failed jobs, or shared hosts.
Facility operational	IT load plus declared cooling and power conversion overhead.	Deployment operations at a named site and time.	Treating a fixed facility factor as universal across load and weather.
Embodied increment	Uniquely allocated manufacturing and construction increment.	Decision about additional equipment or infrastructure.	Adding a full asset lifecycle to every short-run job.
Lifecycle	Declared upstream, operational, replacement, and end-of-life processes.	Technology comparison for a functional unit over a stated life.	Combining a lifecycle total with operational components already inside it.
Primary	Energy resources entering the economy or transformation chain.	Resource supply and economy-wide energy dependence.	Adding conversion output or final use from the same lineage.
Final	Energy products delivered to final users after transformation.	End-use sector demand and intensity.	Interpreting final electricity as primary-resource requirement without a bridge.
Marginal grid	Modeled change in generation, losses, congestion, and emissions caused by an incremental load.	Dispatch, location, timing, and flexible-load decisions.	Equating average facility electricity with the marginal system response.

Table 4: Energy boundaries answer different questions. A bridge can relate them, but a shared joule unit does not make them identical.

Quantity	Native kind	May be normalized by J?	May be added directly to
Useful mechanical work	Physical service, J or MJ	Yes, with matched input boundary.	Only the same physical service under matched type and scope.
Verified task service	Task-equivalent	Yes, for a named task distribution.	Only matched task service. Not tokens, work, money, or trust.
Causal realized value	Money flow or welfare index	Yes, with matched causal energy effect.	Matched value flows under the same functional, price basis, boundary, and horizon.
Settlement risk	Probability or risk measure	It can be displayed beside energy; a ratio needs an interpreted purpose.	Only matched probabilities or risk measures. Not avoided dollars.
Avoided expected loss	Scenario-conditioned money flow	Yes, conditionally, as ALPJ.	Matched expected-loss effects, subject to overlap with recorded services.
Gross protected position	Money stock or exposure	Per-joule display is possible but not a production ratio.	Matched stocks at the same date. Never GDP.
Gross value added	Money flow under SNA production boundary	Yes, as a production-energy indicator.	Matched GVA flows. Not satellite outcomes already represented in GVA.
Composite index	Dimensionless declared view	Only if the index construction specifies the resource role.	Only the same index definition. A different weight vector is a different view.

Table 5: The common denominator does not erase numerator types.

12.2 Why exergy does not solve every incompatibility

Exergy offers a common physical metric for the maximum useful work obtainable relative to an environment. It can improve the description of the energy-to-work portion of a transduction chain. It cannot determine the correctness of an answer, the social usefulness of a task, a protocol’s adversarial robustness, or a counterfactual welfare effect. The synthesis uses exergy as a typed physical node rather than a universal economic numerator.

12.3 Why price does not solve every incompatibility

Market prices encode marginal exchange under particular institutions, property rights, distributions, and market structures. They can value traded flows in a specified accounting context. They may omit externalities, nonmarket production, consumer surplus, distributional standing, tail risks, and rights constraints. A price-weighted scalar is legitimate when its purpose and limitations are stated. Calling it the Joule Standard without qualification would conceal the valuation choice.

13 A synthetic integrated fixture

Consider three feasible systems evaluated under one manifest. The outcome axes are useful work in MJ, verified task service in task-equivalents, and avoided expected loss in 2026 USD.

System	Energy J	Useful work MJ	Task service	Avoided loss USD
Balanced	80	70	70	70
Cognitive	90	30	95	45
Dominated	100	60	60	60

Balanced dominates Dominated because it uses less energy and is better on every outcome. Balanced and Cognitive do not dominate one another. Cognitive produces more task service; Balanced uses less energy and produces more useful work and avoided loss. The Pareto frontier contains both.

After published min-max normalization, a view with weights

$$(\omega_E^{\text{index}}, \omega_W^{\text{index}}, \omega_Q^{\text{index}}, \omega_A^{\text{index}}) = (0.15, 0.15, 0.65, 0.05)$$

can choose Cognitive, while

$$(\omega_E^{\text{index}}, \omega_W^{\text{index}}, \omega_Q^{\text{index}}, \omega_A^{\text{index}}) = (0.30, 0.40, 0.10, 0.20)$$

can choose Balanced. The ranking difference is not a numerical failure. It reports different priorities. The unweighted vector and frontier remain the primary result.

A separate transduction fixture begins with 100 J of facility energy, produces 80 named operations, 40 verified task units, and 20 USD of incremental value under a matched design. The typed yields are

$$\frac{80 \text{ op}}{100 \text{ J}}, \quad \frac{40 \text{ task}}{80 \text{ op}}, \quad \frac{20 \text{ USD}}{40 \text{ task}}.$$

Their exact product is 0.2 USD/J, equal to the terminal ratio. The code rejects a substitution of “80 tokens” for “80 verified tasks” even when both use the display unit “count”. Semantic type is part of the interface.

For an assurance fixture, let exposure be 1000 USD, baseline risk 0.01, intervention risk 0.002, and incremental energy 50 J. Avoided expected loss is 8 USD and $ALPJ = 0.16$ USD/J. These computations do not establish that the risk probabilities are empirically correct. They test the conditional accounting once the scenario inputs are supplied.

13.1 Independent national closures

A small national fixture closes five ledgers independently.

1. **Product energy.** Primary supply is 1000 J. Its disposition is 200 J of direct final primary use, 100 J of primary exports, 400 J of conversion loss, 250 J of final converted energy, and 50 J of converted-energy exports. These disjoint dispositions sum to 1000 J. The 300 J conversion output is displayed as a stage, not added to the primary total.
2. **Component lineage.** A facility meter carries lineage `op-2026-q3` and reports 100 J. A server-manufacture allocation carries lineage `server-manufacture` and reports a 20 J embodied increment. Their scoped 120 J total is permitted because the lineages do not overlap. The fixture does not also add device and cooling components already inside the facility meter.
3. **Production.** GVA entries of 400 and 300 currency units, product taxes of 40, and subsidies of 10 give GDP of 730.
4. **Outcome chain.** The memorandum reports 100 produced, 70 adopted, and 50 realized task units. The quantities reconcile as nested stages and are not summed to 220.
5. **Settlement cash.** Load payments of 1000, generator credits of 850, and network or constraint rent $K = 150$ satisfy $Payment_D - Credit_G = K$. These transfers are not added as three welfare benefits.

Passing one closure does not repair another. A balanced cash settlement can coexist with a duplicated energy lineage, and a balanced physical account can coexist with a repeated production flow.

14 Reference implementation

The reference implementation is `src/joule_standard/synthesis.py`. It uses only the Python standard library and local evidence-status enumeration. Public numeric inputs use `decimal.Decimal`, integers, or decimal strings. Binary floating-point inputs are rejected so a displayed exact fixture cannot silently depend on a binary approximation.

14.1 Executable objects

The module provides:

- **QuantityType** and **TypedQuantity**, with exact semantic, unit, accounting-kind, boundary, and horizon checks;
- **BoundaryManifest**, with a field-by-field direct-comparability gate;
- **TransductionStage**, typed exact path composition, retained stage statuses, and the conservative evidence join in section 4.2;
- **EnergyInventory**, physical-lineage identifiers, energy-stage controls, and separate operational plus embodied increments;
- source-tagged account entries and a minimum double-counting audit;
- **SystemVector**, direction-aware dominance, and deterministic Pareto filtering;
- published normalization rules and explicit scalarization weights;
- matched causal effect ratios and assurance contrasts;
- direct and finite rebound identities; and
- a canonical equation crosswalk for Papers 1 through 9.

14.2 Refusal behavior

Refusal is part of the implementation. The module raises a typed error when an operation would:

- add a money stock to a money flow;
- add tokens to verified task service;
- compose stages whose interface semantics or amounts differ;
- compare manifests with unmatched energy stages or other required fields;
- repeat a component identifier or physical lineage;
- total primary and final energy as independent stages;
- compare Pareto vectors with different axes or units;
- accept scalarization weights that do not sum exactly to one;
- form a causal ratio from unmatched interventions or a zero energy effect; or
- accept a probability outside $[0, 1]$.

These rules do not decide contested empirical questions. They prevent invalid operations after a researcher has supplied the relevant records.

14.3 Adversarial test contract

The dedicated test file contains twenty-four tests. It verifies a valid energy-to-computation-to-service-to-value path, then substitutes an identically displayed but semantically different count at an interface. It attacks stock-flow addition, task-token addition, boundary matching, duplicate energy components, primary-final mixing, repeated account sources, GVA-memorandum overlap, outcome-axis substitution, invalid weight sums, zero-denominator causal ratios, invalid assurance risks, crosswalk omissions, crosswalk misordering, and binary floats. It also checks positive cases: separate operational and embodied increments, Pareto removal of a dominated point, conditional rank reversal, exact avoided loss, and the rebound identities.

The test suite is evidence for program behavior, not evidence that any real-world parameter has been measured correctly.

15 Identification and falsification agenda

A useful standard must say how its claims could fail. The following agenda separates mathematical, software, empirical, and institutional tests.

15.1 Boundary falsification

A comparison fails its direct-comparability claim if a line-item audit reveals different functional units, time intervals, locations, energy stages, failure treatment, embodied scope, counterfactuals, outcome functionals, attribution rules, uncertainty models, or exclusions. Researchers should publish manifests before examining rankings. A blinded boundary audit can then test whether the claimed comparison survives.

Facility energy should be reconciled against device telemetry, power distribution data, and site meters. Discrepancies need an allocation bridge, not a generic overhead factor. Lifecycle claims should publish inventory processes and show that operational energy is not added after already entering the lifecycle total.

15.2 Transduction falsification

Each edge needs an observable or model that could contradict it. Energy-to-work edges can be checked by calibrated physical measurement. Computation-to-service edges can be tested on held-out tasks with blinded scoring and failure retention. Service-to-adoption edges can be tested using workflow logs. Adoption-to-value edges require randomized assignment, a credible quasi-experiment, or explicit partial-identification bounds.

A transduction claim fails if the output population differs from the next stage’s input population, if attrition is silently removed, or if a semantic substitution occurs. A product of average stage yields should be compared with the average pathwise product. Their divergence is evidence against an independence simplification.

15.3 Quality-adjusted intelligence

Quality scores should be challenged by:

- tokenizer substitutions that preserve user-visible content;
- verbosity perturbations that add tokens without adding useful content;
- adversarial task-mix shifts;
- calibration and abstention tests;
- repeated trials that expose reliability and tail failures;
- latency thresholds tied to task usefulness; and
- human or system baselines evaluated under the same functional unit.

A claimed QAIJ advantage is falsified for a target population when it disappears under the prespecified task distribution or fails the boundary gate. A robust result should report a frontier over quality, latency, reliability, and energy rather than one chosen score alone.

15.4 Settlement assurance

An assurance claim should vary adversary resources, network delay, hardware availability, concentration, bribery, reward changes, censorship objectives, confirmation horizon, and asset exposure. Empirical event studies can examine reorganizations, fee shocks, hashrate migration, outages, and concentration, but observed survival alone does not identify counterfactual loss prevention.

The claim “more energy caused more assurance” is falsified if a matched intervention changes energy without improving the prespecified risk measure, or if the improvement is explained by another protocol or market change. Equal energy systems should be used as negative controls against any energy-only security mapping.

15.5 Causal realization

For AI and automation, prespecify numerator and denominator effects. Report first-stage energy effects, outcome effects, compliance, spillovers, interference, and missingness. If the energy effect is near zero, report a ratio-aware confidence set rather than a conventional bounded interval. Compare results with and without low-value, failed, repeated, and abandoned inferences. A value-per-joule estimate that improves only by deleting failures does not represent the deployment.

15.6 Marginal allocation and rebound

Allocation models should be back-tested against realized nodal prices, congestion, curtailment, deadlines, quality, and service delivery. Shadow-price interpretations fail when constraints or nonconvexities omitted from the model bind in operation. Flexible-load claims require measurement of response speed, recovery costs, and foregone service.

Rebound estimation needs a credible source of efficiency variation and a stable service definition. Researchers should report direct service demand, prices, complementary inputs, displaced work, and economy-wide effects over multiple horizons. A decline in device joules per task does not by itself test $\epsilon_{Q^{\text{cog}}, \eta_Q}$.

15.7 National-account audit

Statistical compilers should trace every monetary entry to a source transaction or imputation and every physical entry to a lineage. The following adversarial questions should be answered:

1. Does an AI outcome memorandum repeat sales or output already included in GVA?
2. Is a protected asset position being added to a production flow?
3. Does avoided loss overlap an insurance payment, revaluation, or service fee already recorded?
4. Are primary energy, conversion output, and final use summed?
5. Is embodied energy already contained in a lifecycle total?
6. Do imports, residence, and time-of-recording rules match the monetary core?

7. Does a composite ranking survive plausible normalizations and weights?

16 Claim ledger

Claim	Scope	Status
JS-1	A measurement tuple and expanded boundary manifest are sufficient to state a reproducible conditional value-per-joule estimand. Sufficiency here concerns declaration, not empirical identification.	PROVED
JS-2	Physical efficiency, exergy efficiency, useful work, task service, causal economic realization, and assurance are distinct typed objects.	PROVED
JS-3	Exact typed stage yields telescope when every adjacent amount and type matches. No product-of-expectations claim is included.	PROVED
JS-4	Energy normalization preserves numerator incompatibilities and does not define addition across work, task service, probabilities, stocks, or value flows.	PROVED
JS-5	Non-dominating vectors can reverse order under admissible positive weights. Therefore the framework has no natural common scalar absent declared weights or a value functional.	PROVED
JS-6	Under strict monotonicity and strictly positive weights, a weighted-sum maximizer is Pareto efficient. Unsupported frontier points may remain hidden without convexity.	PROVED
JS-7	Primary, conversion-output, and final-energy appearances from one lineage cannot be added as independent energy inputs.	PROVED
JS-8	AI output already recorded in GVA and protected asset stocks cannot be added to GDP as new production.	PROVED
JS-9	QAIJ is a quality-adjusted task-service measure for a named task distribution, not a tokenizer-invariant measure of all intelligence.	CONDITIONAL
JS-10	ALPJ can express modeled avoided loss per incremental joule under matched risk scenarios. Its empirical value depends on the threat and loss models.	CONDITIONAL

Claim	Scope	Status
JS-11	A ratio of causal value and energy effects requires a matched intervention, counterfactual, boundary, and horizon.	PROVED
JS-12	The direct rebound identity classifies energy change by the demand elasticity. The magnitude of that elasticity in real AI markets is not established here.	PROVED identity; OPEN magnitude
JS-13	Flexible AI and proof-of-work loads can participate in a modeled nodal market when their heterogeneous constraints and settlement rules are represented.	CONDITIONAL
JS-14	The standard-library implementation enforces the documented generic type, boundary, lineage, Pareto, causal-ratio, and crosswalk contracts on its test fixtures.	COMPUTATIONAL

17 Governance and reporting protocol

The Joule Standard can support engineering studies, market analyses, corporate reporting, and public statistics only if revisions remain inspectable. A conforming report should publish the following artifacts.

1. **Manifest.** Every required boundary field in section 3.1, with a persistent identifier.
2. **Energy ledger.** Component identifiers, physical lineages, stages, joules, metering or model source, and uncertainty.
3. **Outcome dictionary.** Semantic types, units, orientations, populations, horizons, and evidence labels.
4. **Transduction graph.** Stage inputs and outputs, interface mappings, missing edges, and evidence status.
5. **Vector results.** Native outcome values, energy, normalized displays if used, and the Pareto frontier.
6. **Scalar views.** Named decision purpose, normalization bounds, weights, aggregation rule, tie rule, and sensitivity analysis.
7. **Double-count audit.** Source-flow and physical-lineage findings, bridges, and unresolved overlaps.
8. **Identification record.** Intervention, counterfactual, estimator, uncertainty, missingness, spillovers, and external-validity limits.
9. **Software record.** Versioned code, tests, exact inputs, and deterministic output where feasible.

Versioning should distinguish corrections to observations from changes in definitions. A changed weight vector creates a new scalar view. A changed energy stage creates a new denominator. A changed task distribution creates a new service estimand. Historical results should not be silently overwritten.

17.1 Minimum publication table

A human-readable summary can remain compact:

Field	Required statement
Question	The decision or descriptive question and functional unit.
Energy	Stage, physical boundary, time, location, allocation, embodied scope, and uncertainty.
Outcome	Semantic type, unit, population, direction, and evidence status.
Counterfactual	The named alternative and identification design.
Result	Native outcome, energy, conditional ratio if meaningful, and uncertainty.
Comparison	Manifest match, bridges, Pareto relation, and sensitivity to declared views.
Nonclaims	The downstream meanings not identified by the study.

18 Limitations

Limitation 18.1 (Type systems cannot validate semantics). The software can detect unequal labels, units, kinds, boundaries, and horizons. It cannot prove that a researcher gave a label its correct empirical meaning. Ontology governance, documentation, audit, and domain review remain necessary.

Limitation 18.2 (No automatic welfare function). The framework can carry a welfare functional and test its accounting interfaces. It cannot infer society's welfare weights from energy use, prices, or statistical variance. Distributional and rights-based constraints may be more appropriate than a compensatory scalar.

Limitation 18.3 (Partial equilibrium). Several paper-level models are partial equilibrium. Prices, hardware supply, labor, institutions, and energy infrastructure can respond endogenously. Rebound and national-account effects require wider system models for economy-wide claims.

Limitation 18.4 (Tail risk and deep uncertainty). Expected loss can conceal fat tails, ambiguity, correlated failures, and irreversible outcomes. Scenario vectors, robust constraints, and stress tests may be preferable to an expected-value scalar.

Limitation 18.5 (Grid physics). The nodal direct-current formulation is useful for economic dispatch and congestion analysis. It omits reactive power, voltage, frequency dynamics, contingency detail, and device-level response. Production deployment requires the relevant engineering studies.

Limitation 18.6 (Benchmark validity). A quality-adjusted task measure is only as valid as its task distribution, scoring, human judgments, and operational match. Distribution shift can change both numerator quality and denominator energy.

Limitation 18.7 (National implementation). The source-tag and type audit is a minimum control. A statistical office needs the full SNA and SEEA compilation process, confidentiality protection, classification bridges, balancing, revisions, and institutional governance.

19 Conclusion

The research program began with a tempting quotient and ended with a boundary discipline. Joules are a common physical accounting unit. They allow systems to be metered, transformation lineages to be reconciled, marginal scarcity to be priced, and outcome ratios to be stated. They do not turn unlike consequences into one substance.

The Joule Standard therefore has a vector core. Useful work remains physical service. Quality-adjusted intelligence remains verified task service on a declared distribution. Realized economic value remains a counterfactual effect under a value functional. Proof-of-work assurance remains a scenario-conditioned risk and mechanism record. GDP remains a production flow, supplemented by physical and outcome memoranda. Every scalar is a conditional view whose normalization, weights, and evidence are open to inspection.

This architecture is less dramatic than a universal energy theory of economic value. It is also more usable. It tells an engineer why device energy is not facility energy, an AI evaluator why tokens are not service, a security analyst why hashrate is not assurance, an economist why adoption is not causal value, a grid operator why flexible loads need nodal constraints, and a statistician why memorandum outcomes are not additions to GDP. Most importantly, it makes disagreement legible. Analysts can dispute a boundary, threat model, task distribution, causal design, or weight vector without pretending that the dispute is about arithmetic.

A Formal reporting schema

A machine-readable implementation may represent one study as the tuple

$$\mathcal{D} = (B, \mathcal{E}, \mathcal{X}, \mathcal{T}, \mathcal{A}, \mathcal{I}, \mathcal{U}),$$

where:

B is the boundary manifest in section 3.1;

$\mathcal{E}$ is the energy inventory with component and lineage identifiers;

$\mathcal{X}$ is the typed outcome-axis dictionary;

$\mathcal{T}$ is the transduction graph;

$\mathcal{A}$ is the set of source-tagged accounting entries;

$\mathcal{I}$ is the identification and uncertainty record; and

$\mathcal{U}$ is the set of optional decision views.

Validation occurs in this order:

1. validate required text, exact numeric domains, and identifiers;
 2. reject duplicate energy components, stages, candidates, and account entries;
 3. validate physical lineages and energy-stage totals;
 4. validate typed interfaces and direct comparisons;
 5. validate source-flow overlaps and integrated-account types;
 6. construct native-unit vectors and the Pareto frontier;
 7. validate normalization, orientation, weights, and tie rules for each optional scalar view;
- and

8. publish refused operations alongside successful calculations.

The order matters for explanation, not for creating a scalar hierarchy. A record that fails the physical-lineage audit should be corrected before its ratio is interpreted. A valid physical ledger can still have an unidentified economic edge.

B Worked audit cases

B.1 Token substitution

Suppose stage one outputs 80 token-count units and stage two declares an input of 80 verified-task-count units. Amount and display unit match. The semantic and accounting kinds do not. Composition is refused. An empirical mapping from tokens to verified tasks may be estimated, but it becomes a new stage with its own evidence status.

B.2 Protected exposure added to production

Suppose a settlement system supports transfers involving a 1 billion USD asset position, charges 2 million USD in service fees, and the relevant business records 1 million USD in GVA. The position is a stock or exposure. The fees and GVA are period flows with distinct accounting roles. The account does not report 1.001 billion USD as new production.

B.3 Lifecycle plus operation

Suppose a lifecycle study allocates 120 J to a functional unit, including 100 J of operation and 20 J of embodied energy. A separate site meter also reports the same 100 J. Adding 120 + 100 repeats operation. A legitimate two-panel report may display the lifecycle total and its components, or a current operational account plus a separately identified embodied increment.

B.4 Near-zero causal denominator

Suppose an intervention effect on value is estimated at 10 USD while its energy effect is 0.1 J with standard error 0.2 J. The point ratio is 100 USD/J, but the denominator is not separated from zero. Reporting the point ratio as a stable productivity estimate is misleading. The study should report the joint effect estimates and a ratio-aware confidence set, which may be unbounded or disjoint.

B.5 Weight-sensitive frontier

Suppose one allocation is better on task service and another is better on useful work, assurance, and energy. Neither is universally superior. The report publishes both frontier points, then

shows the regions of the weight simplex in which each maximizes the declared score. A single preferred view may be selected for a decision, but the frontier remains available for audit.

C Reproducibility map

The numeric labels below are stable contract identifiers, not the source-order positions of pytest functions.

Contract ID	Adversarial condition	Contract
1	Valid three-stage typed path.	Direct ratio equals product of yields.
2	Token count substituted for task count.	Interface type mismatch.
3	Money stock plus money flow; tokens plus tasks.	Typed addition refused.
4	Facility and device manifests.	Energy-stage mismatch reported.
5	Duplicate energy identifier; primary plus final stages.	Both totals refused.
6	Unique facility operation plus embodied increment.	Non-overlapping total accepted.
7	One source flow in GVA and a memorandum.	Repeated-source and overlap findings emitted.
8	Work, task, and money outcome axes.	Per-joule units remain distinct.
9	One dominated and two non-dominated systems.	Only dominated system removed.
10	Scalar weights sum above one.	View refused.
11	Two published positive weight vectors.	Conditional choice reverses.
12	Effect estimates with matched and mismatched boundaries or populations.	Only a fully matched causal ratio is accepted.
13	Zero energy effect.	Finite causal ratio refused.
14	Valid and invalid assurance risks.	Exact result, domain check.
15	Differential and finite rebound fixtures.	Identities reproduced.
16	Canonical crosswalk.	Papers 1 through 9 present exactly once.
17	Missing or misordered crosswalk.	Validation refused.
18	Binary floating-point input.	Exact numeric boundary enforced.
19	Canonical optional-node graph and token attachment.	Immediate output, verified outcome, and token statistic remain distinct.
20	Lineage-complete monetary total and source alias.	Unique sources sum; alias is refused.
21	Ratio of expectations versus mean of unit ratios.	Estimand tags and different exact values retained.
22	High average value with low marginal slope.	Average and marginal families remain separate.

Contract ID	Adversarial condition	Contract
23	Multi-stage path with an open realization edge.	Conservative evidence join retains the open terminal status.
24	Manifests differing only in exclusions.	Exclusion mismatch is reported and direct comparison is refused.

D Checklist for an empirical Joule Standard claim

1. State the decision question before selecting the denominator.
2. Name the functional unit and target population.
3. Publish the physical, temporal, spatial, causal, and institutional boundaries.
4. Choose one coherent energy stage or publish a non-overlapping bridge.
5. Retain idle energy, retries, failures, and shared infrastructure according to a declared rule.
6. Define every numerator in native semantic and accounting units.
7. Draw the transduction graph and label unsupported edges OPEN.
8. Match numerator and denominator interventions for a causal ratio.
9. Report joint uncertainty when the denominator is estimated.
10. Keep task service distinct from tokens and realized value.
11. Keep assurance vectors distinct from energy and gross protected stocks.
12. Trace production memoranda to GVA source flows.
13. Remove dominated alternatives before optional scalarization.
14. Publish normalization, weights, tie rules, and sensitivity regions.
15. State precise nonclaims and conditions that would falsify the result.

References

- [1] M. Long, “Value per Joule: Foundations of an Energy-Normalized Economics,” GrokRxiv:2026.07.value-per-joule, 2026.
- [2] M. Long, “From Exergy to Intelligence: A General Theory of Economic Transduction,” GrokRxiv:2026.07.exergy-to-intelligence, 2026.
- [3] M. Long, “Quality-Adjusted Intelligence per Joule,” GrokRxiv:2026.07.quality-adjusted-intelligence-per-joule, 2026.
- [4] M. Long, “Trust per Joule: The Thermoeconomics of Proof-of-Work Settlement,” GrokRxiv:2026.07.trust-per-joule, 2026.
- [5] M. Long, “Tokens Are Not Value: Economic Accounting for AI Inference,” GrokRxiv:2026.07.tokens-are-not-value, 2026.

- [6] M. Long, “Marginal Value of Energy in an Automated Economy,” GrokRxiv:2026.07.marginal-value-of-energy, 2026.
- [7] M. Long, “The Cognitive Rebound Effect: Why Efficient AI May Increase Energy Demand,” GrokRxiv:2026.07.cognitive-rebound-effect, 2026.
- [8] M. Long, “Energy Markets for Machine Intelligence and Cryptographic Settlement,” GrokRxiv:2026.07.energy-markets-for-machine-intelligence-and- cryptographic-settlement, 2026.
- [9] M. Long, “A Value-per-Joule National Accounting System,” GrokRxiv:2026.07.value-per-joule-national-accounting, 2026.
- [10] M. G. Patterson, “What is energy efficiency? Concepts, indicators and methodological issues,” *Energy Policy*, vol. 24, no. 5, pp. 377–390, 1996. [https://doi.org/10.1016/0301-4215\(96\)00017-1](https://doi.org/10.1016/0301-4215(96)00017-1).
- [11] R. U. Ayres and B. Warr, “Accounting for growth: the role of physical work,” *Structural Change and Economic Dynamics*, vol. 16, no. 2, pp. 181–209, 2005. <https://doi.org/10.1016/j.strueco.2003.10.003>.
- [12] B. Warr, R. Ayres, N. Eisenmenger, F. Krausmann, and H. Schandl, “Energy use and economic development: A comparative analysis of useful work supply in Austria, Japan, the United Kingdom and the US during 100 years of economic growth,” *Ecological Economics*, vol. 69, no. 10, pp. 1904–1917, 2010. <https://doi.org/10.1016/j.ecolecon.2010.03.021>.
- [13] C. E. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal*, vol. 27, pp. 379–423 and 623–656, 1948. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>.
- [14] R. Landauer, “Irreversibility and heat generation in the computing process,” *IBM Journal of Research and Development*, vol. 5, no. 3, pp. 183–191, 1961. <https://doi.org/10.1147/rd.53.0183>.
- [15] P. Mattson et al., “MLPerf Training Benchmark,” *Proceedings of Machine Learning and Systems*, vol. 2, pp. 336–349, 2020. https://proceedings.mlsys.org/paper_files/paper/2020/hash/411e39b117e885341f25efb8912945f7-Abstract.html.
- [16] P. Liang et al., “Holistic Evaluation of Language Models,” *Transactions on Machine Learning Research*, 2023. <https://openreview.net/forum?id=i04LZibEqW>.
- [17] S. Nakamoto, “Bitcoin: A Peer-to-Peer Electronic Cash System,” 2008. <https://bitcoin.org/bitcoin.pdf>.
- [18] J. Garay, A. Kiayias, and N. Leonardos, “The Bitcoin Backbone Protocol: Analysis and Applications,” in *Advances in Cryptology, EUROCRYPT 2015*, pp. 281–310. https://doi.org/10.1007/978-3-662-46803-6_10.

- [19] E. Budish, “The Economic Limits of Bitcoin and the Blockchain,” NBER Working Paper 24717, 2018. <https://doi.org/10.3386/w24717>.
- [20] D. B. Rubin, “Estimating causal effects of treatments in randomized and nonrandomized studies,” *Journal of Educational Psychology*, vol. 66, no. 5, pp. 688–701, 1974. <https://doi.org/10.1037/h0037350>.
- [21] P. W. Holland, “Statistics and causal inference,” *Journal of the American Statistical Association*, vol. 81, no. 396, pp. 945–960, 1986. <https://doi.org/10.1080/01621459.1986.10478354>.
- [22] G. W. Imbens and D. B. Rubin, *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge University Press, 2015. <https://doi.org/10.1017/CB09781139025751>.
- [23] United Nations et al., *System of National Accounts 2025*, adopted standard, pre-edit text consulted 23 July 2026. https://unstats.un.org/unsd/nationalaccount/docs/2025_SNA_Pre-edit.pdf.
- [24] United Nations et al., *System of Environmental-Economic Accounting 2012: Central Framework*, 2014. <https://seea.un.org/content/seea-central-framework>.
- [25] United Nations, *System of Environmental-Economic Accounting for Energy*, 2019. <https://seea.un.org/content/seea-energy>.
- [26] OECD and European Commission Joint Research Centre, *Handbook on Constructing Composite Indicators: Methodology and User Guide*, OECD Publishing, 2008. <https://doi.org/10.1787/9789264043466-en>.
- [27] A. Sen, *Collective Choice and Social Welfare*. Holden-Day, 1970.
- [28] M. Fleurbaey, “Beyond GDP: The quest for a measure of social welfare,” *Journal of Economic Literature*, vol. 47, no. 4, pp. 1029–1075, 2009. <https://doi.org/10.1257/jel.47.4.1029>.