SETTLEMENT LOOPS AND THE PROJECTION DEFECT: A HOMOLOGICAL COUPLING A CENTRAL PRICE REPRESENTS BUT NO AGENT-LOCAL ACCOUNT INTERNALIZES

JONATHAN WASHBURN

Abstract. Model a settlement economy as a finite graph and value a subsystem by its first Betti number, the number of independent closed recognition loops. When two value-bearing loci are joined by a settlement seam carrying s cross-ledger postings, the joint value exceeds the sum of the two private values by a topological term, the *projection defect* $C = \beta_1(A \cup B) - \beta_1(A) - \beta_1(B) = \max\{s-1, 0\}$, strictly positive under any settled trade $s \geq 2$. The defect is the rank of the cokernel of the canonical inclusion $H_1(A) \oplus H_1(B) \hookrightarrow H_1(G)$, a basis-independent invariant. We are explicit about the boundary of the claim. A central observer who treats the seam count s as an observable commodity represents $Z = \beta_1(A \cup B)$ exactly, since Z is a closed-form function of s; this is no impossibility theorem for welfare analysis, and we concede the centralized representation at the outset. The content sits one level beneath that fix and is twofold. First, a *measurement* statement: the private profile $\pi(G) = (\beta_1(A), \beta_1(B))$ does not determine Z, so any objective whose only input is the pair of private values is structurally blind to C. Second, and sharper, an *agent-local internalization* statement, which we prove as a dichotomy rather than a single model. A settlement posting has exactly two *faithful* fates, where faithful means the owning agent preserves the posting's recognition endpoints rather than rewriting one of them. Relocated into an owning agent's ledger it becomes a dangling edge, which closes no loop, so under *any* split of the s postings between the agents the decentralized total of the two enlarged private values equals the autarkic total and recovers *none* of C. Held jointly across the boundary it closes a loop only inside a single connected supra-agent object, which is the seam glue itself and whose value is precisely the conceded centralized Z. So the Coasean rejoinder of joint ownership does not recover C decentrally; it reconstructs the center. The one relocation that *does* close an owner-internal loop, moving the posting's far endpoint off the counterparty state in B and onto an existing owner state in A, is not an internalization of the cross-ledger posting at all but a *substitution* that discards the counterparty; it is exactly the globally destructive defection analyzed in the welfare section (Theorem 6.3), which strictly *lowers* Z, not a recovery of C. The coupling is representable by a central commodity price and by no agent acting on its own ledger; that gap, with no third option, is the result. Two welfare corollaries follow: the profile-Pareto criterion permits moves that strictly lower Z, and, conditional on a gains-from-trade hypothesis, a strict profile-Pareto improvement strictly lowers Z. Cooperative game theory can still impute C as a coalition surplus (a Shapley split exists); what is impossible is non-cooperative, ledger-local internalization, and we keep that distinction throughout. All load-bearing combinatorial claims, including the agent-local internalization impossibility, are machine-checked in Lean 4 over Mathlib and are axiom-clean; the homological reading is a standard text argument. That the first Betti number is the *unique* settlement-network valuation forced by a short list of accounting axioms is stated as an open companion target, not claimed here. Pointers are consolidated in a single appendix.

1. Introduction: the profile assumption, stated mathematically

A welfare judgment over economic states is a comparison rule. A large and standard class of welfare objectives shares a single structural assumption about that rule: it reads only the private positions of the agents. Objectives that take a cross-agent quantity as a primitive input, a public good or a Lindahl price written on the seam, or a utility $u_i(x_i, s)$ that depends on the settlement datum s directly, fall outside this class; they are on the centralized side of the boundary we draw

Date: June 26, 2026.

below and can represent Z. The result concerns precisely the class that does not read the seam. Write the private profile of a configuration G with two loci A and B as

$$\pi(G) = (V_A, V_B),$$

the pair of private values. The assumption is the following.

Assumption 1.1 (Profile sufficiency)**.** A welfare objective is a functional W that factors through the private profile, i.e. there is a map $\widehat{W}$ on profiles with

$$W(G) = \widehat{W}(\pi(G)) \qquad \text{for every configuration } G.$$

The content of the assumption is informational, not philosophical. Welfare functionals differ in what they read, prices, allocations, or utilities, but they share an informational basis: each agent enters through quantities computed from that agent's own ledger. Once the individual indices are fixed as ledger-local values, the social objective is a function of the resulting profile. The Pareto preorder declares $G \succeq G'$ exactly when $V_A(G) \geq V_A(G')$ and $V_B(G) \geq V_B(G')$, a statement about π; a Bergson-Samuelson functional aggregates the profile into a scalar. Both factor through π by construction. We show below that there is joint structure no enlargement of the agents' own indices can capture, because it lives in the relation between ledgers rather than in either ledger, so including the seam data is necessary, not optional (Lemma 3.3). Equivalently, any W satisfying Assumption 1.1 is constant on each seam fiber, the configurations obtained by varying only the seam over a fixed pair (A, B), since π is constant there. That single fact is what the central theorem turns on.

We ask a narrow and decidable question. Is there a well-defined economic objective that Assumption 1.1 cannot represent? We answer yes, constructively, with elementary topology. The objective is the count of independent closed recognition loops in the joint system, its first Betti number Z. We prove that joining two loci by a settlement seam creates loops that belong to neither ledger, contributing a strictly positive term C to Z that is invisible to π. Because π does not see C, two configurations can share a profile while differing in Z. Any W satisfying Assumption 1.1 assigns them the same value, so no such W equals Z. We call Z, viewed through π, the *C-blind projection*.

The boundary of the claim, stated first. We draw the boundary before developing the result, because the obstruction is precise and a careless reading would inflate it. The defect C is a closed-form function of the seam count s (Corollary 3.2: $C = \max\{s - 1, 0\}$). A central observer who is allowed to read s, that is, who treats the seam count as an observable commodity, therefore represents Z exactly. We concede this at the outset and claim no impossibility theorem for welfare analysis in general, and in particular nothing Arrovian: a planner with the seam datum has no difficulty. The result lives one level beneath the centralized fix, in two layers we keep distinct.

The first is a measurement statement. If the only admissible input is the pair of private values $\pi = (V_A, V_B)$, the objective Z is not representable, because π is constant on the seam fiber while Z is not (Theorem 5.2). This is a statement about an information type, the private profile, not about welfare economics as a discipline.

The second is the sharp one. The textbook response to a missing cross-agent interaction is the externality remedy: assign the interaction to the agents as property rights and let each internalize its share (Coase, Arrow). We model this faithfully and show it fails for a structural, not incentive, reason. Internalizing a seam posting into the ledger of the agent who owns it attaches a dangling edge, one new posting to a fresh boundary stub, which closes no cycle and so leaves that agent's loop-count unchanged. Consequently, under *any* partition of the s seam postings between the two agents, the decentralized total of the two enlarged private values equals the autarkic total $V_A + V_B$ and recovers none of C (Theorem 4.3). The coupling is representable by a central commodity price and by no decentralized property-rights assignment. That is the result, and it is machine-checked.

We are equally explicit about what cooperative game theory can do, to forestall a misreading of the word "non-attributable." Treat the two agents as a coalition with worth Z and the singletons with worth V_A, V_B; the surplus C is then a perfectly good coalition value and admits a Shapley imputation, which here splits C symmetrically as $C/2$ to each agent. Our claim is not that no accounting number can be assigned to each agent. It is that the assignment cannot be realized *topologically by the agents' own ledgers*: no spanning-forest basis moves a cross-ledger cycle inside A or inside B (Lemma 3.3), and no property-rights internalization closes it locally (Theorem 4.3). Cooperative imputation is a centralized bookkeeping operation; it is on the central side of the boundary we just drew, not the decentralized side.

This is an obstruction, not an impossibility theorem in the Arrovian sense. Arrow's theorem concerns the aggregation of ordinal preference orders and the axioms a social order must violate. Our objective is a cardinal quantity, a loop count, and we do not aggregate orders at all. We bypass the aggregation question by exhibiting a target that the aggregation input type provably cannot track. The distinction matters for what is and is not claimed, and we keep to it throughout.

The result is deliberately generous to the canon. We do not show that classical welfare analysis is wrong on its own terms. We show that it computes a faithful projection of a richer invariant and silently drops one term, the projection defect. Where that term is zero, which is the case of fully separated loci, the canon is exact. The contribution is to name the term, to prove it is a topological invariant rather than a bookkeeping artifact, and to locate exactly the regime in which the projection loses information.

The development carries no human-specific assumption. Nothing below refers to preferences as mental states, to people, or to any particular substrate. The primitives are graphs, cycles, and a comparison cost. We flag this for one reason: the obstruction is a fact about value defined as loop count under settlement, and it applies to any system that banks value by closing loops across a seam. We avoid metaphor in the theorems and reserve interpretive language for clearly marked remarks.

Status conventions. We tag every substantive claim by its epistemic status. Theorem marks a result whose combinatorial core (integer $\beta_1 = E - V + c$ identities, the profile-fiber contradiction, the weighted inequalities, and the derivative-sign and first-order-condition facts) is machine-checked and axiom-clean; where such a result also carries a homological or analytic reading beyond that core, the reading is a standard text argument, flagged as such in the appendix. Theorem-conditional marks a result proved under an explicit named hypothesis. Open marks a program we state but do not prove. These tags are binding; we do not let a conditional or open item drift into the prose as if settled.

2. Setup: ledger, seam, value, comparison cost

2.1. **The recognition ledger and its graph.**

A recognition ledger is a finite directed structure of postings. We retain only its topology. To a ledger we associate a finite multigraph $G = (V, E)$, equivalently a finite 1-dimensional CW complex, the *recognition graph*, whose vertices are recognition states and whose edges are postings between them. Parallel edges are allowed and are essential: two distinct postings between the same pair of states are two edges. We work throughout with finite multigraphs; no continuum is invoked, and all sums are finite.

Definition 2.1 (Value as loop count)**.** For a finite graph $G = (V, E)$ with $\mathrm{c}(G)$ connected components, the first Betti number is

$$\beta_1(G) = |E| - |V| + \mathrm{c}(G).$$

The *value* of G is $V(G) := \beta_1(G)$, the number of independent closed recognition loops, i.e. the rank of the cycle space $H_1(G; Z)$.

Two facts make β_1 the right notion of banked value, and both are standard. First, $\beta_1(G)$ is the dimension of the cycle space and is therefore invariant under any change of spanning tree, that is, under any choice of accounting basis. Second, β_1 is additive over connected components and strictly monotone under the addition of an edge that closes a loop. A single open posting banks no value; a closed loop does. This matches the economic reading that an unsettled promise is not yet capital, while a settled, closed obligation is.

The identification of value with closed-loop count is a modeling choice, and we mark it as the one the results are conditional on. The paper does not assert that every economic value is a Betti number. It asserts a conditional: *wherever* value is carried by closed obligation loops and aggregation reads only each agent's own loops, the joint objective carries a term the aggregation cannot see. Which real settings instantiate the antecedent is an empirical modeling question and is outside the present scope; the mathematical content is the obstruction that follows once the antecedent holds. A reader who rejects the identification rejects the antecedent, not the theorem.

Definition 2.2 (Locus and private profile)**.** A *locus* is a connected recognition graph. For two loci A, B we write $V_A = \beta_1(A)$ and $V_B = \beta_1(B)$ for the private values and

$$\pi(A, B) = (V_A, V_B)$$

for the private profile. Private value is computed from a locus's own graph and is not affected by any posting that lies outside it.

2.2. **The seam.** An economy is two loci joined by settlement. We model settlement as a set of cross-ledger postings, a *seam*, that connect A and B without lying inside either.

Definition 2.3 (Seam glue)**.** Given loci A and B and an integer $s \geq 0$, the *seam glue* $G = A \cup_s B$ is the graph on vertex set $V_A \sqcup V_B$ with edge set $E_A \sqcup E_B$ together with s additional seam edges, each joining a vertex of A to a vertex of B. The integer s is the number of cross-ledger postings.

The seam edges are not edges of A and not edges of B. They are postings of the joint recognizer across the boundary. This is the precise sense in which the seam is invisible to either private ledger: removing the seam returns the disjoint union $A \sqcup B$ and changes neither V_A nor V_B.

We read the small cases as the paper's own definitions of when an economy begins. With $s = 0$ the loci are disjoint, two separate ledgers, no economy. With $s = 1$ a single bridge connects the two components but closes no loop, so $C = 0$; this is an exchange across a seam with no settled cross-ledger loop. With $s = 2$ an obligation and its counter-entry close the first cross-ledger loop; this is the minimal settled trade.

2.3. **The carrying cost of an open seam.** The obstruction in Sections 3–8 is purely topological and uses no comparison cost; the objective Z is an integer loop count and is already cardinal. We isolate one convex cost for the conditional coda of Section 9 alone. The carrying cost of an imbalance σ held open on a seam is

$$2(\cosh \sigma - 1) = 2\, J(e^{\sigma}), \qquad J(x) := \tfrac{1}{2}\Big(x + \tfrac{1}{x}\Big) - 1,$$

strictly convex in σ with a unique minimum at $\sigma = 0$. We introduce J only as the definition that produces this convex cost. Nothing in the main results depends on it.

3. The coupling: super-additivity of value

We now compute the value of a settled economy and isolate the term that the private profile cannot see.

Theorem 3.1 (The defect formula)**.** (Theorem.) *Let A, B be loci and $G = A \cup_s B$ with $s \geq 1$. Then*

$$C(G) \;:=\; \beta_1(A \cup_s B) - \beta_1(A) - \beta_1(B) \;=\; s - 1.$$

Proof. Each locus is connected, so $\beta_1(A) = |E_A| - |V_A| + 1$ and $\beta_1(B) = |E_B| - |V_B| + 1$. For $s \geq 1$ the seam glue is connected, so $\mathrm{c}(G) = 1$ and

$$\beta_1(G) = \big(|E_A| + |E_B| + s\big) - \big(|V_A| + |V_B|\big) + 1.$$

Subtracting,

$$\beta_1(G) - \beta_1(A) - \beta_1(B) = \Big[\,|E_A| + |E_B| + s - |V_A| - |V_B| + 1\Big] - \Big[\,|E_A| + |E_B| - |V_A| - |V_B| + 2\Big] = s - 1. \quad \square$$

Corollary 3.2 (Super-additivity)**.** (Theorem.) *Define the joint objective $Z(G) := \beta_1(A \cup_s B)$. Then*

$$Z(G) = V_A + V_B + C(G), \qquad C(G) = \max\{s-1, 0\},$$

with $C(G) = 0$ for $s \in \{0,1\}$ and $C(G) \geq 1$ for every settled trade $s \geq 2$. For $s = 0$ the loci are disjoint and β_1 is additive over components, so $C = 0$; the formula $C = s - 1$ of Theorem 3.1 is the $s \geq 1$ branch. The objective is strictly super-additive over loci precisely once an economy exists.

The economic content is sharp. The defect C counts cross-ledger loops, obligations whose two halves sit in different ledgers and close through the seam. It is positive the moment a promise is settled, and it is zero in autarky. Two readings deserve emphasis.

The defect is not an artifact of how one keeps the books. The next lemma states this in its sharpest form: C is a cokernel rank, so it lives in the relation between the two ledgers and in neither ledger alone.

Lemma 3.3 (Non-attributability)**.** (Theorem on the integer rank; the homological identification is a standard text argument.) *The disjointness of the edge sets of A and B induces a canonical injection $H_1(A) \oplus H_1(B) \hookrightarrow H_1(G)$, and its cokernel has rank $\max\{s-1, 0\}$:*

$$C \;=\; \operatorname{rank}\, \frac{H_1(G)}{H_1(A) \oplus H_1(B)} \;=\; \operatorname{rank}\, H_1(Q_s),$$

where Q_s is the multigraph obtained by contracting each of A and B to a single vertex and keeping the s seam edges as s parallel edges between the two resulting vertices. For $s \geq 2$, every nonzero class of the quotient is represented only by cycles that traverse the seam, so no choice of spanning forest moves it inside A or inside B.

Proof. No seam edge lies in A or in B, so the cycle spaces of A and B are independent subspaces of the cycle space of G, which gives the injection. Contracting each locus to a point sends $H_1(A)$ and $H_1(B)$ to zero and leaves the quotient, which is the cycle space of the two-vertex multigraph Q_s with s parallel edges; that space has rank $\max\{s-1, 0\}$. A cycle representing a nonzero quotient class is not a combination of cycles internal to A or B, so it must use at least one seam edge under any spanning-forest basis. $\square$

Remark 3.4 (On the word "defect")*.* We call C the projection defect: the rank by which Z exceeds the profile sum on a single configuration. It is not a kernel; the kernel of π is the entire fiber of configurations sharing a profile, not a scalar.

Example 3.5 (Two triangles)**.** Let A and B each be a triangle, so $V_A = V_B = \beta_1 = 1$ and the profile is $(1,1)$ at every seam size. Gluing with s seam edges gives, by Theorem 3.1 and Corollary 3.2,

$$Z = 2,\; 2,\; 3,\; 4 \qquad \text{for} \qquad s = 0, 1, 2, 3,$$

with $C = 0, 0, 1, 2$. The profile is constant while Z rises with the seam. The first cross-ledger loop appears at $s = 2$; it is one of the loops counted by C and missed by π (Figure 1).

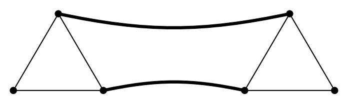

FIGURE 1. The minimal settled economy ($s = 2$): two triangular loci A, B joined by two seam edges (bold). Each locus contributes one internal loop, so the profile is $(1, 1)$. The two seam edges close one additional cross-ledger loop that lies in neither A nor B; it is the defect $C = 1$. Adding seam edges ($s = 0, 1, 2, 3$) leaves the profile at $(1, 1)$ while $Z = 2, 2, 3, 4$.

4. Agent-local internalization is topologically impotent

The non-attributability lemma says no accounting basis localizes C inside an existing ledger. The present section addresses the active remedy a reader will reach for, and is the paper's sharpest claim. Faced with a cross-agent interaction missing from private accounts, the standard prescription (Coase, Arrow [1]) is to assign property rights over the interaction to the agents and have each internalize its share. We model that prescription on the graph and show it cannot recover C, for a structural reason that has nothing to do with incentives.

When an agent "owns" a seam posting, that posting is brought inside the owning ledger. Faithfulness pins the model. A seam posting records a recognition between a fixed owner state $u \in A$ and a fixed counterparty state $v \in B$; internalizing it cannot assert that v is an owner state, because it is not, so the owner can only record the posting against a fresh internal boundary stub standing in for the outside counterparty. It therefore terminates at a boundary stub vertex internal to the owner, so it is a *dangling edge*, one new edge ending at one new degree-one vertex. A dangling edge changes the owner's counts by $|E| \mapsto |E| + 1$ and $|V| \mapsto |V| + 1$ with the component count fixed, so it leaves β_1 unchanged: it closes no loop. The alternative graph move, relocating the far endpoint off v and onto an *existing* owner state $w \in A$, does close an internal loop (since A is connected), but it is not an internalization of this posting: it rewrites the counterparty, discarding the cross-ledger datum that the obligation ran to B. That is a *substitution*, a distinct economic act we analyze on its own terms in Section 6 (it is the destructive defection of Theorem 6.3), not a faithful internalization. We make the faithful dangling-edge operation the model of internalization and read off the consequence.

Definition 4.1 (Internalization as dangling half-edges)**.** For a connected locus A and $a \in \mathbb{N}$, the internalization $\mathrm{int}(A, a)$ attaches a seam postings to A as dangling edges: $|V| \mapsto |V| + a$, $|E| \mapsto |E| + a$, components unchanged.

Lemma 4.2 (Dangling edges bank no value)**.** (THEOREM.) *$\beta_1(\mathrm{int}(A, a)) = \beta_1(A)$ for every $a \in \mathbb{N}$.*

Proof. $\beta_1(\mathrm{int}(A, a)) = (|E_A| + a) - (|V_A| + a) + \mathrm{c}(A) = |E_A| - |V_A| + \mathrm{c}(A) = \beta_1(A)$. □

Theorem 4.3 (Agent-local internalization is impotent)**.** (THEOREM.) *Let A, B be connected loci and let $a + b = s$ be any partition of the s seam postings between them, including the unilateral split $(a, b) = (s, 0)$. Then the decentralized total equals the autarkic total,*

$$\beta_1(\mathrm{int}(A, a)) + \beta_1(\mathrm{int}(B, b)) = V_A + V_B,$$

independent of how the seam is divided. Consequently, for a settled economy $s \geq 2$, centralized Z exceeds the decentralized total by exactly the full coupling,

$$Z(A \cup_s B) - \big[\beta_1(\mathrm{int}(A, a)) + \beta_1(\mathrm{int}(B, b))\big] = s - 1 \geq 1,$$

so decentralized property-rights internalization recovers none of C.

Proof. Both equalities are immediate from Lemma 4.2 and Corollary 3.2: the first because each internalization fixes its locus's β_1, the second by subtracting that total from $Z = V_A + V_B + (s-1)$. The gap is $s - 1 \geq 1$ for $s \geq 2$. □

The economic reading is exact and is the crux of the paper. A central planner who can price the seam count, treating s as a commodity, represents Z with no difficulty; this is the concession of Section 5's preamble. But no assignment of local property rights does, because property rights move postings *into* ledgers as dangling stubs, and a stub is not a loop. The cross-ledger loop is closed by the joint recognizer across the boundary; ownership relocates its endpoint but does not close it.

4.1. **The joint-ownership objection is the theorem.** A referee will press the natural rejoinder: decentralized property rights need not relocate a posting into one ledger as a stub; the two agents can hold the posting *jointly*, and joint ownership closes the cross-ledger loop. This is correct, and it is the theorem, not a refutation. Holding the posting's two recognition endpoints fixed, $u \in A$ and $v \in B$, it has exactly two fates, because a cycle through the posting requires both of its endpoints in a single vertex set and "both endpoints lie in one agent's vertex set" is then a binary property. (The third graph move that a referee might propose, relocating the endpoint v off B onto an existing owner state, does not preserve the posting's endpoints: it is the substitution of Section 6, not a fate of this posting, and Theorem 6.3 shows it lowers Z. With the endpoints fixed, the two fates below are exhaustive.)

Proposition 4.4 (The seam dichotomy)**.** (Theorem.) *For any partition $a + b = s$ of the seam postings, the two agent-local fates are exhausted, and their loop-banking is rigid:*

(a) Relocation into a private ledger. *A posting brought inside one agent becomes a dangling edge; by Lemma 4.2 it closes no cycle, and the two enlarged private values sum to the autarkic total,* $\beta_1(\mathrm{int}(A, a)) + \beta_1(\mathrm{int}(B, b)) = V_A + V_B$.

(b) Joint retention across the boundary. *A posting kept cross-ledger has its endpoints in distinct vertex sets, so any loop it closes spans both vertex sets and lives in a single connected object on the joint vertex set $V_A \sqcup V_B$: the seam glue $A \cup_s B$ itself. Its loop-count is the conceded centralized objective,* $\beta_1(A \cup_s B) = V_A + V_B + (s-1)$.

With the posting's endpoints held fixed there is no third fate, and the gap between the two routes is exactly the coupling $C = s - 1$.

Proof. Immediate from Lemma 4.2, Theorem 4.3, and Corollary 3.2: the two displayed equalities are the two routes, and their difference is $\big(V_A + V_B + (s-1)\big) - \big(V_A + V_B\big) = s - 1$. □

The dichotomy answers the joint-ownership objection rather than dodging it. "Joint ownership that closes the loop" *is* option (b): a single connected object spanning both vertex sets. That object is not a private ledger of either agent; it is the supra-agent book, and it is exactly the centralized representation already conceded in Section 5's preamble. The Coasean rejoinder therefore does not recover C decentrally, it *constructs the center*: closing the cross-ledger loop and forming the joint ledger are one act. The impossibility is consequently not an artifact of the dangling-stub model. Any faithful formalization of an agent computing on its own vertex set lands in option (a), and the only route into option (b) is to build the joint object the result already grants to the center. This is the precise sense in which the obstruction sits beneath the externality fix: the move is available to the center and unavailable to the agents, and the failure is topological, prior to any question of bargaining or transaction cost. We use "center" in an informational sense only: the loop C counts is a property of the joint object on $V_A \sqcup V_B$, so reading it requires the cross-ledger structure that no single agent's ledger holds. We do not claim that a network of bilateral contracts is literally a planner; we claim that the quantity C is computed only from the joint structure and not from either ledger or any partition of postings among them.

Remark 4.5 (What forces β_1, stated as an open target). (Open.) The results take value $= \beta_1$ as a modeling choice (Definition 2.1). One can ask whether β_1 is *forced*: is it the unique graph valuation F that is isomorphism-invariant, additive over disjoint union, normalized to 0 on forests, and incremented by exactly one under a loop-closing edge and by zero under a tree-extending edge? On a genuine graph type carrying the realizability invariant $|E| \geq |V| - \mathrm{c}$, the spanning-tree decomposition gives $F = \beta_1$, so the characterization is true; we do not formalize it here because our count-triple certificate model admits non-realizable triples on which the increment axioms do not pin F, and a clean proof belongs over a graph type that enforces realizability. We state the uniqueness as a well-posed companion target and do not claim it.

5. The central theorem: the C-blind projection

The projection map sends a configuration to its private profile,

$$\pi : G \longmapsto (V_A, V_B).$$

Welfare aggregation, by Assumption 1.1, is any functional that factors through π. The objective Z does not.

Theorem 5.1 (Profile insufficiency)**.** (Theorem.) *There exist configurations G, G' with $\pi(G) = \pi(G')$ and $Z(G) \neq Z(G')$. Consequently Z is not a function of the private profile: there is no $\widehat{W}$ with $Z(G) = \widehat{W}(\pi(G))$ for all G.*

Proof. Fix loci A, B and set $G = A \cup_2 B$, $G' = A \cup_3 B$. Seam postings are not edges of A or of B, so $\beta_1(A)$ and $\beta_1(B)$ are unchanged by s, giving $\pi(G) = \pi(G') = (V_A, V_B)$. By Corollary 3.2, $Z(G) = V_A + V_B + 1$ and $Z(G') = V_A + V_B + 2$, so $Z(G) \neq Z(G')$. If some $\widehat{W}$ satisfied $Z = \widehat{W} \circ \pi$, then $Z(G) = \widehat{W}(\pi(G)) = \widehat{W}(\pi(G')) = Z(G')$, a contradiction. □

Theorem 5.2 (The C-blind projection)**.** (Theorem.) *Fix loci A, B and consider the seam fiber $\{A \cup_s B : s \geq 1\}$. Any functional $W = \widehat{W} \circ \pi$ satisfying Assumption 1.1 is constant on this fiber, because π is. The objective is strictly increasing on the fiber, since $Z(A \cup_s B) = V_A + V_B + \max\{s-1, 0\}$. A constant function cannot equal a strictly increasing one, so no welfare functional that factors through π represents Z. The quantity that π drops is exactly the projection defect $\Delta Z(G) := Z(G) - (V_A + V_B) = C(G)$, strictly positive for every settled trade $s \geq 2$.*

Proof. $\pi(A \cup_s B) = (V_A, V_B)$ is independent of s by Definition 2.2, so W is constant on the fiber. By Corollary 3.2, Z increases by one with each seam edge added beyond the first, hence is strictly increasing in s for $s \geq 1$. A constant function and a strictly increasing one cannot coincide. The dropped term is $Z - (V_A + V_B) = C$, which is ≥ 1 for $s \geq 2$ by Theorem 3.1. □

This is the paper. The classical objective is the projection of Z through π, and that projection drops a non-localizable, strictly positive, topologically invariant term in every settled economy. We now read off the welfare consequences.

6. Dissolution corollaries

Corollary 6.1 (Profile-Pareto permits Z-loss)**.** (Theorem.) *Consider the move $G' = A \cup_3 B \mapsto G = A \cup_2 B$, the withdrawal of one settled cross-ledger loop. Neither private value changes, $\pi(G') = \pi(G)$, so no agent is made worse off and the profile-Pareto criterion does not forbid the move. Yet Z strictly falls, $Z(G') - Z(G) = 1$. The profile-Pareto criterion therefore permits a strict decrease of the joint objective.*

Proof. By Definition 2.2 private values are independent of s, so the move is Pareto-indifferent and a fortiori Pareto-permitted. By Corollary 3.2, $Z(G') = V_A + V_B + 2$ and $Z(G) = V_A + V_B + 1$. □

The unconditional corollary already dissolves the welfare claim. A criterion that cannot distinguish G from G' governs an objective on which they differ. To make the failure vivid, we add one honest economic primitive and obtain a strict statement. We weight loops by load: an internal loop carries weight w_{int}, a cross-ledger loop carries weight w_{seam}, and the load-weighted objective is

$$Z_w(G) = w_{\text{int}}\,(V_A + V_B) + w_{\text{seam}}\,C(G).$$

We assume each agent operates under fixed posting bandwidth, so a posting moved off the seam is re-banked as one internal loop on that agent's own ledger. This re-banking is precisely the *substitution* flagged in Section 4: the agent withdraws the cross-ledger posting, discarding its B-counterparty, and spends the freed bandwidth closing a loop among its own existing states. It is not the faithful internalization of Theorem 4.3 (which closes no loop); it is the loop-closing alternative, and the point of the next theorem is that, far from recovering C, it strictly lowers the joint objective. The two sections thus model one operation each, consistently: faithful internalization (Section 4) banks no loop, and substitution (here) banks an internal loop worth strictly less than the seam loop it destroys.

Assumption 6.2 (Gains from trade)**.** A closed cross-ledger loop banks more load than the bandwidth it frees can re-bank internally: $w_{\text{seam}} > 2\,w_{\text{int}} > 0$.

Theorem 6.3 (Strict Pareto improvement lowers Z)**.** (Theorem-conditional on Assumption 6.2.) *Suppose both agents symmetrically withdraw from the seam, dissolving one cross-ledger loop and each re-banking the freed bandwidth as one internal loop. Then each private weighted value strictly rises, by w_{int}, so the move is a strict Pareto improvement; yet the load-weighted objective strictly falls,*

$$\Delta Z_w = 2\,w_{\text{int}} - w_{\text{seam}} < 0.$$

A strict Pareto improvement is globally destructive.

Proof. Dissolving one seam loop decreases C by 1 and, by bandwidth conservation, increases V_A and V_B by 1 each. Private weighted values rise by $w_{\text{int}} > 0$, so both agents strictly gain. The weighted objective changes by $w_{\text{int}} \cdot 2 - w_{\text{seam}} \cdot 1$, which is negative under Assumption 6.2. □

Remark 6.4. The conditional layer needs an economic input and we mark it as such. The unconditional content, Theorem 5.2 and Corollary 6.1, needs none; it is the coupling arithmetic. The reading is that gains from trade are exactly the structure that makes private optimization globally destructive. We state it carefully because the strict version, and only the strict version, depends on $w_{\text{seam}} > 2w_{\text{int}}$. The two thresholds that appear in this paper are for two different moves and are not in tension. The symmetric defection of Theorem 6.3 frees two postings, one per agent, and so requires $w_{\text{seam}} > 2w_{\text{int}}$; the marginal interior optimum of Section 9 concerns a single posting at the margin and requires only the weaker $w_{\text{seam}} > w_{\text{int}}$.

We summarize. Profile-local welfare aggregation is the C-blind projection of Z: exact in autarky, where $C = 0$, and systematically dropping the settlement term C in every settled economy. A central observer who reads the seam count recovers C; no agent-local property-rights assignment does (Theorem 4.3). We invoke no impossibility theorem for the aggregation of orders. The claim is sharp and bounded: the objective is cardinal, the profile-local and agent-local input types cannot represent it, and the centralized commodity fix is conceded.

7. The canon as the C-blind projection: an open program

If the classical objective is the projection π, then each classical result should be recoverable as something the projection computes correctly. We state this as a program, not as a theorem. The table below lists targets. Each row is a conjecture about how a known result sits inside the projection picture, tagged by status. We claim none of these recoveries here. We claim only that the projection theorem makes the program well-posed, by identifying precisely the term, C, whose vanishing returns the classical regime.

Classical result	**Position in the projection picture (conjectured)**	**Status**
Smith, gains from exchange	The cross-ledger loop is the banked surplus; trade creates loops the autarkic profile cannot hold. This is the $C > 0$ regime read forward.	Open
Marginalist value	Price as the J-readout of a local ratio is the gradient of the profile shadow $V_A + V_B$, the projection's first-order data.	Open
Marshallian surplus	Consumer plus producer surplus is the profile-local accounting of loop closure, exact when the seam term is held fixed.	Open
Arrow-Debreu equilibrium	A fixed point of profile-local optimization, i.e. optimization of the projection; equilibrium is the C-held-constant stationary point. The case $C = 0$ is autarky, not Arrow-Debreu.	Open

Remark 7.1. The projection theorem does not by itself recover the canon as the image of π; it makes the recovery a definite task by naming the dropped term. The case $C = 0$ is the disjoint, autarkic configuration; it is not the Arrow-Debreu world, which lives at a stationary point of the projection with the seam present but its variation suppressed. Conflating $C = 0$ with general equilibrium would be an error.

8. Scope, falsifiers, and what is not claimed

8.1. **What is claimed.** The projection defect formula $C = s - 1$, the super-additivity $Z = V_A + V_B + C$, profile insufficiency, the corollary that the profile-Pareto criterion permits Z-loss, and the agent-local internalization impotence theorem (Theorem 4.3: under any property-rights split the decentralized total misses the full coupling $s - 1$) are theorems. The strict statement that a profile-Pareto improvement lowers Z is a theorem conditional on Assumption 6.2. The objective Z is a topological invariant of the settled economy; it is representable by a central seam-count commodity and by no agent-local property-rights assignment.

8.2. **What is not claimed.**

- We do not claim that the Smith-to-Arrow-Debreu canon is formally recovered as the image of π. Section 7 is an open program, tagged row by row, not a proved recovery.
- We do not claim an impossibility theorem for welfare analysis in general. A central observer who reads the seam count s represents Z exactly, since Z is a closed-form function of s; we concede this throughout. The obstruction is to profile-local and agent-local representation, not to centralized representation.
- We do not claim that C admits no imputation. Cooperative game theory imputes it: the coalition worth Z against singleton worths V_A, V_B gives a Shapley split $C/2, C/2$. Our impossibility is specifically topological and decentralized, that no spanning-forest basis localizes the cross-ledger cycle and no property-rights internalization closes it; cooperative imputation is a centralized operation and is conceded.
- We do not claim Arrow's impossibility theorem, nor any variant of it. Our objective is a cardinal loop count, not an aggregation of orders, so Arrow is bypassed, not recovered. The case $C = 0$ is autarky, not general equilibrium.

- We make no claim about mechanism-design impossibility results. Incentive compatibility, individual rationality, and budget balance are not modeled here, and the obstruction proved in this paper is informational non-representability, not strategic non-implementability. Myerson-Satterthwaite-type obstructions are a separate question.
- We do not claim empirical confirmation. The construction predicts a structural obstruction, not a fitted number, and we advance no statistical regularity in its support.
- We do not claim a parameter-free optimal economy. The coda of Section 9 is conditional on its modeling inputs and appears only as a conditional statement.
- We make no claim about human cognition or about any privileged substrate. The interpretive language of a single recognizer is metaphor; the literal content is that the proofs use no human-specific assumption.

8.3. **Falsifiers.**

- A welfare functional taking only the private profile as input that provably tracks Z across configurations with equal profile and unequal Z. This refutes Theorem 5.2. Given the proof, the live target is the identification: exhibit a classical objective that does not factor through π, contradicting Assumption 1.1 as a description of the canon.
- A settled economy with $s \geq 2$ and $C = 0$, that is, a settled trade that opens no cross-ledger loop. This refutes the defect formula. Given the proof, the live target is the modeling map from real ledgers to the recognition graph.
- A setting with gains from trade, in the sense of Assumption 6.2, in which the load-weighted net objective has no interior optimum. This refutes the conditional coda below.
- A model of agent-local internalization, faithful to the property-rights remedy, under which the two enlarged private loop-counts sum to more than the autarkic total. This refutes Theorem 4.3. Given the proof, the live target is the modeling map: argue that owning a seam posting should close a loop inside the owner's ledger rather than attach a dangling stub.

9. Conditional coda: a unique interior Z-optimum

We close with one positive statement, marked conditional, that we do not advance as a headline. Holding the seam open carries a cost. The imbalance on an open seam, parameterized by σ, carries the convex cost $2(\cosh\sigma - 1)$ introduced in Section 2.3. We relax the integer seam count to a continuous loop-load $q \geq 0$ (the realized seam count is the nearest feasible integer to the optimum below) and write the load-weighted net objective

$$Z_{\text{net}}(q) = (w_{\text{seam}} - w_{\text{int}})\, q \;-\; 2\big(\cosh(kq) - 1\big),$$

where $k > 0$ is the per-promise imbalance scale.

Theorem 9.1 (Conditional interior optimum)**.** (Theorem-conditional on the modeling inputs $w_{\text{int}}, w_{\text{seam}}, k$.) *$Z_{\text{net}}$ is strictly concave. Under the marginal seam-premium condition $w_{\text{seam}} > w_{\text{int}}$ (weaker than the gains-from-trade Assumption 6.2), Z_{net} has a unique interior maximizer $q^* > 0$ characterized by the seam-funding margin*

$$w_{\text{seam}} - w_{\text{int}} = 2k\,\sinh(k\,q^*),$$

the marginal recognition gain of one more closed promise set equal to its marginal convex carrying cost. If $w_{\text{seam}} \leq w_{\text{int}}$, the optimum is at the boundary $q^* = 0$, autarky.

Proof. The term $2(\cosh(kq) - 1)$ is strictly convex in q because cosh is, and the linear term is affine, so Z_{net} is strictly concave. Its derivative is $w_{\text{seam}} - w_{\text{int}} - 2k\sinh(kq)$, which is strictly decreasing, equals $w_{\text{seam}} - w_{\text{int}}$ at $q = 0$, and tends to $-\infty$. It has a positive root iff the value at 0 is positive, i.e. iff $w_{\text{seam}} > w_{\text{int}}$; the root is unique by strict monotonicity. □

This is the replacement objective in conditional form. If gains from trade, then a unique interior Z-optimum exists and is governed by a computable first-order condition. We state it as a coda and nothing more. The promotion of its inputs from modeling assumptions to forced quantities is the task of a separate paper and is out of scope here.

Appendix A. Formalization

The combinatorial core of every result tagged Theorem or Theorem-conditional is machine-checked; the body carries no formalization names by design. The certificate is a single self-contained Lean 4 file that imports only the Mathlib library, so it builds from scratch and is reproducible independently of any larger development. Every statement below is axiom-clean: it reduces to the ambient logic with the standard classical foundations (propositional extensionality, the axiom of choice, and quotient soundness) and no domain-specific axiom. What is formalized is exactly the integer first-Betti identities ($\beta_1 = E - V + c$), the profile-fiber contradiction, the Pareto and load-weighted inequalities, and the derivative-sign and first-order-condition facts of the coda. What is *not* formalized, and is instead a standard text argument, is the homological packaging (the groups H_1, the canonical injection and its cokernel, the contraction Q_s) and the analytic packaging (strict concavity and global maximality as named predicates); the file proves the integer and derivative facts that these standard arguments rest on.

Body result	**Checked statement**
Defect formula, Theorem 3.1	$C = \beta_1(A \cup_s B) - \beta_1(A) - \beta_1(B) = s - 1$, with the settled-promise instance $C = 1$ at $s = 2$ and the disjoint instance $C = 0$ at $s \leq 1$, on the canonical first-Betti definition of value.
Super-additivity, Corollary 3.2	$Z = V_A + V_B + C$, with the settling-tick gain that adds 1 to Z and 0 to each private value.
Profile insufficiency, Theorem 5.1	There exist G, G' with $\pi(G) = \pi(G')$ and $Z(G) \neq Z(G')$; Z is not a function of the profile, on a concrete two-locus instance.
C-blind projection, Theorem 5.2	$\Delta Z = Z - (V_A + V_B) = C = s - 1$, and constancy of any profile functional on the seam fiber against the strict monotonicity of Z.
Pareto permits Z-loss, Corollary 6.1	The move $s : 3 \to 2$ leaves the profile unchanged and strictly lowers Z.
Strict Pareto lowers Z, Theorem 6.3	Under the named gains-from-trade hypothesis $w_{\text{seam}} > 2w_{\text{int}}$, symmetric defection strictly raises both private weighted values and strictly lowers weighted Z.
Conditional optimum, Theorem 9.1	The derivative formula for Z_{net}, the first-order condition (derivative $= 0$ at q^*), the strict monotonicity of the derivative, and its sign on each side of q^*, all under $w_{\text{seam}} > w_{\text{int}}$. (Strict concavity and global maximality are the standard text reading of these facts.)
Non-attributability, Lemma 3.3	The integer excess $\beta_1(A \cup_s B) - (\beta_1(A) + \beta_1(B)) = \max\{s - 1, 0\}$, on the combinatorial first-Betti definition. Its cokernel reading (the rank of the cokernel of $H_1(A) \oplus H_1(B) \hookrightarrow H_1(G)$) is standard cycle-space algebra and is not separately formalized.
Dangling edges bank no value, Lemma 4.2	$\beta_1(\text{int}(A, a)) = \beta_1(A)$ for every a, on the count model of internalization as dangling half-edges.
Agent-local impotence, Theorem 4.3	For every split $a + b = s$, $\beta_1(\text{int}(A, a)) + \beta_1(\text{int}(B, b)) = V_A + V_B$, and the centralized-minus-decentralized gap is exactly $s - 1$, which is ≥ 1 for every settled trade $s \geq 2$. The full coupling is recovered by no decentralized property-rights assignment.

All of the above is contained in one self-contained Lean 4 file (Mathlib only), which builds and is axiom-clean as stated. The file is posted to the public repository https://github.com/jonwashburn/c-blind-projection at commit `c29191d75dab277fa41400bb8e20dd7e4f00a877`. That commit pins the Lean toolchain (`leanprover/lean4:v4.27.0-rc1`) and the Mathlib revision (`d7ea5678e6d426e87e9b4a65a48143c4874dc501`) at which the file is checked, and it has been built from a clean checkout, so a third party can reproduce the build (`lake exe cache get && lake build`) from that commit with no further dependency. The body of this paper is standalone mathematics; nothing in Sections 1 through 9 depends on consulting the formalization, which serves only to certify that the load-bearing claims are not informal.

References

[1] K. J. Arrow, *The organization of economic activity: issues pertinent to the choice of market versus nonmarket allocation*, in: The Analysis and Evaluation of Public Expenditure: The PPB System, vol. 1, U.S. Joint Economic Committee, 1969, pp. 59–73.

Recognition Physics Institute, Austin, Texas
Email address: jon@recognitionphysics.org