

Value per Joule: Foundations of an Energy-Normalized Economics

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

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

23 July 2026

Abstract

Dividing an economic outcome by energy is easy. Deciding what the quotient means is harder. Existing work already reports gross domestic product per unit of energy, studies useful work and exergy conversion, constructs energy efficiency indicators, and joins physical with monetary accounts. This paper does not claim those ideas as new. It gives a formal account of the declarations needed before a value-per-joule statistic can support comparison. The central object is a measurement tuple $\mathcal{M} = (b, c, W, \tau, e, a, \pi)$: system boundary, counterfactual, value functional, time horizon, energy convention, attribution rule, and uncertainty model. Conditional on this tuple, a system s has

$$\text{VPJ}_{\mathcal{M}}(s) = \frac{\mathbb{E}_{\pi}[\Delta W(s \mid c)]}{E_b^e(s; \tau)}.$$

The quotient has a clear unit and a reproducible interpretation only when every component is declared. Several results delimit what the statistic can do. A common positive rescaling of a value unit preserves rankings but changes cardinal magnitudes. Heterogeneous outcome vectors that do not dominate one another admit opposing rankings under admissible positive weights. Enlarging a physical boundary, changing a counterfactual, altering the time horizon, or reallocating a shared burden can also reverse a ranking. These are not defects that better arithmetic removes. They identify the normative and accounting choices on which the result depends. The paper supplies a comparability gate, an uncertainty protocol, worked counterexamples, and a standard-library Python reference implementation. The proposed framework is a discipline for making conditional ratios legible, not an energy theory of value and not a universal ranking of industrial, computational, monetary, or cognitive systems.

Keywords: energy productivity; value measurement; system boundaries; counterfactuals; index numbers; lifecycle accounting; rank reversal.

Evidence status. Definitions and algebraic results are labeled PROVED when their assumptions appear in the statement. Numerical fixtures are COMPUTATIONAL. Empirical attribution claims are CONDITIONAL. No empirical sector ranking is attempted.

Contents

1	The quotient and the question	4
2	Related work and intellectual boundary	4
2.1	Energy productivity and useful work	4
2.2	Lifecycle assessment and environmental-economic accounts	5
2.3	Welfare, index numbers, and composite indicators	6
2.4	Counterfactuals and attribution	6
3	Primitives and scope	6
3.1	Why seven fields	8
4	Well-posedness and the comparability gate	8
4.1	A practical gate	8
5	Scale transformations and what invariance means	9
6	Non-invariance and rank reversal	10
6.1	Valuation weights	10
6.2	System boundaries	11
6.3	Counterfactuals	11
6.4	Time horizons	12
6.5	Shared burdens and attribution	12
6.6	No universal scalar order	12
7	Worked examples	13
7.1	Value weights	13
7.2	Operational and lifecycle boundaries	13
7.3	Physical efficiency versus realized value	14
7.4	The danger of a gross ratio	14
8	Uncertainty and ratio estimands	14
8.1	Three different ratios	14
8.2	Sources of uncertainty	15
8.3	Recommended reporting	15
9	Measurement protocol	15
9.1	Step 1: state the decision and functional unit	15
9.2	Step 2: draw the physical boundary	16
9.3	Step 3: predeclare the counterfactual	16
9.4	Step 4: define the value functional	16
9.5	Step 5: set horizon and attribution	17
9.6	Step 6: specify uncertainty before estimation	17
9.7	Step 7: apply the comparability gate	17

10 Interpretation and policy use	18
10.1 What a high ratio can mean	18
10.2 Average and marginal questions	18
10.3 Private and social functionals	18
11 Implementation and executable evidence	18
12 Limitations and nonclaims	19
12.1 Empirical scope left open	20
13 Research implications	20
14 Conclusion	20
A Proof details and extensions	21
A.1 Weight regions for two outcomes	21
A.2 Strictly positive weights	21
A.3 Boundary threshold with unequal values	21
A.4 Discounted horizon reversal	22
A.5 Ratio decomposition	22
B Boundary manifest template	22
C Comparability worksheet	22
D Reproduction record	23
E Claim ledger	23
F Nonclaim checklist	23

1 The quotient and the question

Energy ratios have a reassuring appearance. A numerator is divided by joules, and the result seems to say how much an activity produces from a scarce physical input. The arithmetic can be correct while the comparison is false. One study may count electricity at a device. Another may count electricity at a facility, including cooling and power conversion. A third may add manufacturing energy. Their denominators are all measured in joules, but they describe different systems. The same problem appears above the line. Revenue, consumer surplus, avoided loss, completed tasks, and useful mechanical work are different quantities. Calling each one “value” does not make their units match.

The difficulty is old. Patterson’s survey of energy efficiency indicators distinguishes thermodynamic, physical, economic-thermodynamic, and economic measures, each suited to a different question [1]. Ayres and Warr argue that converted useful work, rather than raw energy alone, matters for historical production [2]. Warr and coauthors estimate useful work over a century in four economies [3], and Warr and Ayres later connect useful work with information in a growth model [4]. National statistical practice already publishes energy intensity and integrated environmental-economic accounts [8, 9]. Thus neither energy-normalized output nor the economic importance of useful work begins here.

This paper asks a narrower question: what must be fixed before a ratio of economic value to energy is interpretable, reproducible, and comparable? The answer is a seven-part measurement tuple. The tuple separates physical accounting from valuation and causal attribution. It also makes non-invariance visible. A reported ranking is conditional on the tuple unless an analyst proves robustness over a declared set of alternatives.

Three distinctions organize the argument.

1. Physical conversion efficiency compares commensurate physical inputs and outputs. Economic value per joule places a declared value functional in the numerator. The second does not follow from the first.
2. A ratio can be well defined for one study without being comparable to a ratio from another study. Reproducibility and cross-study comparability are separate gates.
3. Dependence on a value functional is not mere measurement noise. When outcomes are heterogeneous, a scalar ranking contains a weighting judgment.

The paper contributes four things. First, it defines the tuple and a typed ratio. Second, it proves constructive rank-reversal and non-invariance results. Third, it turns those results into a measurement protocol. Fourth, it maps the formal objects to executable, unit-checked examples. The claims are modest by design. A careful conditional statistic is more useful than a sweeping quotient whose numerator and denominator move whenever the application changes.

2 Related work and intellectual boundary

2.1 Energy productivity and useful work

The reciprocal pair “energy intensity” and “energy productivity” has long been used at sector and economy scale. If Y is a monetary output and E is energy, then Y/E is commonly called

energy productivity, while E/Y is energy intensity. The label does not settle whether Y is gross output, value added, welfare, or a volume index, nor whether E is primary or final energy. Patterson documents exactly this diversity [1]. The International Energy Agency likewise treats indicator construction as a matter of decomposition, activity, structure, and end-use efficiency rather than one universal number [5]. Ang’s practical guide to logarithmic mean Divisia decomposition shows how an aggregate change can be separated into activity, structure, and intensity terms [22].

Useful-work economics moves the physical account closer to end-use services. Ayres and Warr model production using capital, labor, and useful work [2]. Warr et al. construct long historical series by mapping energy carriers through conversion devices to useful work [3]. Warr and Ayres extend the account to information and communication technologies [4]. These contributions matter for the denominator and for the physical stages that precede economic outcomes. They do not imply that useful work, information, and welfare have a common natural unit. Their thermodynamic framing sits in a longer ecological-economics tradition that includes Georgescu-Roegen’s account of entropy and economic process [23].

The present framework is compatible with that literature. Exergy, final energy, and useful work can appear as separately typed nodes in an accounting chain. The framework’s proposed contribution is to join physical boundaries, counterfactual value, attribution, and uncertainty in one comparison contract. It does not rename existing energy productivity as a new theory.

2.2 Lifecycle assessment and environmental-economic accounts

ISO 14040 places goal and scope definition before lifecycle inventory and interpretation [6]. A functional unit provides the reference for comparison, while a system boundary determines which unit processes belong to the product system. ISO 14044 adds requirements and guidance for lifecycle inventory and critical review [7]. Rebitzer and coauthors survey the corresponding framework, goal-and-scope, and inventory practice [21]. Those principles carry directly to energy denominators. An operational wall-energy figure and a lifecycle energy figure are not rival estimates of one hidden number. They answer different scoped questions.

The System of Environmental-Economic Accounting Central Framework integrates physical flows, environmental assets, and economic accounts [8]. SEEA-Energy organizes energy resources and products in physical supply and use tables [9]. Its economic-account side is consistent with the production and accounting structure of the System of National Accounts [25]. Transformation inputs and outputs can both appear in a gross table, so analysts need bridge tables or net presentations when forming an intensity denominator. Adding primary fuel to the electricity made from that fuel would count one flow twice.

Our tuple borrows the discipline of goal, scope, boundary, and allocation from lifecycle assessment, then adds an explicit counterfactual and value functional. It also keeps primary, final, operational, and lifecycle energy conventions distinct. Paper 9 of the larger research program develops the national-account implications; this paper confines itself to the measurement foundation.

2.3 Welfare, index numbers, and composite indicators

Gross domestic product is a production measure, not a complete welfare functional. Fleurbaey reviews the long search for welfare measures beyond GDP and finds multiple defensible approaches rather than one settled substitute [10]. Sen’s treatment of collective choice supplies part of the social-choice background to that measurement problem [24]. Fleurbaey and Blanchet give a book-length treatment of welfare and sustainability measurement [11]. Stiglitz, Fitoussi, and Durand similarly argue for dashboards and distributional information alongside aggregate production [12].

Index-number theory shows that aggregation formulas embody assumptions about the underlying aggregator. Diewert’s exact and superlative indices do not evade this point; they make the connection between a formula and a flexible aggregator explicit [13]. The OECD and Joint Research Centre handbook treats normalization, weighting, aggregation, uncertainty, and sensitivity analysis as separate design steps for composite indicators [14]. Rankings can change when weights or aggregation rules change.

Value per joule inherits these concerns. Dividing a composite value index by energy does not make the weights objective. The denominator can sharpen the resource question, but it cannot choose a social welfare functional.

2.4 Counterfactuals and attribution

An incremental-value numerator requires a baseline. The potential-outcomes tradition distinguishes causal effects from observed associations [15, 16]. A unit cannot be observed both under an intervention and under its alternative, so identification rests on design and assumptions. Imbens and Rubin give a systematic treatment of randomized and observational designs [17]. Manski shows how partial identification can report informative bounds when point identification is not defensible [18].

The tuple therefore includes both a counterfactual c and an attribution rule a . The counterfactual states what would happen otherwise. The attribution rule states how evidence connects observed data to that contrast, and how shared outputs or burdens are assigned. Keeping these fields separate prevents a frequent error: writing down a causal-looking numerator while estimating only an observational ratio.

3 Primitives and scope

Definition 3.1 (System). A system s is an intervention, technology, organization, or bounded process for which an analyst defines a functional unit, an outcome record, and an energy inventory over a time interval.

The word “system” does not imply a closed thermodynamic system. Most economic applications are open systems that exchange matter, energy, information, and money with their surroundings. A boundary is an accounting choice around part of that open system.

Definition 3.2 (Outcome vector). For system s , let $x(s) \in \mathcal{X}$ denote a vector of consequences measured in their native units. Components may include tonnes moved, service hours, correctly

completed tasks, dollars of producer surplus, or avoided expected loss. Components with different units remain distinct until a value functional maps them to a scalar.

Definition 3.3 (Value functional). A value functional $W : \mathcal{X} \rightarrow \mathcal{V}$ maps an outcome record to a declared scalar value space $\mathcal{V}$. The codomain may be real currency of a stated price year, a money-metric welfare measure, or a named index. The functional includes its population, standing, discounting, distributional treatment, and treatment of external effects.

A functional contains more than a unit label. “2026 dollars” still leaves open whose willingness to pay counts, whether costs are netted, which externalities enter, and how future values are discounted. Two studies can share a currency and use different functionals.

Definition 3.4 (Counterfactual increment). Let c denote an explicit alternative to system s . The incremental value is

$$\Delta W(s \mid c) = W(x(s)) - W(x(c)),$$

or the corresponding unit-level causal contrast when outcomes are stochastic.

Definition 3.5 (Energy inventory). For boundary b , energy convention e , system s , and horizon τ , let $E_b^e(s; \tau) \in \mathbb{R}_{>0}$ be a non-overlapping energy total in joules. The superscript records the convention, such as device operational, system operational, facility operational, lifecycle, primary, final, or marginal-grid energy.

Strict positivity excludes division by zero. A net energy exporter can still be studied, but its gross inputs and exported energy must be recorded separately. A signed “net joule” denominator would change the interpretation and is outside the present definition.

Definition 3.6 (Measurement tuple). A value-per-joule measurement tuple is

$$\mathcal{M} = (b, c, W, \tau, e, a, \pi),$$

where:

1. b is the physical and organizational system boundary;
2. c is the counterfactual;
3. W is the value functional;
4. τ is the time horizon and discounting interval;
5. e is the energy convention;
6. a is the causal and joint-production attribution rule;
7. π is the uncertainty model.

Definition 3.7 (Value per joule). For a fixed tuple $\mathcal{M}$ and positive energy inventory,

$$\text{VPJ}_{\mathcal{M}}(s) = \frac{\mathbb{E}_{\pi}[\Delta W(s \mid c)]}{E_b^e(s; \tau)}.$$

If energy is itself stochastic, the analyst must declare whether the estimand uses expected numerator over expected denominator, an expected unit-level ratio, or another functional. These are generally different.

Remark 3.8 (Dimensional type). If W returns 2026 US dollars and energy is in joules, then $\text{VPJ}_{\mathcal{M}}$ has unit 2026 US dollars per joule. If W returns successful task equivalents, the unit is successful task equivalents per joule. The two ratios cannot be compared without an additional valuation map.

3.1 Why seven fields

Each tuple component blocks a distinct ambiguity. Boundary b says what equipment, infrastructure, upstream process, and downstream process belong to the study. Counterfactual c determines incrementality. Functional W defines the numerator. Horizon τ controls when benefits and energy enter. Convention e distinguishes stages of energy accounting. Attribution a allocates shared causes, assets, and burdens. Uncertainty model π states what is random, how dependence is represented, and what interval accompanies the point estimate.

None can be inferred reliably from the others. A lifecycle boundary does not select a counterfactual. A randomized counterfactual does not decide how to allocate a shared data center. A price-year label does not disclose a welfare functional. The tuple is intentionally redundant with a good study protocol: its purpose is to make omissions visible.

4 Well-posedness and the comparability gate

Definition 4.1 (Internally well posed). A reported $VPJ_{\mathcal{M}}(s)$ is internally well posed if:

1. every tuple component is declared;
2. numerator and denominator refer to the same functional unit and horizon;
3. the energy inventory is positive and contains no overlapping flows;
4. the value unit and energy unit are stated;
5. the uncertainty estimand matches the reported estimator.

Definition 4.2 (Direct comparability). Two results are directly comparable when their functional units match and their tuples are equal, except for representational transformations that have been proved order preserving. A documented bridge may replace exact equality for an energy convention or value unit when the bridge is applied to both results.

Proposition 4.3 (Exact tuple matching is an equivalence relation). *On the set of internally well-posed results, the relation “has the same functional unit and measurement tuple” is reflexive, symmetric, and transitive.*

Proof. Equality of the functional unit and each tuple component is reflexive, symmetric, and transitive. Their finite conjunction has the same properties. $\square$

PROVED. The result is elementary, but it matters operationally. It lets a registry partition results into comparison classes before anyone sorts their numerical values.

4.1 A practical gate

Before ranking two systems, an analyst should answer the following questions.

1. Do the functional units describe the same service?
2. Do boundaries include compatible physical components?
3. Are energy stages identical or connected by a documented bridge?
4. Do counterfactuals describe the same alternative?
5. Does W use the same population, prices, weights, and discount rule?
6. Do horizons match?

7. Are idle energy, failures, shared infrastructure, and embodied burdens allocated under compatible rules?
8. Are uncertainty intervals based on compatible estimands?

If one answer is no, the comparison is not automatically useless. It becomes a sensitivity exercise or an OBSTRUCTED direct comparison. The analyst can report each result in its own class, derive a bridge, or show a range over plausible tuples. What is not justified is a point ranking that suppresses the difference.

Table 1: Examples of mismatches that preserve the word “joule” but change the estimand.

Field	Result A	Result B
Boundary	accelerator device	facility meter including cooling
Energy convention	final electricity	primary energy equivalent
Counterfactual	no automation	older automated system
Value functional	private operating profit	social surplus net of emissions
Horizon	one benchmark run	five-year lifecycle
Attribution	marginal energy	average shared-infrastructure al- location
Uncertainty	conditional bootstrap	engineering tolerance only

5 Scale transformations and what invariance means

The word “invariant” needs care. A numerical ratio in dollars per joule changes when dollars become cents. A ranking need not change. Cardinal invariance and ordinal invariance are different properties.

Theorem 5.1 (Common positive value-unit rescaling). *Fix b, c, τ, e, a, π . Let $W'(x) = \lambda W(x)$ for one $\lambda > 0$, applied to every system. Then*

$$\text{VPJ}_{\mathcal{M}'}(s) = \lambda \text{VPJ}_{\mathcal{M}}(s).$$

For any systems r, s ,

$$\text{VPJ}_{\mathcal{M}}(r) > \text{VPJ}_{\mathcal{M}}(s) \iff \text{VPJ}_{\mathcal{M}'}(r) > \text{VPJ}_{\mathcal{M}'}(s).$$

Proof. Linearity of subtraction and expectation gives $\mathbb{E}[\Delta W'] = \lambda \mathbb{E}[\Delta W]$. The energy denominator is unchanged. Division yields the first equality. Multiplication by a common positive scalar preserves strict and weak order. $\square$

PROVED. Converting dollars to cents changes every magnitude by 100 but preserves a ranking. This is the appropriate currency-rescaling invariance test. Using different exchange-rate or purchasing-power adjustments for different systems is not the transformation in the theorem.

Proposition 5.2 (Affine translations cancel only under matched contrasts). *If $W'(x) = \lambda W(x) + \kappa$, with common $\lambda > 0$ and the same κ applied to both the intervention and its counterfactual, then $\Delta W'(s | c) = \lambda \Delta W(s | c)$. If the translation differs between intervention and baseline, cancellation fails.*

Proof. Under a common translation,

$$W'(x(s)) - W'(x(c)) = \lambda W(x(s)) + \kappa - \lambda W(x(c)) - \kappa = \lambda \Delta W(s | c).$$

Unequal translations leave their difference in the numerator. $\square$

The proposition explains why the counterfactual is part of unit handling. Incremental value is invariant to a common origin shift, but a gross revenue ratio need not represent incremental value at all.

6 Non-invariance and rank reversal

6.1 Valuation weights

Suppose a system and its counterfactual produce m heterogeneous outcomes. Write the componentwise increment as $\Delta x(s | c) = x(s) - x(c)$. For nonnegative weights $w \in \mathbb{R}_{\geq 0}^m$, consider the linear functional $W_w(x) = w^\top x$. Define the incremental-outcome-per-energy vector

$$z(s) = \frac{\Delta x(s | c)}{E(s)},$$

where division acts componentwise.

Theorem 6.1 (Scalarization rank reversal). *Let systems r and s have positive energy denominators. If neither $z(r)$ nor $z(s)$ weakly dominates the other, then there exist nonnegative, nonzero weight vectors w and v such that*

$$\text{VPJ}_w(r) > \text{VPJ}_w(s) \quad \text{and} \quad \text{VPJ}_v(r) < \text{VPJ}_v(s).$$

Proof. Failure of weak dominance in both directions gives coordinates i and j with $z_i(r) > z_i(s)$ and $z_j(r) < z_j(s)$. Choose w as the i -th coordinate vector and v as the j -th coordinate vector. Both are nonnegative and nonzero. Then $w^\top z(r) > w^\top z(s)$, while $v^\top z(r) < v^\top z(s)$. $\square$

PROVED. The theorem does not say every small change of weights reverses a ranking. It says a universal scalar order is unavailable for non-dominating vectors unless the admissible valuation class is restricted. The restriction may be ethically or institutionally justified, but it must be published.

Corollary 6.2 (Robust ranking under componentwise dominance). *If $z(r) \geq z(s)$ componentwise, then $w^\top z(r) \geq w^\top z(s)$ for every nonnegative w . If at least one strictly better component receives positive weight, the ranking is strict.*

Proof. Every component of $z(r) - z(s)$ is nonnegative, so its inner product with a nonnegative vector is nonnegative. Strictness follows when one positive difference has positive weight. $\square$

Dominance is the strongest defensible cross-weight conclusion. When it fails, a Pareto frontier communicates more than a single league table.

6.2 System boundaries

Theorem 6.3 (Boundary expansion can reverse a ranking). *Let two systems have positive values V_r, V_s and operational energies $E_r, E_s > 0$. Suppose*

$$\frac{V_r}{E_r} > \frac{V_s}{E_s}.$$

If an admissible expanded boundary assigns additional energy $K > 0$ to system r and none to s , then the ranking reverses whenever

$$K > \frac{V_r E_s}{V_s} - E_r.$$

Proof. The expanded-boundary ranking favors s exactly when

$$\frac{V_r}{E_r + K} < \frac{V_s}{E_s}.$$

All denominators and values are positive. Cross multiplication and rearrangement give the stated threshold. $\square$

PROVED, conditional on the accounting admissibility of the boundary expansion. The theorem is not permission to add arbitrary energy to a disfavored system. The added burden must belong to the declared expanded boundary under a consistent allocation rule. Examples include manufacturing energy for dedicated hardware or a facility load omitted from a device measurement.

Remark 6.4 (Symmetric expansions). Both systems may receive added burdens. A reversal occurs when

$$\frac{V_r}{E_r + K_r} < \frac{V_s}{E_s + K_s}.$$

No qualitative result follows from the word “lifecycle” alone. The quantities K_r and K_s , their lifetimes, and their allocation across functional units determine the comparison.

6.3 Counterfactuals

Proposition 6.5 (Gross outcomes do not identify incremental rankings). *On an unrestricted scalar outcome space, let observed outcomes satisfy $Y_r > Y_s$. For any desired strict ordering of incremental values, there exist counterfactual outcomes $Y_r(0)$ and $Y_s(0)$ that produce it through $\Delta_r = Y_r - Y_r(0)$ and $\Delta_s = Y_s - Y_s(0)$.*

Proof. Choose any target increments δ_r, δ_s with the desired ordering, then set $Y_r(0) = Y_r - \delta_r$ and $Y_s(0) = Y_s - \delta_s$. The observed outcomes remain fixed while the increments equal the targets. $\square$

PROVED as a non-identification construction. Empirical counterfactuals cannot be chosen after seeing the desired answer. The construction shows why observed revenue per joule alone does not identify incremental value per joule. Design, institutional knowledge, or defensible bounds must restrict the counterfactual class.

6.4 Time horizons

Proposition 6.6 (Horizon reversal). *There exist systems r, s with equal energy such that r has higher value per joule at horizon τ_1 , while s has higher value per joule at a longer horizon τ_2 .*

Proof. Let both systems use one joule at time zero. Let r produce value 10 at time one and no later value. Let s produce value 6 at time one and 8 at time two. At horizon one, the ratios are 10 and 6. At horizon two, with no discounting, they are 10 and 14. Positive discounting preserves a reversal for a nonempty range of discount factors. $\square$

PROVED. Delayed maintenance costs, learning effects, degradation, and asset life can generate the same structure. A horizon should therefore be chosen from the decision problem, not from whichever endpoint favors a technology.

6.5 Shared burdens and attribution

Proposition 6.7 (Allocation-induced reversal). *Let systems r and s have equal positive value V , dedicated energies $E_r < E_s$, and a shared energy burden $H > 0$. Allocate fraction $\alpha \in [0, 1]$ of H to r and the remainder to s . If*

$$H > E_s - E_r,$$

then there are admissible allocations under which either system ranks first.

Proof. At $\alpha = 0$, denominators are E_r and $E_s + H$, so r ranks first. At $\alpha = 1$, denominators are $E_r + H$ and E_s . The condition implies $E_r + H > E_s$, so s ranks first. $\square$

PROVED. An admissible allocation still requires a rationale, such as measured marginal use, capacity reservation, run time, or causal responsibility. Equal splitting, revenue shares, and peak-capacity shares are not interchangeable merely because they sum to the same total.

6.6 No universal scalar order

Theorem 6.8 (Tuple-relative ordering). *Consider a domain that contains at least two systems with non-dominating outcome-per-energy vectors and permits the coordinate valuations used in theorem 6.1. No complete scalar ranking of that domain is invariant to all admissible value functionals.*

Proof. Assume a complete invariant scalar ranking exists. For the two non-dominating systems, completeness ranks one weakly above the other or treats them as tied. Theorem 6.1 supplies one admissible valuation that ranks the first strictly above the second and another that ranks it strictly below. Either strict ranking contradicts a fixed tie, and one contradicts either fixed strict order. Thus the assumed invariant ranking does not exist. $\square$

PROVED under the stated valuation class. The theorem is deliberately weaker than a social-choice impossibility theorem. It needs no claim about collective preference aggregation. It only records a fact about scalarizing heterogeneous outcomes.

7 Worked examples

7.1 Value weights

Consider two stylized systems that each use 100 joules. Relative to a zero-vector counterfactual, system A adds nine verified task completions and two units of avoided loss. System B adds three verified task completions and eight units of avoided loss. Their incremental-outcome-per-energy vectors are

$$z(A) = (0.09, 0.02), \quad z(B) = (0.03, 0.08).$$

If the functional assigns one index point to each completed task and zero to avoided loss, A ranks first. If it assigns one point to each avoided-loss unit and zero to task count, B ranks first. Neither vector dominates.

Table 2: A valuation-weight reversal. Values are stylized and carry no empirical sector claim.

System	Energy (J)	Tasks	Avoided loss	Preferred by
A	100	9	2	task-only functional
B	100	3	8	loss-only functional

The correct response is not to average the dimensions by default. One can publish both components, state a reason for weights, and show a sensitivity region. For linear weights $w_T, w_L \geq 0$, A ranks first when

$$0.09w_T + 0.02w_L > 0.03w_T + 0.08w_L,$$

or $w_T > w_L$. B ranks first when $w_T < w_L$, and they tie when the weights are equal. The threshold is transparent.

7.2 Operational and lifecycle boundaries

Now give both systems the same realized value of ten currency units. Under an operational boundary, A uses 40 joules and B uses 50. Their ratios are 0.25 and 0.20 currency units per joule, so A ranks first. An expanded boundary assigns 100 joules of embodied energy to A and 20 to B. The lifecycle denominators become 140 and 70. The ratios become approximately 0.0714 and 0.1429, so B ranks first.

Table 3: Boundary reversal with equal numerator.

System	Value	Operational E	Embodied E	Operational VPJ	Lifecycle VPJ
A	10	40	100	0.2500	0.0714
B	10	50	20	0.2000	0.1429

This example is COMPUTATIONAL. It verifies the construction in theorem 6.3; it is not a lifecycle assessment of a real technology. An empirical study would need inventories, lifetimes, utilization, allocation, geography, and uncertainty.

7.3 Physical efficiency versus realized value

Suppose heat system H converts 90 percent of final electricity into useful heat, while system L converts 70 percent. On a matched physical functional unit and boundary, H is more physically efficient. Now suppose H serves a process that often produces inventory with no buyer, while L serves a temperature-sensitive repair whose completion avoids a costly outage. L can have higher expected realized value per joule despite lower heat-delivery efficiency.

There is no paradox. Physical efficiency answers how much useful heat leaves the device for a given physical input. Economic VPJ also depends on the use, counterfactual, prices or welfare weights, reliability, and attribution. The physical result remains true within its domain. It is one input to the economic account, not a complete valuation.

7.4 The danger of a gross ratio

Two sites report annual revenue and electricity:

$$\frac{Y_A}{E_A} = 12 \text{ currency/MJ}, \quad \frac{Y_B}{E_B} = 9 \text{ currency/MJ}.$$

The first site appears better. But suppose the intervention studied at A replaced a process that already earned 11 currency/MJ, while B replaced one that earned 3 currency/MJ, on matched energy bases. The incremental comparisons are one and six currency/MJ. Gross output and incremental contribution answer different questions.

The numerical subtraction alone does not establish causality. The baselines would need evidence. A randomized rollout, credible natural experiment, engineering model, or partial-identification bound could support them. The example's role is to make the estimand distinction concrete.

8 Uncertainty and ratio estimands

8.1 Three different ratios

Let V and $E > 0$ be random. Analysts may encounter:

$$\theta_1 = \frac{\mathbb{E}[V]}{\mathbb{E}[E]}, \quad \theta_2 = \mathbb{E} \left[\frac{V}{E} \right], \quad \theta_3 = \frac{\sum_i V_i}{\sum_i E_i}.$$

The third is a sample ratio estimator for the first under common sampling conditions. The second is the mean of unit-level ratios. In general, $\theta_1 \neq \theta_2$. If high-value observations also use more energy, an unweighted mean of unit ratios can answer a particularly different question from an aggregate ratio. In the notation of definition 3.7, V is $\Delta W(s | c)$ and E is $E_b^e(s; \tau)$; this subsection relaxes the definition's default treatment of the denominator as fixed.

Proposition 8.1 (Expectation does not commute with division). *There exist positive random energies E and values V such that*

$$\mathbb{E}[V/E] \neq \mathbb{E}[V]/\mathbb{E}[E].$$

Proof. Let $V = 1$ always and let E equal 1 or 2 with equal probability. Then

$$\mathbb{E}[V/E] = \frac{1}{2}(1 + \frac{1}{2}) = \frac{3}{4}, \quad \mathbb{E}[V]/\mathbb{E}[E] = 1/(\frac{3}{2}) = \frac{2}{3}.$$

□

PROVED. The uncertainty component π must name the estimand before it names a confidence-interval algorithm.

8.2 Sources of uncertainty

A useful uncertainty ledger separates:

1. measurement error in meters, prices, and outcome records;
2. sampling variation across tasks, days, sites, or users;
3. model uncertainty in counterfactual and causal assumptions;
4. scenario uncertainty in future demand, lifetime, and discounting;
5. normative uncertainty over weights and social standing;
6. boundary uncertainty over omitted or shared processes.

The first two may yield to repeated measurement and conventional intervals. The last four often require sensitivity analysis, scenario sets, or bounds. A narrow bootstrap around one selected model does not represent uncertainty about the model or value functional. Saltelli and coauthors provide a general framework for global sensitivity analysis when uncertain inputs interact [19].

8.3 Recommended reporting

For an empirical ratio, report the numerator and denominator separately before the quotient. Include their covariance when both are estimated from the same sample. If the denominator is noisy or can approach zero, ordinary symmetric intervals can fail. The present framework requires positive energy and recommends Fieller-type or bootstrap methods only after checking their assumptions [20]. When causal identification is weak, report bounds on ΔW and propagate those bounds through the positive denominator.

Sensitivity to tuple choices is not summarized adequately by one standard error. A robust report contains:

- a statistical interval conditional on the chosen tuple;
- a scenario range over defensible boundaries and horizons;
- a weight region or Pareto statement for heterogeneous outcomes;
- an identification statement for the counterfactual;
- a list of unresolved exclusions.

9 Measurement protocol

9.1 Step 1: state the decision and functional unit

Start with the decision the metric will inform. “Which server is more efficient?” is incomplete. A functional unit might be one thousand requests meeting a named accuracy and latency

threshold, one tonne-kilometre delivered, or one year of a settlement service under a named threat model. The unit should describe a comparable service, not an easy intermediate count.

Record the population and deployment context. A benchmark mixture, geographic site, and service-level constraint can all change outcomes and energy. If the functional unit cannot be matched, stop before forming a comparative ratio.

9.2 Step 2: draw the physical boundary

List included equipment and processes. Mark the meter location. Record idle energy, warm-up, failures, retries, power conversion, cooling, networking, and shared infrastructure as included, excluded, or separately allocated. For a lifecycle study, add manufacturing, construction, maintenance, and end-of-life processes without overlapping the operational inventory.

Choose one energy convention. Device operational energy is appropriate for a low-level engineering question. Wall energy fits many deployed computing comparisons. Facility operational energy adds site overhead. Lifecycle energy adds non-overlapping embodied burdens. Primary and final energy require a bridge through conversion losses. Marginal-grid energy is indexed by location and time. These conventions should appear as enum-like labels in data, not as prose buried in a methods appendix.

9.3 Step 3: predeclare the counterfactual

Name what would happen without the system. “No energy use” is rarely a credible baseline when another process would supply the service. Specify whether the alternative is no service, manual work, an incumbent technology, delayed action, or a portfolio response.

Then state the identification strategy. Random assignment can identify an average effect for the assigned population under its assumptions. Observational designs need exchangeability, timing, exclusion, or structural assumptions. Engineering counterfactuals need validation. If none is credible, an OPEN or partially identified numerator is more honest than a precise gross ratio.

9.4 Step 4: define the value functional

Write W as an auditable calculation. For a monetary functional, state:

- currency, price year, and conversion method;
- whether the measure is revenue, profit, surplus, or cost saving;
- whose costs and benefits count;
- treatment of taxes, transfers, externalities, and distribution;
- discount rate and terminal value;
- treatment of risk and catastrophic loss.

For a nonmonetary index, state component units, normalization, weights, and aggregation rule. Publish the component vector. A composite should be a view of the data, not its only surviving form.

9.5 Step 5: set horizon and attribution

The horizon should include relevant benefits, operating energy, degradation, maintenance, and replacement. If horizons differ across systems, annualize or discount with a documented bridge. Do not count a full asset benefit against one day of energy or a full manufacturing burden against one unusually small batch.

For joint production, identify dedicated and shared flows. Prefer causal or measured marginal allocation when available. Otherwise show results under more than one defensible rule. Capacity share, run time, revenue share, mass, and equal allocation answer different questions. The allocation rule belongs in the tuple because proposition 6.7 shows it can decide the ranking.

9.6 Step 6: specify uncertainty before estimation

List random variables, dependence, sampling unit, missingness, and estimator. Predeclare whether the target is a ratio of expectations or an expectation of ratios. Record meter accuracy and calibration, data exclusions, and failed observations. Run sensitivity over boundary, baseline, weights, horizon, and allocation, as well as sampling noise.

9.7 Step 7: apply the comparability gate

Compare tuples field by field. If a bridge is used, publish the equation, source data, and uncertainty. Report a scalar ranking only inside a comparison class. Outside it, report a vector, Pareto frontier, scenario range, or OBSTRUCTED comparison.

Table 4: Minimum machine-readable result record.

Field	Required content
result identifier	immutable study and scenario identifier
functional unit	quantified service and quality threshold
boundary	included and excluded components
counterfactual	alternative and identification design
value functional	formula, unit, standing, prices, weights
horizon	start, end, discounting, lifetime treatment
energy convention	named stage and meter or inventory source
attribution rule	joint outputs and shared energy allocation
uncertainty model	estimand, random variables, intervals, scenarios
numerator	estimate, unit, interval, provenance
denominator	estimate, joules, interval, provenance
status	proved, conditional, computational, obstructed, or open

10 Interpretation and policy use

10.1 What a high ratio can mean

Inside a matched tuple, a higher ratio means more expected declared incremental value per included joule. That sentence is intentionally repetitive. Remove “declared,” “incremental,” or “included,” and the claim grows beyond the measurement.

A high ratio can result from a larger causal benefit, a smaller energy inventory, or both. The decomposition matters for policy. A low-energy, low-value activity and a high-energy, high-value activity can have the same ratio but different scale, risk, and constraints. Ratios should therefore be reported beside absolute energy and absolute value.

10.2 Average and marginal questions

An average lifecycle VPJ is not the marginal value of the next joule. The latter depends on time, location, constraints, and displacement. Average metrics may support benchmarking or historical accounting. Marginal allocation requires a response function and an opportunity cost. Paper 6 of this program develops that optimization problem.

The distinction also blocks a policy error. A sector with high average value per joule need not be the best recipient of one additional joule if it is at capacity, has diminishing returns, or cannot use energy at the available node and time. Conversely, a low average ratio does not prove that every marginal use is low value.

10.3 Private and social functionals

A firm’s profit functional can be useful for operating decisions. A social welfare functional may add consumer surplus, environmental damages, public risk, and distribution. Neither should be mislabeled as the other. Taxes and prices may internalize some external effects but not all.

Reporting both functionals can reveal the source of disagreement. If a system ranks well under private profit and poorly under a social functional, the conflict is substantive. Hiding it inside one “value” number removes information that policy needs.

11 Implementation and executable evidence

The companion module `src/joule_standard/foundations.py` uses only the Python standard library. Decimal arithmetic rejects binary floating-point input at the public quantity boundary. Energy is stored in joules and must be strictly positive. Kilowatt-hours convert by the exact identity

$$1 \text{ kWh} = 3,600,000 \text{ J.}$$

The `MeasurementTuple` data class has fields for the seven declarations in definition 3.6. `ValueFunctional` carries a name and output unit. `IncrementalValue` carries an exact amount and unit. The `value_per_joule` function rejects a numerator whose unit differs from the declared functional. A `ValuePerJoule` comparison rejects unequal tuples or value units.

Incremental outcome vectors, with a zero vector as the counterfactual in the published fixtures, and nonnegative linear valuations implement theorem 6.1. The fixtures reproduce the weight reversal in table 2 and the boundary reversal in table 3. Tests also check:

1. exact joule and kilowatt-hour conversion;
2. presence of all seven tuple fields;
3. positive-energy and unit mismatch rejection;
4. refusal to compare different boundaries;
5. common currency-rescaling order invariance;
6. strict valuation-weight rank reversal;
7. strict boundary rank reversal;
8. missing outcome-dimension rejection.

These tests are COMPUTATIONAL. The algebraic theorems have separate proofs in the paper. Passing a finite test suite does not prove a universal statement, and the implementation does not estimate empirical values on its own.

12 Limitations and nonclaims

Limitation 12.1 (No natural welfare unit). The framework requires W ; it does not derive a unique W from thermodynamics. Value is not conserved like energy. Money, utility, capability, rights, security, and task success do not become one physical dimension because each is divided by joules.

Limitation 12.2 (No energy theory of value). Energy is necessary for material production and computation, but energy input alone does not determine exchange value or welfare. Scarcity, preferences, institutions, knowledge, location, timing, complementary inputs, and distribution matter. This paper does not revive a caloric theory of price.

Limitation 12.3 (No universal sector ranking). The rank-reversal results rule out the paper’s own use as a universal league table for artificial intelligence, cryptographic settlement, manufacturing, transport, health, or other sectors. Cross-sector comparison requires a common functional unit or an explicit welfare functional and matched boundaries.

Limitation 12.4 (No causal effect from an observational ratio). Observed value divided by observed energy is descriptive. It becomes an incremental causal ratio only under a design or identification argument that supports the counterfactual numerator. Adding more decimal places does not fix confounding.

Limitation 12.5 (No automatic lifecycle superiority). A broader boundary can be relevant, but breadth is not accuracy by itself. Lifecycle inventories contain allocation, lifetime, geography, and data-quality choices. Operational and lifecycle figures should be labeled and used for their respective questions.

Limitation 12.6 (No sufficiency of the ratio). Even a well-posed VPJ statistic omits absolute scale, budget constraints, nonconvexities, tail risk, rights, feasibility, and distribution unless W and the decision model include them. A ratio is one statistic, not a complete social decision procedure.

12.1 Empirical scope left open

This foundation paper does not estimate a national energy-productivity series, an AI benchmark, or proof-of-work assurance. Those applications need domain specific functional units and data. Their empirical status remains OPEN here. The absence is deliberate: a synthetic example should not be mistaken for a measured sector claim.

13 Research implications

The tuple provides a common interface for later work without forcing common units. A typed transduction model can track energy, computation, immediate output, verified outcome, adoption, and realized value as separate stages. Quality-adjusted machine intelligence can use task success per wall joule while stopping short of realized business value. Proof-of-work analysis can measure scenario-conditioned settlement assurance without calling energy itself “trust.” Marginal allocation can operate on response curves rather than average ratios. National accounts can publish vectors and bridge tables before any composite.

Several research questions remain.

1. Which tuple components can be standardized across domains, and which must remain application specific?
2. How should partial identification of the numerator interact with inventory uncertainty in the denominator?
3. What dominance or robust-order regions survive over a publicly defensible class of welfare weights?
4. How should shared digital and physical infrastructure be allocated when marginal use, reservation capacity, and causal responsibility disagree?
5. Which provenance format lets auditors reproduce every bridge without exposing confidential unit-level data?

The most useful near-term output is not one global number. It is a registry of typed results whose comparison classes, bridges, and sensitivity regions are machine readable.

14 Conclusion

Value per joule is meaningful only after the analyst says whose value, relative to what alternative, over which period, inside which boundary, under which energy convention, with what attribution rule, and under what uncertainty model. The tuple $\mathcal{M} = (b, c, W, \tau, e, a, \pi)$ records those choices.

Once the tuple is fixed, the quotient is ordinary and useful. A common positive currency rescaling preserves order. Matched tuples define clear comparison classes. Numerator and denominator can be audited separately. The statistic can support engineering, operating, investment, or policy analysis within its declared scope.

The same formalism explains the limit. Non-dominating outcomes reverse under admissible weights. Boundaries, counterfactuals, horizons, and shared-burden allocations can reverse rankings too. These dependencies are not an excuse to abandon measurement. They are

instructions for honest measurement: publish the tuple, retain the outcome vector, test sensitivity, and decline comparisons that fail the gate.

A Proof details and extensions

A.1 Weight regions for two outcomes

Let the normalized difference between two systems be $d = z(r) - z(s) = (d_1, d_2)$, with $d_1 > 0 > d_2$. Normalize nonnegative weights so $w_1 + w_2 = 1$. System r ranks first when

$$w_1 d_1 + (1 - w_1) d_2 > 0.$$

Solving gives

$$w_1 > \frac{-d_2}{d_1 - d_2}.$$

The right side lies strictly between zero and one. Thus the admissible simplex contains a nonempty region favoring each system and one tie point. This gives a complete sensitivity diagram for two components.

With $m > 2$, the robust-order region is the intersection of the weight simplex with a half-space:

$$\mathcal{W}_{r \succeq s} = \{w \in \mathbb{R}_{\geq 0}^m : \mathbf{1}^\top w = 1, w^\top [z(r) - z(s)] \geq 0\}.$$

Publishing this polytope is often more informative than publishing one selected weight vector.

A.2 Strictly positive weights

Theorem 6.1 uses coordinate vectors, which permit zero weights. If the admissible class requires every component to have positive weight, the reversal still holds for sufficiently small perturbations. Let e_i select a coordinate where r is better. Define

$$w_\epsilon = (1 - (m - 1)\epsilon)e_i + \epsilon \sum_{k \neq i} e_k$$

for $0 < \epsilon < 1/m$. Continuity of the inner product preserves the strict inequality for all sufficiently small ϵ . Apply the same construction at a coordinate where s is better.

A.3 Boundary threshold with unequal values

For positive V_r, V_s, E_r, E_s , a burden pair $K_r, K_s \geq 0$ favors s when

$$V_r(E_s + K_s) < V_s(E_r + K_r).$$

Solving for K_r gives

$$K_r > \frac{V_r}{V_s}(E_s + K_s) - E_r.$$

The expression is a sensitivity threshold, not an empirical estimate. A study can place distributions or inventory intervals on K_r and K_s , then report the probability or scenario share on either side.

A.4 Discounted horizon reversal

In proposition 6.6, let the time-two discount factor be $\delta \in (0, 1]$. System B’s long-horizon value is $6 + 8\delta$. It exceeds A’s value of 10 whenever $\delta > 1/2$. Thus the reversal persists for discount rates whose one-period factor exceeds one half. The result does not select a discount rate. It identifies the threshold at which the decision changes.

A.5 Ratio decomposition

For two tuples that differ in both numerator and denominator, a logarithmic decomposition can be useful when values remain positive:

$$\log \frac{VPJ_2}{VPJ_1} = \log \frac{V_2}{V_1} - \log \frac{E_2}{E_1}.$$

This identity separates value change from energy change. It does not attribute causality, and it fails when incremental value is zero or negative. Those cases should be reported directly rather than repaired by an arbitrary offset.

B Boundary manifest template

An empirical study can use the following template.

1. **Study identifier.** Immutable name, version, date, and owner.
2. **Decision.** The choice the statistic is meant to inform.
3. **Functional unit.** Quantity, quality threshold, location, and service conditions.
4. **Physical system.** Diagram or inventory of included processes.
5. **Temporal interval.** Measurement window, benefit horizon, lifetime, and discounting.
6. **Geography.** Site, grid node or region, and market jurisdiction.
7. **Counterfactual.** Alternative process and evidence supporting it.
8. **Energy stage.** Device, system, facility, lifecycle, primary, final, or marginal grid.
9. **Meter and inventory.** Instrument, calibration, sampling rate, missingness, and data lineage.
10. **Idle and failed work.** Inclusion and allocation.
11. **Embodied energy.** Included assets, lifetime, utilization, and allocation.
12. **Value functional.** Formula, unit, prices, standing, external effects, distribution, and risk.
13. **Attribution.** Causal design and joint-production allocation.
14. **Uncertainty.** Estimand, stochastic model, scenarios, and sensitivity ranges.
15. **Exclusions.** Every known omitted process or consequence.

C Comparability worksheet

The disposition “stop” means stop the direct point ranking, not stop the research. A mismatch can motivate a new measurement or a useful sensitivity analysis.

Table 5: A worksheet for two candidate results.

Check	Result A	Result B	Disposition
Functional unit			match, bridge, stop
Boundary			match, scenario
Counterfactual			re-estimate
Value functional			weight region
Horizon			discount bridge
Energy convention			energy bridge
Attribution			sensitivity
Uncertainty estimand			restate

D Reproduction record

The reference tests run with Python 3.11 or later and pytest. From the repository root:

```
python -m pytest tests/test_foundations.py
```

The code uses no runtime dependency outside the standard library. Tests use pytest as the runner. Exact decimal-like inputs must be integers, strings, or `Decimal` values; public constructors reject binary floats. This is a small guard against pretending that a display-rounded decimal is exact.

The two rank-reversal fixtures are intentionally transparent. They do not depend on random seeds, hidden data, or fitted parameters. Their purpose is to make the existence proofs executable. Any empirical application should live in a separate module with source provenance and should not replace these fixtures with paper-specific hard-coded acceptance checks.

E Claim ledger

F Nonclaim checklist

For clarity, the following statements are outside the paper’s claims:

1. GDP divided by energy is new.
2. Useful-work economics is new.
3. Joules are a currency, utility unit, or conserved measure of value.
4. Physical conversion efficiency determines economic value.
5. The framework supplies a politically neutral welfare functional.
6. Artificial intelligence, Bitcoin, manufacturing, and transport have a natural universal ranking.
7. An observational value-energy ratio identifies a causal effect.
8. Lifecycle accounting eliminates allocation choices.
9. A finite Python test suite proves the general theorems.
10. A high average ratio proves that the next joule should be allocated to the same system.

Table 6: Claims made in this paper and their evidence status.

Identifier	Claim	Status
JS-C001	VPJ is well defined only relative to a declared tuple and positive, typed denominator.	PROVED
JS-C002a	Non-dominating outcome vectors admit valuation-weight reversal.	PROVED
JS-C002b	An admissible boundary expansion can reverse a ranking.	PROVED
JS-C002c	Gross outcomes do not identify incremental-value rankings.	PROVED
JS-C002d	Horizon and shared-burden allocation can reverse rankings.	PROVED
JS-COMP-1	The published Python fixtures reproduce currency scaling and two rank reversals.	COMPUTATIONAL
JS-EMP-1	Any real sector ranking under this framework.	OPEN

References

- [1] M. G. Patterson. What is energy efficiency? Concepts, indicators and methodological issues. *Energy Policy*, 24(5):377 to 390, 1996. [https://doi.org/10.1016/0301-4215\(96\)00017-1](https://doi.org/10.1016/0301-4215(96)00017-1).
- [2] R. U. Ayres and B. Warr. Accounting for growth: The role of physical work. *Structural Change and Economic Dynamics*, 16(2):181 to 209, 2005. <https://doi.org/10.1016/j.strueco.2003.10.003>.
- [3] B. Warr, R. U. 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*, 69(10):1904 to 1917, 2010. <https://doi.org/10.1016/j.ecolecon.2010.03.021>.
- [4] B. Warr and R. U. Ayres. Useful work and information as drivers of economic growth. *Ecological Economics*, 73:93 to 102, 2012. <https://doi.org/10.1016/j.ecolecon.2011.09.006>.
- [5] International Energy Agency. Energy End-uses and Efficiency Indicators Data Explorer. <https://www.iea.org/data-and-statistics/data-tools/energy-end-uses-and-efficiency-indicators-data-explorer>. Accessed 23 July 2026.
- [6] International Organization for Standardization. ISO 14040:2006, Environmental management: Life cycle assessment, principles and framework. <https://www.iso.org/standard/37456.html>.

- [7] International Organization for Standardization. ISO 14044:2006, Environmental management: Life cycle assessment, requirements and guidelines. <https://www.iso.org/standard/38498.html>.
- [8] United Nations, European Commission, Food and Agriculture Organization, International Monetary Fund, OECD, and World Bank. *System of Environmental-Economic Accounting 2012: Central Framework*. United Nations, 2014. https://unstats.un.org/unsd/envaccounting/seeaRev/SEEA_CF_Final_en.pdf.
- [9] United Nations Statistics Division. *System of Environmental-Economic Accounting for Energy*. <https://seea.un.org/content/seea-energy>.
- [10] M. Fleurbaey. Beyond GDP: The quest for a measure of social welfare. *Journal of Economic Literature*, 47(4):1029 to 1075, 2009. <https://doi.org/10.1257/jel.47.4.1029>.
- [11] M. Fleurbaey and D. Blanchet. *Beyond GDP: Measuring Welfare and Assessing Sustainability*. Oxford University Press, 2013. <https://doi.org/10.1093/acprof:oso/9780199767199.001.0001>.
- [12] J. E. Stiglitz, J. P. Fitoussi, and M. Durand. *Beyond GDP: Measuring What Counts for Economic and Social Performance*. OECD Publishing, 2018. <https://doi.org/10.1787/9789264307292-en>.
- [13] W. E. Diewert. Exact and superlative index numbers. *Journal of Econometrics*, 4(2):115 to 145, 1976. [https://doi.org/10.1016/0304-4076\(76\)90009-9](https://doi.org/10.1016/0304-4076(76)90009-9).
- [14] OECD, European Union, 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>.
- [15] D. B. Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5):688 to 701, 1974. <https://doi.org/10.1037/h0037350>.
- [16] P. W. Holland. Statistics and causal inference. *Journal of the American Statistical Association*, 81(396):945 to 960, 1986. <https://doi.org/10.1080/01621459.1986.10478354>.
- [17] 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>.
- [18] C. F. Manski. *Identification for Prediction and Decision*. Harvard University Press, 2007. <https://www.hup.harvard.edu/books/9780674026537>.
- [19] A. Saltelli, M. Ratto, T. Andres, F. Campolongo, J. Cariboni, D. Gatelli, M. Saisana, and S. Tarantola. *Global Sensitivity Analysis: The Primer*. Wiley, 2008. <https://doi.org/10.1002/9780470725184>.

- [20] E. C. Fieller. Some problems in interval estimation. *Journal of the Royal Statistical Society, Series B*, 16(2):175 to 185, 1954. <https://doi.org/10.1111/j.2517-6161.1954.tb00159.x>.
- [21] G. Rebitzer et al. Life cycle assessment, part 1: Framework, goal and scope definition, inventory analysis, and applications. *Environment International*, 30(5):701 to 720, 2004. <https://doi.org/10.1016/j.envint.2003.11.005>.
- [22] B. W. Ang. The LMDI approach to decomposition analysis: A practical guide. *Energy Policy*, 33(7):867 to 871, 2005. <https://doi.org/10.1016/j.enpol.2003.10.010>.
- [23] N. Georgescu-Roegen. *The Entropy Law and the Economic Process*. Harvard University Press, 1971. <https://doi.org/10.4159/harvard.9780674281653>.
- [24] A. Sen. *Collective Choice and Social Welfare*. Holden-Day, 1970. Expanded edition, Harvard University Press, 2017. <https://www.hup.harvard.edu/books/9780674919211>.
- [25] European Commission, International Monetary Fund, OECD, United Nations, and World Bank. *System of National Accounts 2008*. <https://unstats.un.org/unsd/nationalaccount/docs/SNA2008.pdf>.