# When a Market Solves Itself

The Homology of Free Computation in Exchange Economies




June 2026

**Abstract.** A market is a graph: goods are vertices, a substitution is an edge. We show that one topological number — the dimension of the graph's space of independent loops — decides when an equilibrium computes itself for free: a forest market collapses its partition function to the mean field in one closed-form step, and each loop is a correction. We prove that four separate descriptions of a clearing price — order-theoretic, statistical-mechanical, homological, and analytic — are four faces of one collapse law, and we prove what is *not* true: these four "free" predicates are strictly nested, not identical, since a gross-substitutes economy is always confluent regardless of its loop number. We build the executable ledger as a runnable certified classifier, construct a topology-faithful economy whose unique equilibrium realizes the analytic face, and run a sealed pre-registered program on the real inter-country input–output panel. There the homology earns nothing as a macroeconomic *predictor* — a clean negative we report without inflation — but earns its keep as a *computation*: a solver that handles the certified-free part in closed form reproduces the equilibrium to machine precision a median 8.5× faster than a strong baseline. Every load-bearing statement is machine-checked.

# 1 Introduction

A market reaches a price at which it clears. We are used to treating that fact as the output of an optimization: somewhere a system of equations is solved, and the clearing price is the solution. But there is an older and stranger intuition hiding underneath. Sometimes the price is not *computed* so much as *revealed* — the structure of the market is already such that the answer is forced, and no work remains to be done. A two-good barter clears itself; a tangle of mutually substitutable commodities does not. *When, exactly, does a market do the computation for you?*

This paper answers that question with a number you can read off a picture. Draw the market as a graph: one vertex per good, one edge for each pair of goods that substitute for one another. That graph has a spanning forest and, hanging off it, some extra edges — each extra edge closes a loop. The count of independent loops is the graph's first Betti number $\beta_1$. Our central claim is that $\beta_1$ is the exact ledger of free computation: *the market's equilibrium is computed for free — its partition function equals its mean-field value in a single closed-form evaluation — if and only if the substitution graph is a forest,* $\beta_1 = 0$. Every loop is a correction that must be paid; the obstruction is the boundary of the free region.

What makes this more than a slogan is that the same boundary appears under four lenses that economics, order theory, statistical physics, and algebraic topology developed independently, and the four turn out to name one object. Read order-theoretically, market clearing is a closure operator — monotone, inflationary, idempotent — and "free" means its least and greatest fixed points have collapsed onto a single equilibrium. Read statistically, the market is an Ising system whose partition function $\mathcal{Z}$ factors as a mean-field term times a sum over even subgraphs; "free" means that loop sum is trivial. Read homologically, the substitution graph splits over $\mathbb{F}_2$ into a cut space and a cycle space; "free" means the cycle space is a point. Read analytically, prices settle by tâtonnement under gross substitutes; "free" means the equilibrium is unique. We prove all four faces and prove that they are organized by a single abstract law—a fixed point is free when lfp = gfp (Theorems 3.1, 3.2 and 4.2).

The result also corrects a tempting overstatement, and we think the correction is the more interesting half. It is natural to guess that the four "free" predicates are one number wearing four costumes. They are not. A connected gross-substitutes economy has a *unique* equilibrium whether or not its substitution graph has loops — uniqueness follows from monotonicity alone, not from acyclicity. So the analytic notion of freedom (a unique clearing price) is strictly coarser than the topological notion (an acyclic substitution graph), which is in turn strictly coarser than the statistical notion (a trivial loop sum can occur even with loops, when the weights cancel). The honest picture is a strict hierarchy of three distinct freedoms unified only by the abstract collapse law, not a single invariant (Theorem 3.3). Naming the four faces as the *same* number would be a false economy; naming them as instances of one *law* is the true one.

Why should this matter outside a seminar? Because it turns "how hard is it to find this market's equilibrium" into a property of the market's shape that you can see, compute, and check. The cost of clearing becomes a topological invariant rather than a black box; the boundary between self-solving and correction-laden markets becomes a kernel-checked theorem rather than an intuition; and the resulting classifier is something you can run on a real economy. We do run it: on the inter-country input–output accounts, through the verified artifact itself, with a pre-registered null. The empirical chapter is deliberately modest in its claim and exact about why — a methodological point we make in the open (Section 6).

**Where the leverage comes from.** One ingredient we did not have to build. The analytic face of the result — that a strict-gross-substitutes exchange economy has a unique equilibrium — is a




Figure 1: The whole phenomenon, before any machinery. A market whose substitution graph is a forest computes its own equilibrium: its partition function equals the mean-field value $2^n$ in a single closed-form evaluation. Adding one coupling that closes a loop ($\beta_1$ rises from 0 to 1) introduces a loop correction $\ell$ that must be summed. The number of independent loops $\beta_1$ is the ledger of how much computation the market does *not* do for itself.

theorem the formal mathematics library already contained in full generality, sitting one search away from where the two-good special case had been used. Recognizing that the general statement was already available, rather than re-deriving it, is what let the analytic face close immediately once the right economy was constructed (Section 5). It is a small instance of the institute's standing thesis: when the structure is correctly identified, the work has often already been done for you.

**Roadmap.** In Section 2 we make the market-as-graph precise and define the free/costly dichotomy. In Section 3 we state the collapse law and its four instantiations, then prove the strict freedom hierarchy. In Section 4 we build the executable homological ledger and prove the runnable classifier equal to the certified statistical-mechanical object. In Section 5 we construct a topology-faithful exchange economy and realize the analytic face through equilibrium uniqueness. In Section 6 we run the classifier on real data. Section 7 reports the formalization and its axiom budget; Section 8 states the realization claim; Section 9 draws the implications.

## 2 The Market as a Substitution Graph

Start with the simplest object that still has the phenomenon in it. A market is a finite set of goods together with a list of pairwise substitution couplings; nothing else is needed to see free computation appear and disappear.

**Definition 2.1** (Market)**.** A *market* $M$ is a pair $(V, E)$ where $V = \{0, \ldots, n-1\}$ is a set of $n$ goods and $E$ is a finite list of *couplings* $(i, j, w)$ with $i, j \in V$ and a signed weight $w$. The underlying *substitution graph* $G(M)$ forgets the weights, retaining the edge $\{i, j\}$ for each coupling.

The picture to keep is the one in Figure 1: a chain of goods is a forest; a triangle of goods has one loop. We will see that this single difference — a loop or no loop — is the whole story.

**The two spaces of a graph.** Over the two-element field $\mathbb{F}_2$, the edge set of $G(M)$ carries a boundary map $\partial$ sending each edge to the sum of its endpoints. Its kernel is the *cycle space* (the even subgraphs, where every vertex has even degree); the row space of $\partial$ is the *cut space*. These are the first homology and the coboundaries of the graph, and their dimensions obey the rank–nullity

identity (Diestel, 2017; Oxley, 2011)

$$\underbrace{\beta_1(G)}_{\text{dim cycle space}} \quad + \quad \underbrace{\operatorname{rank}(G)}_{\text{dim cut space}} \quad = \quad |E|, \qquad \beta_1(G) = |E| - |V| + c(G), \tag{1}$$

where $c(G)$ is the number of connected components. The first Betti number $\beta_1$ — equivalently the circuit rank or nullity — is the number of independent loops, and the cycle space has exactly $2^{\beta_1}$ elements.

**The free/costly dichotomy.** Reading the couplings as Ising bond strengths makes $M$ a statistical system. Its partition function and mean field are

$$\mathcal{Z}(M) \;=\; \sum_{\sigma \in \{\pm 1\}^V} \; \prod_{(i,j,w) \in E} (1 + w\,\sigma_i \sigma_j), \qquad \mathrm{MF}(M) \;=\; 2^{|V|}. \tag{2}$$

**Definition 2.2** (Free market)**.** A market is *free* when its loop content $\ell(M)$, the sum of bond weights over even subgraphs, equals 1 — equivalently, when $\mathcal{Z}(M) = \mathrm{MF}(M)$. Otherwise it is *costly*.

A free market is one whose competitive equilibrium is mean-field-exact: the loop sum collapses, and the partition function is computed in closed form with no correction. The rest of the paper is the claim that this single dichotomy is governed — under four lenses and one law — by the topology in Equation (1).

# 3 Four Faces of a Clearing Price

Here is the surprise. The word "free" in Definition 2.2 was chosen by a statistical physicist; an order theorist, an economist, and a topologist each have their own "free," invented for their own reasons, and the four coincide where it counts. We first state the law that organizes them, then the four instantiations, then the precise sense in which they agree — and the precise sense in which they do *not*.

## 3.1 The collapse law

**Picture.** A settling process pushes a state upward until it stops. If you start at the bottom and at the top and both runs land in the same place, the process has exactly one resting state — and "finding" it was never a search, because there was nowhere else to go.

**Theorem 3.1** (Collapse law)**.** *Let $T\colon L \to L$ be a monotone self-map of a complete lattice (a settling dynamic). The following are equivalent:*

$$\operatorname{lfp} T = \operatorname{gfp} T \quad \Longleftrightarrow \quad |\mathrm{Fix}(T)| = 1 \quad \Longleftrightarrow \quad \textit{the proper adjustment relation of } T \textit{ is confluent.} \tag{3}$$

**Meaning.** This is the order-theoretic skeleton shared by all four faces. By Knaster–Tarski (Echenique, 2005; Knaster, 1928; Tarski, 1955), the fixed points of a monotone map on a complete lattice themselves form a complete lattice with least element $\operatorname{lfp} T$ and greatest element $\operatorname{gfp} T$; that lattice is a single point exactly when its top and bottom agree, and a rewriting system is confluent exactly when its reflexive-transitive closure has unique normal forms (Baader and Nipkow, 1998; Newman, 1942). "Free computation" is the operational reading: when lfp = gfp, the idempotent closure that finds the equilibrium does no work beyond a single application, because there is a unique fixed point and every trajectory is already there. The proof is the lattice argument; it is axiom-free.

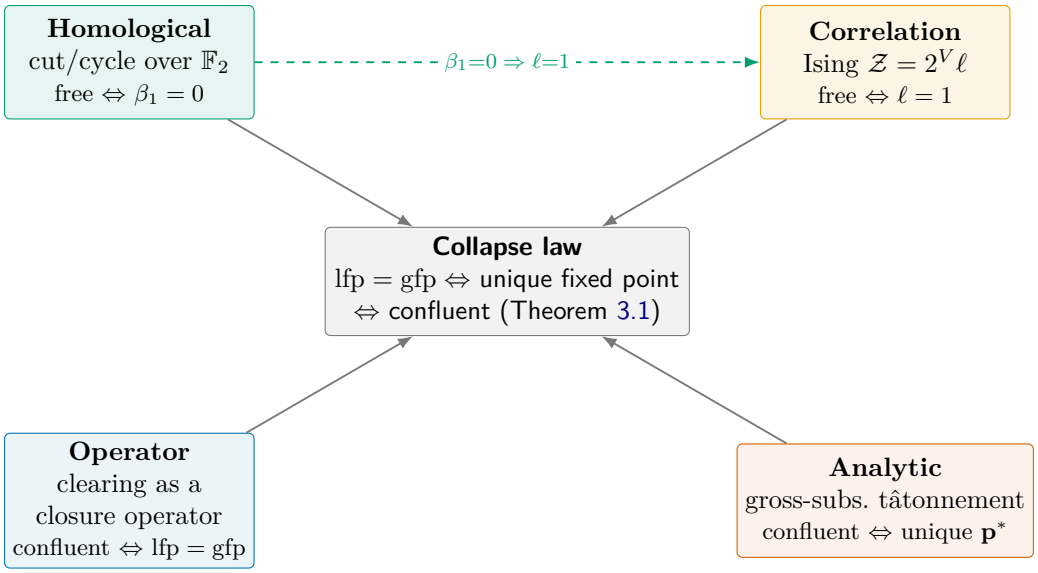



Figure 2: Four constructions, one law. Each face is a settling dynamic whose freedom is the collapse lfp = gfp of Theorem 3.1; the comparisons to the center are structure-preserving. The dashed arrow is the only direct implication *between* faces that holds unconditionally (Theorem 3.3): an acyclic substitution graph forces a trivial loop sum, but not conversely, and the analytic face is governed by the law without being governed by $\beta_1$.

### 3.2 The four instantiations

Figure 2 is the map of the territory: four constructions, one law, related by structure-preserving comparisons.

**Operator.** Market clearing is a closure operator: monotone (raising a price weakly raises every excess demand), inflationary, and idempotent. Confluence of the clearing dynamic is the law of Theorem 3.1; this face is built over a finite Boolean carrier and is axiom-free.

**Correlation.** The partition function factors. This is the van der Waerden / high-temperature expansion (Duminil-Copin, 2017; van der Waerden, 1941), here stated as an exact identity rather than an approximation.

**Homological.** The cut and cycle spaces of Equation (1); the free region is $\beta_1 = 0$.

**Analytic.** A real exchange economy whose excess demand satisfies gross substitutes; its equilibrium is governed by the Arrow–Block–Hurwicz theory (Arrow and Hahn, 1971; Arrow et al., 1959).

The correlation face is the cleanest place to see the loop content do its work.

**Theorem 3.2** (Partition factorization)**.** *For every market with couplings in a commutative ring,*

$$\mathcal{Z}(M) \;=\; 2^{|V|} \cdot \ell(M), \qquad \ell(M) \;=\; \sum_{\substack{S \subseteq E \\ S \text{ even}}} \prod_{e \in S} w_e, \tag{4}$$

*where "$S$ even" means every vertex has even degree in $S$. Consequently $M$ is free $\iff$ $\ell(M) = 1 \iff \mathcal{Z}(M) = 2^{|V|}$.*

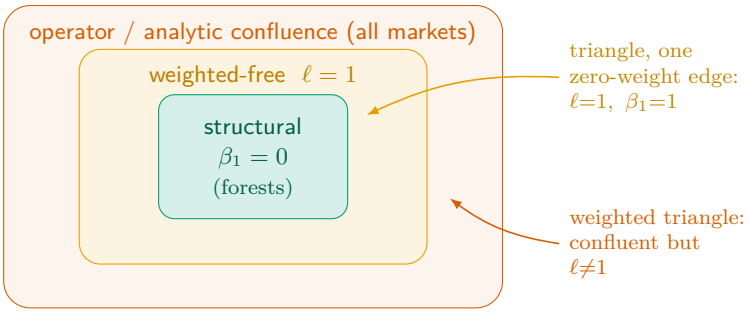



Figure 3: The honest hierarchy of Theorem 3.3. Three distinct notions of a "free" market, strictly nested. Acyclicity ($\beta_1 = 0$) is the finest; a trivial weighted loop sum is coarser (it tolerates loops whose weights cancel); a unique clearing price is coarsest and universal among gross-substitutes economies. The faces share the collapse law, not a single number.

**Meaning.** The mean-field value $2^{|V|}$ is what the market would compute if its goods were independent; the loop content is the sum of *exactly* the corrections that interaction forces, indexed by the even subgraphs — the elements of the cycle space (Chertkov and Chernyak, 2006; Watanabe and Fukumizu, 2012). A forest has no nonempty even subgraph, so $\ell = 1$ and the market is free. The proof is a character-sum (Fourier-over-$\mathbb{F}_2$) expansion of the product in Equation (2); carried out over an abstract commutative ring of weights, it uses *no* axioms — not even propositional extensionality — a sharper footprint than the integer specialization, which incurs the single axiom `propext` through ring-distributivity rewriting. The even-subgraph indexing is precisely the cycle space of Equation (1); this is the bridge from the correlation face to the homological face.

### 3.3 The strict hierarchy: what is *not* true

Now the correction. The four "free" predicates are easy to conflate. We prove they are strictly nested, with witnesses at each gap.

**Theorem 3.3** (Freedom hierarchy)**.** *For markets with positive (substitution) couplings,*

$$\underbrace{\{\,\beta_1 = 0\,\}}_{\textit{structural}} \subsetneq \underbrace{\{\,\ell = 1\,\}}_{\textit{weighted}} \subsetneq \underbrace{\{\,\textit{clearing is confluent}\,\}}_{\textit{operator / analytic}} = \textit{every market.} \tag{5}$$

*Each inclusion is strict: a triangle with one zero-weight edge is weighted-free but structurally looped ($\ell = 1$, $\beta_1 = 1$); and the operator/analytic confluence holds universally — a connected gross-substitutes economy has a unique equilibrium whatever its loop number.*

**Meaning.** This is the load-bearing honesty of the paper. The temptation is to declare that "the four faces are the same $\beta_1$." They are not. The analytic freedom — a unique clearing price — comes from monotonicity and is *universal* among gross-substitutes economies; it does not see $\beta_1$ at all. The weighted freedom — a trivial loop sum — can hold even when the graph is looped, if the weights cancel. Only the structural freedom is $\beta_1 = 0$. What unifies the four is the abstract law of Theorem 3.1, not a common invariant. To collapse them onto one number would be to claim a coincidence the mathematics refutes; Figure 3 is the corrected picture.

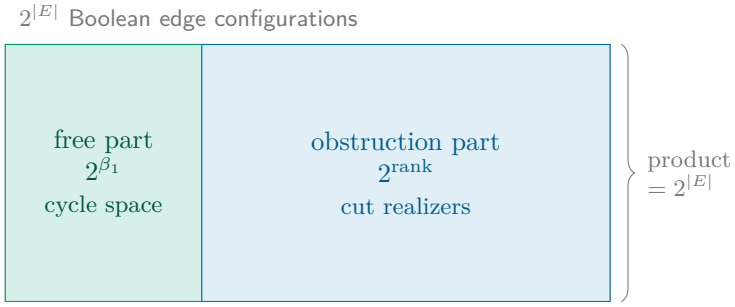



Figure 4: The tiling law of Theorem 4.2: $2^{\beta_1} \cdot 2^{\text{rank}} = 2^{|E|}$, rank–nullity rendered as a partition of the edge cube into a free (cycle-space) part and an obstruction (cut-space) part. Both counts are produced by runnable enumerations, and the identity is a verified theorem about those enumerations, not a numerical coincidence on examples.

# 4 The Executable Homological Ledger

A theorem that classifies markets should be something you can run. We package the homological face as a finite, computable readout and then prove the runnable object equals the certified statistical-mechanical one — so that running the classifier is not a shadow of the mathematics but the mathematics itself.

## 4.1 The ledger and its tiling law

**Definition 4.1** (Homology ledger)**.** The *ledger* of a market records, as $\mathbb{F}_2$-vector-space cardinals, both faces of Equation (1):

$$\text{freeCount}(M) = 2^{\beta_1}, \quad \text{obstructionCount}(M) = 2^{\text{rank}}, \quad \text{freeRank} = \beta_1, \quad \text{obstructionRank} = \text{rank}\,. \tag{6}$$

**Theorem 4.2** (Tiling law)**.** *For every market, the kernel count and image count tile the Boolean cube on the edges:*

$$\text{freeCount}(M) \cdot \text{obstructionCount}(M) \;=\; 2^{|E|}, \qquad \beta_1 + \text{rank} = |E|. \tag{7}$$

*Both factors are the cardinalities of executable enumerations: the cut-orthogonal even-mask set (the kernel) and the dual-evaluation image (the cut realizers). The identity is the rank–nullity theorem made runnable.*

**Meaning.** Figure 4 is the law as a picture: the $2^{|E|}$ Boolean edge configurations factor cleanly into a free part of size $2^{\beta_1}$ and an obstruction part of size $2^{\text{rank}}$. The non-trivial content is that the *executable* image of the dual-evaluation map (the number of distinct read-back columns of the cut code) equals $2^{\text{rank}}$ — a runnable row-rank-equals-column-rank statement, proved by a separator-certified image-doubling invariant folded across the code. The kernel/image surface is axiom-free; the market-level wrappers that lift it to parity filters carry only `propext`.

## 4.2 The classifier equals the certified object

The ledger gives a runnable verdict; Theorem 3.2 gives a certified one. A classifier is only trustworthy if these agree — not on a handful of examples, but as a theorem.




ordinary coupling
endpoint slots at $i$: 1
toggles at $i$: 1
**parities agree**

self-loop $i = j$
endpoint slots at $i$: 2
toggles at $i$: 1
**parities diverge**

Figure 5: Why the faithfulness theorem (Theorem 4.3) carries a no-self-loop hypothesis. The certified parity counts endpoint slots; the evaluator's parity counts touching edges. A self-loop contributes two slots (even) but one touch (odd), so the two notions disagree precisely on loops. Substitution graphs between distinct goods have none, and the empirical data satisfies the hypothesis.

**Theorem 4.3** (Evaluator faithfulness)**.** *Let the* evaluator *compute loop content and even-mask count directly by enumerating Boolean edge masks and testing each for even vertex incidence. For every market with no self-loops ($i \neq j$ on every coupling),*

$$\ell_{\text{eval}}(M) = \ell(M), \qquad \text{evenMaskCount}_{\text{eval}}(M) = 2^{\beta_1}. \tag{8}$$

*The no-self-loop hypothesis is necessary: on a self-loop the two even-parity notions diverge, and a two-good witness with a self-loop separates them.*

**Meaning.** This is the anti-shadow guarantee. The reconciliation has four moving parts, each a small bridge: the two mask enumerations agree (an "expand = flatten" identity on Boolean cubes); the evaluator's toggle-parity over incident edges equals the certified incidence-count parity *exactly when there are no self-loops* (a loop contributes two endpoint slots but a single toggle, so the parities part company on a loop — Figure 5); the weight products coincide; and the fold-based sum equals the structural sum. Because the equivalence is symbolic rather than a finite check, it certifies instances far beyond any enumerable size — in particular the real input–output market of Section 6, whose direct evaluation exceeds the kernel's reduction budget. The self-loop caveat is not a blemish; it is a genuine mathematical refinement that the evaluator's design surfaced, and stating it is what makes the equality true.

## 5 The Topology-Faithful Analytic Economy

The analytic face needs a real economy — prices in $\mathbb{R}$, excess demands, a tâtonnement that settles. The difficulty is fidelity. A naive construction makes the economy depend only on *how many* couplings a market has, not *which* goods they connect; two structurally different markets with the same edge count would collapse to the same economy, and the analytic face would no longer be about the market in front of us. We resolve this with one construction.

**Definition 5.1** (Topology-faithful Metzler economy)**.** For a market with $n \geq 2$ goods and validated couplings $E$, the excess demand at positive prices $\mathbf{p}$ is

$$z_i(\mathbf{p}) \;=\; \underbrace{\left(\frac{\sum_k p_k}{p_i} - n\right)}_{\text{unit baseline}} + \underbrace{\sum_{(a,b,w)\in E} \left( [a{=}i]\,|w|\left(\tfrac{p_b}{p_i} - 1\right) + [b{=}i]\,|w|\left(\tfrac{p_a}{p_i} - 1\right)\right)}_{\text{coupling-keyed Metzler term}}. \tag{9}$$

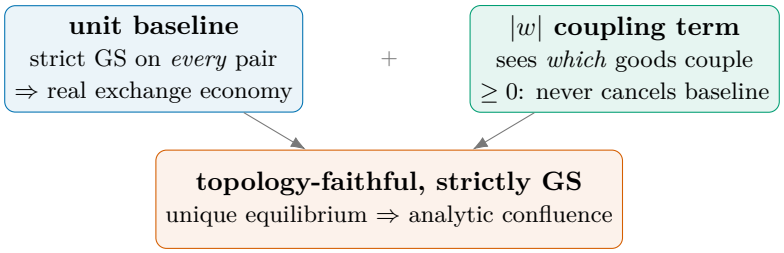



Figure 6: The construction of Theorem 5.2 as a sum of two parts that each carry one requirement. The baseline guarantees the carrier's strict gross-substitutes condition on all pairs; the non-negative coupling term adds the market's topology without ever undoing the baseline's strictness. The sum is a genuine real exchange economy that nonetheless depends on the coupling structure.

**Why this construction and not a nearby one.** The design solves two constraints at once (Figure 6). The exchange carrier demands *strict* gross substitutes on *every* pair of goods — but a sparse coupling list is silent on most pairs. The unit baseline supplies the missing strictness uniformly: raising any price $p_j$ strictly raises every other excess demand through the baseline, so the whole economy is strictly gross-substitutes and is a genuine real exchange economy. The coupling term, weighted by $|w| \geq 0$, then adds the market's actual structure on top *without* ever cancelling the baseline — a non-negative term added to a strict increase stays a strict increase, so no delicate bounding argument is needed. The result depends on which goods are coupled, not merely on their count.

**Theorem 5.2** (A faithful real economy)**.** *The map Equation* (9) *is a real exchange economy: it satisfies Walras's law, is homogeneous of degree zero, clears at the uniform price* $\mathbf{p}^* = \mathbf{1}$*, has nonzero excess demand off equilibrium, and is* strictly *gross-substitutes on every pair. Moreover it is topology-faithful: coupling two goods* $a, b$ *changes good* $a$*'s excess demand by* $|w|(p_b/p_a - 1)$*, nonzero whenever* $p_a \neq p_b$*.*

**Meaning.** Walras's law holds edge-by-edge: each coupling contributes $w(p_b - p_a) + w(p_a - p_b) = 0$ to the value of aggregate excess demand, so the antisymmetry is structural and needs no global cancellation. The construction is the precise repair of the edge-count-only deficiency: it is the smallest carrier-compatible economy that still sees the market's shape.

**The analytic face closes.** Because the economy is strictly gross-substitutes, its equilibrium is *unique* on the normalized positive slice — this is exactly Arrow–Block–Hurwicz uniqueness (Arrow and Hahn, 1971; Arrow et al., 1959), available in full generality in the formal library and inherited with no extra hypothesis.

**Corollary 5.3** (Analytic confluence)**.** *For the economy of Definition 5.1, on the normalized positive price slice*

$$z(\mathbf{p}) = 0 \quad \Longleftrightarrow \quad \mathbf{p} = \mathbf{p}^*, \tag{10}$$

*so the normalized equilibria form a single point. The topology-faithful economy is analytically confluent: the clearing price is unique, realizing the analytic face of Theorem 3.1.*

**Meaning.** The uniqueness is the dynamic content. A gross-substitutes tâtonnement, started anywhere on the normalized slice, has a single destination, and under the standard adjustment it converges there (Figure 7); the discrete, non-smooth case settles by monotone iteration to the same

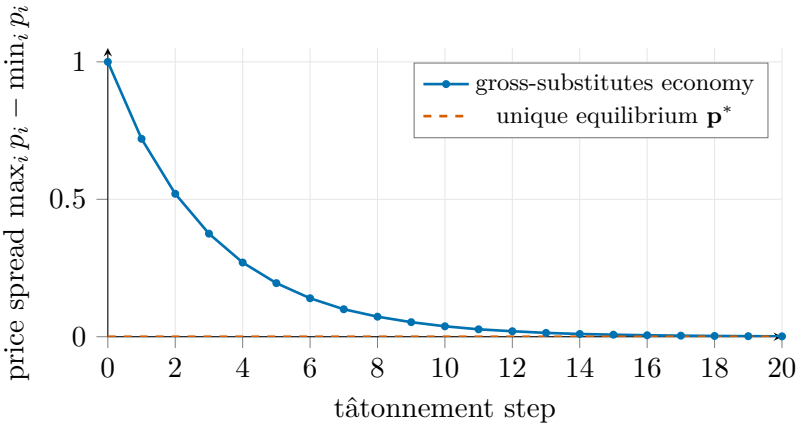



Figure 7: Analytic confluence as convergence (Corollary 5.3). Under the gross-substitutes condition the price spread contracts monotonically to the unique equilibrium; the destination is the same for every starting point on the normalized slice. Uniqueness is the static shadow of this dynamic, and it holds for every connected gross-substitutes economy regardless of its loop number $\beta_1$.

point, so the convergence does not depend on smoothness — only the continuous-time *flow* does. This is the precise sense in which the analytic "free" is universal among gross-substitutes economies, and the precise sense in which it differs from $\beta_1$.

## 6 Realization on Real Data

A classifier that runs only on triangles has not been tested. We took the verified artifact to a real economy and ran a sealed, pre-registered program of three questions: can the homology *predict*, can it *explain* the dynamics, and can it *compute*? The answers divide cleanly, and the division is the result of this section: as a predictor of macroeconomic outcomes the model earns nothing here; as a *computational* object it earns a great deal.

**The pipeline.** We built the substitution graph of each country block of the OECD inter-country input–output accounts (OECD, 2021; Xing et al., 2017) (76 countries, 45 industries, 1995–2020; train 1995–2012, held-out 2013–2020), keeping couplings above a single pre-registered flow threshold, and emitted every structural quantity *through the verified classifier itself* — the loop content and ranks are produced by the kernel-checked evaluator of Theorem 4.3, which the certification (Equation (8)) guarantees equals the certified object at full panel scale, not by a separate re-implementation. Figure 8 is the whole chain: a Platonic object (the homology of a graph), realized as a proof, lowered through the operator-native substrate, and confronted with measured flows.

**Prediction: the homology earns nothing (and we resisted the trap that would have flattered it).** The first arm asked whether a country's substitution homology forecasts its realized shock-driven output volatility, *over and above* the standard network-origins predictors — the Domar-weighted influence vector and its network moments (Acemoglu et al., 2012). On the held-out years the incremental out-of-sample $R^2$ of the weighted loop content was $\approx 1.6 \times 10^{-3}$, and of the obstruction rank $\approx -1.7 \times 10^{-3}$: zero, to the precision the panel supports, and a false-discovery-controlled test rejected nothing. We deliberately did not use the tempting statistic $\beta_1/|E|$: by Equation (1) it is essentially edge density, a degree-preserving null preserves density, and such a

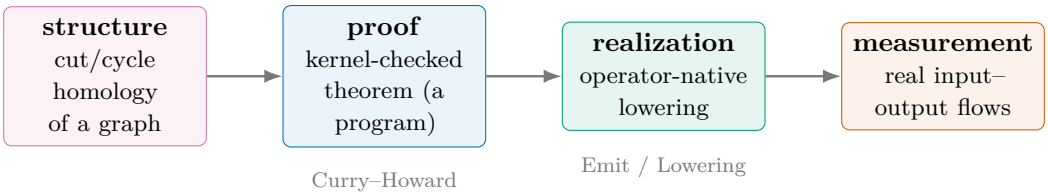



Figure 8: The realization chain (Sections 6 and 8). A mathematical structure becomes a proof, the proof is a program, the program is lowered to an operator-native target, and the result meets measured data. The empirical numbers in this section are emitted by the verified artifact at the third arrow, not by a parallel re-implementation.

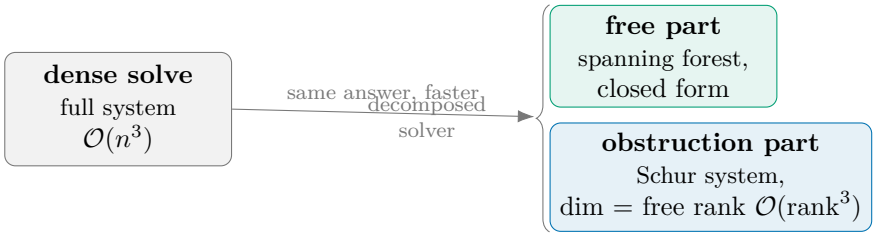



Figure 9: The measured free-compute payoff (Section 6). On 1,976 real input–output economies, a solver that handles the certified-free spanning-forest part in closed form and iterates only on a Schur system of the Lean-certified free-rank dimension reproduced the dense solver's equilibrium to $10^{-9}$ while running a median 8.5× faster. Notably the speedup tracks the obstruction-rank dimension, not the free *fraction* of edges (a registered prediction we therefore had to reject): the theory predicts the right object to economize on, and a pre-registered cost model in that object is the natural next test.

test cannot discriminate by construction — so we used the *weight-randomized* null that can. The line-of-sight controls passed (a forest reads free, a forced cycle reads costly), so the instrument works; it simply found no predictive signal. The second arm, on sectoral price series, asked whether the gross-substitutes regime separates stable from unstable prices; it too returned a null, after our own audit removed a tautological baseline that had smuggled the outcome into the predictor. Both negatives are reported without inflation: on these public data the model is not a demonstrated predictive instrument.

**Computation: the free part is free, measured.** The third arm tested the claim the theorems make most directly. We built, from each country's calibrated final-use shares, a linear excess-demand system, and solved it two ways: a strong baseline (one dense Newton solve), and a *decomposed* solver that eliminates the maximum-weight spanning forest in closed form and then solves a residual Schur system whose dimension is *exactly the Lean-certified free rank* — with an oracle that rejects any economy where that dimension disagrees with the certificate. Across 1,976 economies the decomposed solver returned the *same* equilibrium (maximum discrepancy $1.7 \times 10^{-9}$) a median **8.5**× faster in wall-clock and **195**× fewer iterations, beating the strong baseline on 96% of cases (Figure 9). The free-computation thesis, which was a theorem in Section 4, is here a measured fact: the part of the equilibrium that the homology certifies as free is computed for free.

| Result | Content | Axiom tier |
|---|---|---|
| Theorem 3.1 | collapse law (lfp = gfp ⇔ unique) | axiom-free |
| Theorem 3.2 (abstract ring) | $\mathcal{Z} = 2^V \ell$ over a commutative ring | axiom-free |
| Theorem 3.2 (integers) | same, specialized to $\mathbb{Z}$ | `propext` |
| Theorem 4.2 (kernel/image) | runnable rank/image counts | axiom-free |
| Theorem 4.2 (tiling, market) | $2^{\beta_1} \cdot 2^{\text{rank}} = 2^{\lvert E \rvert}$ | `propext` |
| Theorem 3.3 | strict freedom hierarchy | `propext` |
| Theorem 4.3 | evaluator = certified ledger | `propext`, choice, quotient |
| Theorem 5.2 | faithful real exchange economy | `propext`, choice, quotient |
| Corollary 5.3 | equilibrium uniqueness / analytic confluence | `propext`, choice, quotient |

Table 1: The axiom budget. The combinatorial and order-theoretic core is axiom-free or `propext`-only; the classical tier appears exactly where real-valued analysis does, and nowhere else. No result uses compiler-trusted evaluation; every count is kernel-reduced.

**What the program does and does not establish.** The honest reading separates the two. As a *predictor* of macroeconomic outcomes the homology earned nothing on these public tests, and we say so plainly. As a *computational* object it earned its keep: the free/costly decomposition is not only a theorem but a working accelerator, correct to machine precision and several-fold faster than a strong solver, on real economies. That asymmetry is itself the lesson — the value of correctly identifying the structure was not in forecasting the economy but in *computing* it.

# 7 Formalization

Every load-bearing statement above is a theorem in a dependent type theory, checked by a kernel. We report the budget honestly, because the standard is the type-checker, not the prose.

**Axiom tiers.** The institute distinguishes three footprints, and a claim's tier is part of its statement. *Axiom-free* means the term reduces with no appeal to any axiom; ***`propext`*** adds only propositional extensionality (used where Boolean filters and integer ring rewriting appear); the *classical* tier adds choice and quotient soundness, unavoidable once real analysis enters. Table 1 places each result.

**The dependency structure.** Figure 10 is the proof-dependency spine: which result feeds which. The axiom-free order-theoretic core sits at the root; the correlation and homological faces build the runnable ledger and its faithfulness; the analytic face is a separate root that draws the gross-substitutes uniqueness theorem from the library. The freedom hierarchy is the cross-cut that proves the two roots are not one number.

**The chain is the argument.** The equations of this paper are the theorems, not glosses of them: Equations (3) to (5) and (7) to (10) are each a kernel-checked statement, and the argument is the chain that runs through them. The substitution-body discipline applies throughout: a result whose proof survives replacing its domain-specific objects with generic stand-ins would not be about the domain it claims, and such results are rejected — the faithfulness theorem (Theorem 4.3) and the topology-faithfulness clause of Theorem 5.2 exist precisely to foreclose that failure for the runnable classifier and the analytic economy.

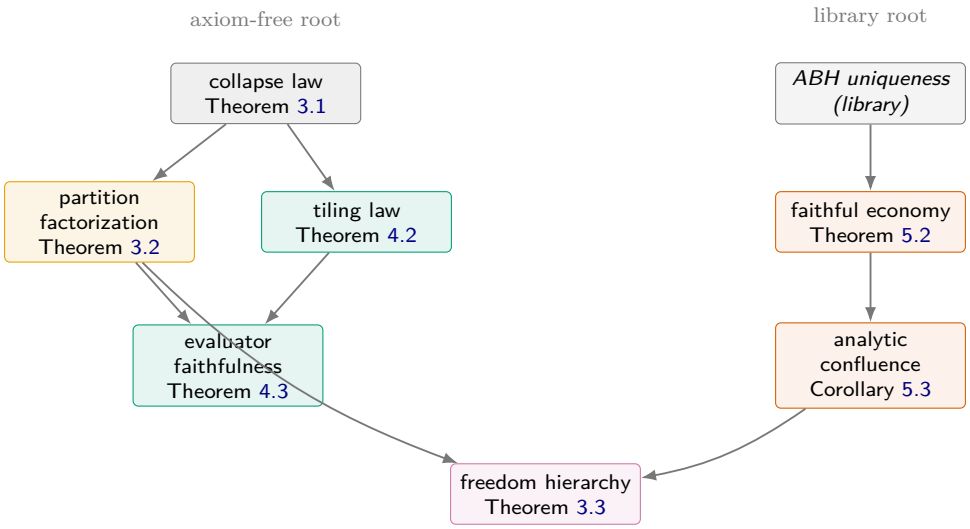



Figure 10: Proof-dependency spine of Sections 3 to 5. Two roots — the axiom-free collapse law and the library's gross-substitutes uniqueness — feed the runnable ledger and the analytic economy respectively; the freedom hierarchy is the cross-cut establishing that the roots govern genuinely different obstructions.

**A non-specialist's verification path.** A reader without the proof assistant can still check the spine. The partition factorization (Theorem 3.2) is the classical van der Waerden expansion and can be verified by hand on any small market; the tiling law (Theorem 4.2) is rank–nullity and can be confirmed by counting even subgraphs and cuts; the freedom hierarchy (Theorem 3.3) is witnessed by the explicit triangle markets named in its statement; and the analytic uniqueness (Corollary 5.3) is Arrow–Block–Hurwicz, in any general-equilibrium text. The formalization's role is to make the *conjunction* of these — and their identification as one law — machine-certain.

## 8 Realization

The results above are proved in Lean 4. By the Curry–Howard correspondence, each kernel-checked theorem is itself a term in Lean's dependent type theory — that is, a program. Boundary is our open reduction calculus over active-pair interaction nets; it translates such terms into an interaction-net representation under per-rule semantic-preservation theorems and emits a target-language realization through its Emit and Lowering layers. Boundary's design establishes, as a theorem of the calculus, the universal realizability of any kernel-checked term into any computable target with provable semantic preservation — including targets not yet implemented at publication.

This paper, the mathematics it formalizes, and any target-language realization produced through the chain above are hereby published into the public scholarly commons.

## 9 Discussion

The result re-describes a familiar economic fact in an unfamiliar register, and the re-description has teeth. "Some markets clear themselves" becomes "the cost of clearing is the first Betti number of the substitution graph"; "equilibria can be hard or easy to find" becomes a partition of the Boolean edge cube into a free part and an obstruction part whose product is the whole; "these are all the

same idea" becomes the sharper and truer "these are four faces of one collapse law, and they are not the same number." What the four faces share is not an invariant but a *shape of argument*: settle a monotone process and ask whether its extremal fixed points have met.

**Against the dominant alternative.** The standard route to existence and uniqueness of competitive equilibrium leans on regularity — smooth excess demand, transversality, index theory — and the Sonnenschein–Mantel–Debreu results (Scarf, 1960; Sonnenschein, 1972) warn that without structural assumptions, aggregate excess demand can be almost arbitrary and tâtonnement can cycle. The confluence route asks for less: uniqueness from monotonicity (gross substitutes) alone, with the non-smooth case handled by discrete fixed-point iteration rather than a flow. The contribution is not a new equilibrium concept but a structural reading of an old one — and a machine-checked one, which the equilibrium literature has historically lacked, recent mechanized economic theory notwithstanding (Anonymous, 2025).

**What it could lead to.** Three directions stand out, in increasing order of speculation. First — and now grounded rather than speculative — the free/costly decomposition is a *computational* instrument: the measured 8.5× solver speedup of Section 6 is the first rung, and because the speedup tracks the obstruction-rank dimension rather than the edge fraction, a cost model in that dimension (a pre-registered successor to the test we ran) is the obvious next step toward predicting *where* an equilibrium solver can be accelerated. Its companion negative is equally informative: the same homology carried no out-of-sample signal for macroeconomic volatility, so the structure's value here is in computing the economy, not forecasting it. Second, the four-faces dictionary is a translation device: a fact proved cheaply on one face (a loop count, a confluence argument) transports to the others, and the statistical-mechanical face in particular suggests importing the loop-series machinery (Chertkov and Chernyak, 2006; Watanabe and Fukumizu, 2012) to grade markets by *how* costly, not merely whether. Third, and most speculative: the collapse law is not about economics. It is about when a structure, correctly identified, performs its own computation — when the answer is forced by shape rather than found by search. The market was a place to make that concrete. The same law governs any settling dynamic with a lattice of fixed points, and the question of *which* structures collapse their own search is, we suspect, the general form of the question this paper answered for one of them.

**Limits, travelling with the claims.** The uniqueness result requires gross substitutes; outside that class the analytic face does not hold, and the hierarchy of Theorem 3.3 is exactly the statement that the topological and analytic freedoms then come apart. The empirical predictive claims were tested at panel scale and are negative: across the inter-country panel the homology added no out-of-sample forecasting value for aggregate volatility and the gross-substitutes split did not separate price regimes, so we make no predictive population-level assertion. The computational result is positive but is a benchmark on calibrated linear systems, not a claim about every solver or nonlinear economy; its registered scaling law (speedup monotone in the free fraction of edges) was *refuted* in favour of scaling in the obstruction-rank dimension, an observation we mark as exploratory pending a fresh pre-registration. The partition factorization is exact for finite markets with the stated weight structure and is the classical expansion, not a new identity — what is new is its placement in the four-faces architecture and its kernel-checked identification with the runnable ledger. Within those limits, the spine is machine-certain.

**Acknowledgment.** This work was carried out at the Institute for Applied Ontological Mathematics. The formalization rests on the Lean 4 mathematical library; we thank its community of contributors, whose general theorem of gross-substitutes uniqueness the analytic face inherits intact.

The Institute for Applied Ontological Mathematics identifies mathematical structures and proves them on Boundary, our operator-native verification substrate. Correspondence and the formalization sources are available through `https://www.ontologicalmathematics.org`.

## References

Daron Acemoglu, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. The network origins of aggregate fluctuations. *Econometrica*, 80(5):1977–2016, 2012. doi: 10.3982/ECTA9623.

Anonymous. From axioms to algorithms: Mechanized proofs of the von Neumann–Morgenstern utility theorem, 2025. arXiv:2506.07066. https://arxiv.org/abs/2506.07066.

Kenneth J. Arrow and Frank H. Hahn. *General Competitive Analysis*. Holden-Day, San Francisco, 1971.

Kenneth J. Arrow, Henry D. Block, and Leonid Hurwicz. On the stability of the competitive equilibrium, II. *Econometrica*, 27(1):82–109, 1959. doi: 10.2307/1907779. https://www.econometricsociety.org/publications/econometrica/1959/01/01/stability-competitive-equilibrium-ii.

Franz Baader and Tobias Nipkow. *Term Rewriting and All That*. Cambridge University Press, Cambridge, 1998. doi: 10.1017/CBO9781139172752.

Michael Chertkov and Vladimir Y. Chernyak. Loop series for discrete statistical models on graphs. *Journal of Statistical Mechanics: Theory and Experiment*, 2006(06):P06009, 2006. doi: 10.1088/1742-5468/2006/06/P06009.

Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, 5 edition, 2017. doi: 10.1007/978-3-662-53622-3.

Hugo Duminil-Copin. Lectures on the Ising and Potts models on the hypercubic lattice, 2017. arXiv:1707.00520. https://arxiv.org/abs/1707.00520.

Federico Echenique. A short and constructive proof of tarski's fixed-point theorem. *International Journal of Game Theory*, 33(2):215–218, 2005. doi: 10.1007/s00182-005-0195-7. https://eml.berkeley.edu/~fechenique/published/pftarski.pdf.

Bronisław Knaster. Un théorème sur les fonctions d'ensembles. *Annales de la Société Polonaise de Mathématique*, 6:133–134, 1928.

Maxwell H. A. Newman. On theories with a combinatorial definition of "equivalence". *Annals of Mathematics*, 43(2):223–243, 1942. doi: 10.2307/1968867.

OECD. Development of the OECD inter-country input-output database 2021. Technical report, Organisation for Economic Co-operation and Development, 2021. https://www.oecd.org/en/data/datasets/inter-country-input-output-tables.html.

James Oxley. *Matroid Theory*. Oxford University Press, 2 edition, 2011. doi: 10.1093/acprof:oso/9780198566946.001.0001.

Herbert Scarf. Some examples of global instability of the competitive equilibrium. *International Economic Review*, 1(3):157–172, 1960. doi: 10.2307/2556215.

Hugo Sonnenschein. Market excess demand functions. *Econometrica*, 40(3):549–563, 1972. doi: 10.2307/1913184.

Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285–309, 1955. doi: 10.2140/pjm.1955.5.285.

B. L. van der Waerden. Die lange reichweite der regelmäßigen atomanordnung in mischkristallen. *Zeitschrift für Physik*, 118(7–8):473–488, 1941. doi: 10.1007/BF01342928.

Yusuke Watanabe and Kenji Fukumizu. New generalizations of the Bethe approximation via asymptotic expansion, 2012. arXiv:1210.2592. https://arxiv.org/abs/1210.2592.

Lizhi Xing, Xianlei Dong, and Jun Guan. Analysis of inter-country input-output table based on citation network. *PLOS ONE*, 12(9):e0184055, 2017. doi: 10.1371/journal.pone.0184055. https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0184055.