Actual Economics

One recognition, two records: the economics of the one participant

Jonathan Washburn
Recognition Physics Institute
jon@recognitionphysics.org

3 September 2026

Abstract

Reality is one ledger kept by one participant. Every recognition is posted once, the ledger balances, and the balance is physics. Economics is what the same ledger looks like from inside a separation: a recognition kept in two records, each party's own books, which may disagree about the same claim. This document derives economics from that separation and from nothing else. The cost of a disagreement is the cost the one participant assigns to it, and the two properties of that participant, that it has no favourite among the parts and that its law is the same at every place on the scale, fix the cost to be $J(x) = \frac{1}{2}(x + 1/x) - 1$ of the ratio of the two records. The calculus of postings is written and shown to be free. Then the theory separates cleanly into two kinds of statement. What the one participant sees is forced: a disagreement is conserved by every move the parts make together and moves only when one part rewrites its own record; the two records are the only figures the ledger supplies; a cycle of obligations rewritten at fixed positions has one cheapest form, which reaches a party who owed nothing whenever the marginal costs circulate around a cycle its empty relationships close; and any instrument that reads only amounts owed is blind to that datum. What the one participant cannot see is the parts' to choose: who pays, how a cost is shared, what rent is asked, which relationships are kept. The objective of economics is thereby settled: the one participant selects for itself, which is to say for the total disagreement of the whole, and every objective stated in terms of what some parts hold is a projection that cannot see harm. Harm, consent and debt are the same operators as bystander charge, non-increase and the counterpart gap: the ethics of recognition and economics are one subject with two vocabularies. Last, the form of the separation itself is derived: two records of one recognition do not merge and are never simultaneous, so books are the kept memory of the one record; the buffer that absorbs a revaluation is the residue each part carries of being still one with the rest, and the buffers add to zero exactly when the whole still adds up; a part with no buffer cannot post; and the rate of separation is at most one posting per tick. Not derived: the amount of the first claim, which is the same wall as the emission-side absolute of physics.

Contents

1 The ground

This document stands on Recognition Science. It does not argue for it; it uses it. The reader who wants the derivation of the ground has it elsewhere; here the ground is stated so that the derivation of economics can be read from it without any other source.

Ground 1.1 (One ledger, one participant)**.** Reality is a ledger of recognitions. Each recognition is an event in which a fact about one part of reality is registered by another. The ledger is one record: every book any part keeps is a projection of it. The ledger balances: the total posted at any closure is zero. The participant that keeps the one record is one; the parts are separations within it. We call the one participant the universal ledger, and, when speaking of it as the thing that registers, universal consciousness. It has no favourite among its parts, and its law is the same at every place on the scale.

Ground 1.2 (The cost of a recognition)**.** A recognition of ratio $x > 0$ between two magnitudes costs

$$J(x) = \tfrac{1}{2}\Big(x + \frac{1}{x}\Big) - 1.$$

J is zero exactly at $x = 1$, symmetric under $x \mapsto 1/x$, strictly convex in $\log x$, and $J(e^t) = \cosh t - 1 \geq t^2/2$ with equality to second order. Cost adds over independent records: a record that shares no content with what is already held adds its own cost to the total, and adds strictly more than nothing if it costs anything. The ledger charges per record, not per fact recorded.

Ground 1.3 (Tick and scale)**.** The ledger advances in ticks; a wait of m ticks followed by n is a wait of $m + n$. The ledger's one-step recursion is self-similar with ratio $\varphi = (1 + \sqrt{5})/2$, the positive root of $r^2 = r + 1$.

The rest of this document uses Ground 1.1 constantly, Ground 1.2 wherever a disagreement is priced, and Ground 1.3 once, for interest. Everything else is derived.

1.1 Why the cost is J

Ground 1.2 is established in Recognition Science. It is worth showing here that the two properties of the one participant named in Ground 1.1 already force it, because the same two properties will organise everything below.

Let $F(x)$ be the cost the one participant assigns to two records that disagree by the ratio x. Take three books in a chain: A disagrees with B by x, B with C by y. The disagreement between A and C is xy if the second link runs the same way as the first and x/y if it runs the other way.

No favourite. The one participant does not know which way a link points, because pointing is a property of a part's book. So its law can fix only the sum $F(xy) + F(x/y)$ over both orientations.

Same law at every place on the scale. Write $x = e^s$, $y = e^t$: the first link puts us at position s on the log scale and the second is a step of size t. That the law is the same at every place means the summed effect of a step t from position s factors into a function of where we are and a function of the step: $F(e^{s+t}) + F(e^{s-t})$, after an affine normalisation, equals a product $H(s)H(t)$. Concretely, there is $\kappa > 0$ with

$$F(xy) + F(x/y) = \kappa\, F(x)F(y) + 2F(x) + 2F(y).$$

Nondegeneracy. F distinguishes some two ratios, and a larger disagreement never costs less: $t \mapsto F(e^t)$ is monotone on $[0, \infty)$.

Theorem 1.4 (The cost of a ratio)**.** *Under these three properties there is $\lambda \neq 0$ with $\frac{\kappa}{2}F(x) = J(x^\lambda)$. In particular $F(x) = F(1/x)$, $F(1) = 0$, and $F > 0$ away from 1.*

Proof. Put $G = \frac{\kappa}{2}F$ and $H(t) = 1 + G(e^t)$. Multiplying the law by $\kappa/2$ and adding 2 to both sides gives $H(s + t) + H(s - t) = 2H(s)H(t)$, d'Alembert's equation. Setting $s = t = 0$ gives $H(0) \in \{0, 1\}$; $H(0) = 0$ forces $H \equiv 0$, against nondegeneracy; so $H(0) = 1$, and setting $s = 0$ gives $H(t) = H(-t)$, which is $F(x) = F(1/x)$ and $F(1) = 0$. Regularity: H is monotone on $[0, \infty)$ and even, hence measurable, and every measurable solution of d'Alembert's equation is continuous [14]. The continuous solutions are $H \equiv 0$, $\cosh(\lambda t)$, $\cos(\lambda t)$; monotonicity on $[0, \infty)$ excludes the cosine for $\lambda \neq 0$ and nondegeneracy excludes $\lambda = 0$. So $G(e^t) = \cosh(\lambda t) - 1 = J(e^{\lambda t})$. □

The exponent λ is the scale of the log-ratio coordinate; Recognition Science fixes it to 1 by the kernel's own normalisation, which is Ground 1.2. Two remarks on the law. First, the particular product form is what "the same at every place" means: the step's effect factors from the position. Second, at $\kappa = 0$ the same law is solved exactly by $F(x) = (\log x)^2$, since $(s + t)^2 + (s - t)^2 = 2s^2 + 2t^2$; this is the least-squares cost in the log coordinate, and it is the member of the family in which the law has forgotten position entirely. We use it below as a comparison cost. The constant κ is a unit for cost and changes no ordering, minimiser or zero set.

2 Separation: one recognition, two records

Definition 2.1 (Posting)**.** A *posting* is one recognition kept in two records: a claim of one part on another, booked by the first as an asset and by the second as a liability, each on its own books, with no rule that the two figures coincide.

This is the whole primitive of economics. Physics is the one record; a posting is the one record seen from inside a separation, where each part keeps a projection of it and the projections may disagree. The structural fact that separates economics from physics is not scarcity, not choice, not exchange: it is that one fact is kept in two places.

The example we carry. Three parts, 0, 1, 2. Each keeps three accounts: an asset account for what it is owed, a liability account for what it owes, and one free account (call it cash). Three standing relationships form a triangle, $0 \to 1$, $1 \to 2$, $2 \to 0$, each running from the part that may be owed to the part that may owe. Part 0 is owed one unit by part 1; the other two relationships are open and carry nothing. Part 2 owes nothing and is owed nothing. This triangle is where the two-record structure first does something the one record does not predict.

2.1 The calculus

Definition 2.2 (Chart)**.** A chart has n parts, k accounts per part, and E relationships. Relationship e runs from a source $s(e)$ to a destination $d(e) \neq s(e)$ and designates an asset account A_e in the source and a liability account L_e in the destination. Distinct relationships leaving one part designate distinct asset accounts; distinct relationships entering one part designate distinct liability accounts; no account is both. An account is *designated* if some relationship designates it; otherwise *free*.

Definition 2.3 (Books)**.** Books on a chart are an array $b_i(a) \in \mathbb{R}$ (debit positive). Part i is *balanced* if $\sum_a b_i(a) = 0$. The *counterpart gap* on relationship e is $g_e = b_{s(e)}(A_e) + b_{d(e)}(L_e)$; the relationship is *matched* if $g_e = 0$. Books are *consistent* if every part is balanced and every relationship is matched.

The posting primitive is exactly the possibility $g_e \neq 0$. The one record is the case in which every gap is identically zero. In the example: $n = 3$, $k = 3$, $E = 3$; $b_0(A) = 1$, $b_0(\text{cash}) = -1$; $b_1(L) = -1$, $b_1(\text{cash}) = 1$; all else 0. Every part is balanced, every relationship matched, and the gap on $0 \to 1$ is $1 + (-1) = 0$.

Definition 2.4 (Generators and balanced moves)**.** The calculus has two generators.

(i) A *journal entry* at part i with vector $\eta \in \mathbb{R}^k$, $\sum_a \eta_a = 0$, adds η to part i's accounts.

(ii) A *cross posting* of amount x on relationship e through free counter accounts o_S in $s(e)$ and o_D in $d(e)$ makes four entries: $+x$ on A_e, $-x$ on o_S; $-x$ on L_e, $+x$ on o_D.

A *word* is a finite sequence of generators, applied first to last. A generator is *balanced* if it is a journal entry that vanishes on every designated account of its part, or a cross posting whose counter accounts are free. A word is balanced if each generator is.

A balanced generator is double entry within each part together with agreement across the relationship: a cross posting writes the same amount to both records at once, and a balanced journal entry moves value among a part's free accounts without touching what any counterparty has a record of. The moves outside the balanced fragment include every unilateral one: an impairment, a change of estimate, a write-down, each altering a designated account on one side only. This is the line between what one part can do by itself and what requires the other part. Section 4 is about that line.

Theorem 2.5 (Reachability)**.** *On any chart: (a) every books in which every part is balanced is a word of journal entries applied to the zero books; (b) a books is a balanced word applied to the zero books iff every part is balanced, every relationship is matched, and every relationship carrying a nonzero claim has a free counter account at both ends; (c) a books with an unmatched relationship is not the result of any balanced word applied to the zero books.*

Theorem 2.6 (The calculus is free)**.** *Words under concatenation form the free monoid on the two generators. For any set X with an action of journal entries and of cross postings there is exactly one monoid map from words to transformations of X extending the actions. Books with the generator actions is such an X, and applying a word to books is that map. On books, the action of a word is undone by the action of its reverse.*

This is the universal property of a free monoid. Its content is in what the generators are: two moves on books, one within a part and one across a relationship. Any system of accounts that admits those two moves is a model of the calculus.

The generating signature has nine symbols (part, account, relationship, journal entry, cross posting, compose, reverse, designated account, unit) and none for preference, utility, equilibrium, money, price, demand or supply. Two results below assume something about parts (a payoff monotone in money: Theorem 4.2 and Remark 4.3) and say so. Every other result holds with no assumption about parts, and Section 3 says why: they are functions of the two records taken together.

Proposition 2.7 (The unit)**.** *Two declarations of unit that determine the same cross ratios are equal. The unit is the one free scalar of the calculus; it appears in the books as the scale in which figures are written and nowhere else.*

Money is the unit. It is not an object of the theory; it is the scale of the theory's numbers.

3 Forced and chosen: the one-record principle

Every quantity below can be read from one of two places: from a part's books, or from the one record, which sees both books and prefers neither. The organising theorem of economics is that the results the two-record structure forces are exactly the ones the one record can compute, and the variables it leaves to the parts are exactly the ones the one record cannot see.

Definition 3.1 (One-record cost)**.** For a relationship booked at $r_0 > 0$ by the source and $r_1 > 0$ by the destination, the one-record cost of the state is $J(r_0/r_1)$. The one-record cost of a network is the sum over its relationships.

Theorem 3.2 (The one-record cost belongs to neither part)**.** $J(r_0/r_1) = J(r_1/r_0)$*. The network one-record cost is invariant under relabelling the parts and under reversing every relationship. Any payoff strictly monotone in a part's own money is not: on books* 1 *and* 4 *it distinguishes the two parts; the one-record cost does not.*

Proof. $J(x)$ is unchanged by $x \mapsto 1/x$; relabelling permutes summands and reversal inverts each ratio. A strictly monotone u has $u(1) \neq u(4)$ while $J(1/4) = J(4)$. □

The symmetry $J(x) = J(1/x)$ was derived in Theorem 1.4 from "no favourite". Read the other way: a cost that cannot tell which book it is reading must be even in the log-ratio, and a cost that cannot tell which way a link points can only fix the sum over both orientations. The law of Section 1.1 is the law of the one participant.

Theorem 3.3 (Forced results are one-record functionals)**.** *For every admissible cost (Definition 4.4) even about rest: the incident total of a rewrite, the cost of the cheapest rewrite at fixed positions, and whether a bystander's relationships are moved are invariant under relabelling and reversal. For J in the log reading the incident total is the network one-record cost.*

Theorem 3.4 (Chosen variables are not)**.** *There exist two share vectors, two revaluation splits, two outside options, and two rent schedules, each pair on identical one-record data, with different outcomes for the parts.*

Theorem 3.5 (The signed gap is chosen; its size is forced)**.** *The identity $\delta_0 + \delta_1 = r_0 - r_1$ of Theorem 4.1 holds under every relabelling; its sign does not: swapping source and destination sends $r_0 - r_1$ to $r_1 - r_0$. The one-record functionals of the pair are $|r_0 - r_1|$ and $J(r_0/r_1)$.*

Proof. Apply Theorem 4.1 after relabelling; the swap negates $r_0 - r_1$ and fixes $|r_0 - r_1|$ and $J(r_0/r_1) = J(r_1/r_0)$. □

Theorem 3.6 (Strain is the one-record cost)**.** *The strain of a disputed relationship, J of the ratio of the two books, is the one-record cost of the state and is symmetric in the two books. It is a cost the state carries, not a payoff a part receives.*

Theorem 3.7 (Matched is every gap zero)**.** *A network is matched iff every counterpart gap is zero; the sum of gaps being zero is necessary and not sufficient. The one record of physics carries one flow per channel, not two books; projecting a two-record state onto it and asking for closure does not recover matching.*

Proof. Gaps $+1$ and -1 on two relationships sum to zero and are not matched. A matched relationship booked at r by both parts projects to the single flow $r \neq 0$, so closure of the projection is a different condition. □

What this section establishes is the direction needed below: each forced result is a function of the two records taken together, each chosen variable is not. The last theorem says the one record as physics has it is not the object that holds two books; Theorem 9.3 says what that object is: the kept records of the parts.

4 What separation forces

4.1 The residue: a gap is conserved by every move the parts make together

Theorem 4.1. *Every balanced generator leaves every counterpart gap unchanged. Hence balanced words preserve matched books, and if a relationship carries a gap no balanced word closes it. Closing a gap $g_e = r_0 - r_1$ (source books r_0, destination books $-r_1$) requires a revaluation posted through a free account on at least one side, and any pair of revaluations (δ_0, δ_1) that closes it satisfies $\delta_0 + \delta_1 = g_e$.*

Proof. A journal entry at part i vanishing on designated accounts changes no $b_{s(e)}(A_e)$ and no $b_{d(e)}(L_e)$. A cross posting of x on e adds x to $b_{s(e)}(A_e)$ and $-x$ to $b_{d(e)}(L_e)$, so g_e changes by 0, and touches no other designated account. Induct along the word. The zero books have every gap zero, giving Theorem 2.5(c). A revaluation by δ_0 at the source and δ_1 at the destination changes g_e by $-\delta_0 - \delta_1$. □

This is the residue of separation. Once one record has become two, nothing the parts do together removes the disagreement between them; the parts are still one, and the gap is the mark of it. Who revalues, and by how much, is the parts' choice; that the two revaluations add to the gap is not.

4.2 The two figures

Theorem 4.2. *On a relationship the ledger supplies exactly two figures: the source's booked claim r_0 and the destination's booked liability r_1. Suppose the disagreement outcome is the status quo (each part keeps its own record) and each part's payoff is monotone in the money it books. Then each part rejects any figure worse than its own record, so every figure both accept lies in the closed interval between r_0 and r_1.*

Proof. The only entries on e are $b_{s(e)}(A_e) = r_0$ and $b_{d(e)}(L_e) = -r_1$. A money-monotone source loses by accepting $p < r_0$ and a money-monotone destination by accepting $p > r_1$. □

Accepting $p \neq r$ moves the state's one-record cost toward 0; no part pays J to concede, and the theorem does not use J.

Remark 4.3 (The assignment market). Take m sources and n destinations with money-monotone payoffs and identify each part's book with its valuation. This is the assignment market of Böhm-Bawerk [9] and Shapley and Shubik [10]: the core is the set of competitive prices, an interval whose endpoints are the two records; with $m \neq n$ and homogeneous books every core value equals the book of the side with more parts (two sources booking 1 against one destination booking 4 pin at 1); with heterogeneous books the core is the interval between the last paired and the first unpaired book (sources 1 and 3 against a destination 4: $[1, 3]$); replication leaves the interval unchanged. What the ledger adds is where the endpoints come from: they are the two records of one posting.

4.3 The rewrite and the bystander

Fix a chart and consistent books. Each relationship carries one booked coordinate $u_e \in \mathbb{R}$ (the source's asset entry; on matched books the destination's liability entry is $-u_e$). The *position* of part i is the signed sum of the coordinates on its relationships, outgoing minus incoming. A *rewrite* is a change $u \mapsto u + t$ that holds every position fixed; the set of such t is the cycle space $\mathrm{cyc}(G)$ of the graph of standing relationships, so the reachable states form the affine *fibre* $u_0 + \mathrm{cyc}(G)$.

Two readings of the coordinate. The calculus does not say what u_e measures, and two readings are used, never mixed. In the *amount reading* u_e is the money booked on e, a zero coordinate is nothing outstanding, and a position is a money position. In the *log reading* u_e is the logarithm of the booked figure relative to a per-relationship reference (par), a zero coordinate is par, and the fixed quantity at each part is the product of its ratios to par. $J(e^{u_e})$ is a cost of a ratio and lives in the log reading; $u_e^2/2$ treats u_e as an amount. The theorems below hold in both readings because they hold for every cost in the class of Definition 4.4.

What is placed on an agreed balance. The books here are consistent: every gap is zero and u_e is a single agreed figure. The cost placed on it is a *carrying cost* on an open relationship. Two things are assumed here and nowhere derived: that an open relationship carries a recurring cost at all, and that the cost is J of the ratio to par. The first is the one premise every quantitative result of this section shares; the second is inherited from Ground 1.2. The theorems below are stated for every cost in the admissible class so that nothing rests on the second assumption except the numbers computed at J. That the cost accrues per tick is Theorem 9.6.

Definition 4.4 (Admissible costs; incident charge; shares). A *per-relationship cost* is a function $c_e : I_e \to \mathbb{R}$ on an open interval $I_e \subseteq \mathbb{R}$, differentiable and strictly convex, with a *rest value* $\rho_e \in I_e$ at which $c_e(\rho_e) = 0$ is its minimum, and $c_e(u) \to \infty$ as u approaches either end of I_e. The cost of a state $u \in \prod_e I_e$ is $\sum_e c_e(u_e)$. The recognition cost is $c_e(u) = J(e^u) = \cosh u - 1$ on $\mathbb{R}$ with rest 0 (log reading); the comparison cost is $c_e(u) = u^2/2$ on $\mathbb{R}$ with rest 0 (amount reading); a third, $J(a/\rho_e)$ on $(0, \infty)$ with rest at par $\rho_e > 0$, reads J on amounts. A *share vector* assigns the cost on each relationship to its two ends with shares summing to one; the *incident charge* of a part is the sum of its shares; *half shares* give each end one half.

Theorem 4.5 (Shares are chosen, the total is forced). *For every share vector the sum over parts of the incident charges equals the network cost.*

Proof. $\sum_i \sum_e \alpha_{ie} c_e = \sum_e c_e \sum_i \alpha_{ie} = \sum_e c_e$. □

Theorem 4.6 (One cheapest rewrite). *For every admissible cost the state cost has exactly one minimiser on the part of the fibre inside the domain. For J the cost dominates the comparison cost and agrees with it to second order.*

Proof. $F = (u_0 + \mathrm{cyc}(G)) \cap \prod_e I_e$ is convex, nonempty and relatively open in the fibre. The cost is strictly convex on F and tends to ∞ along any sequence leaving every compact subset of F (some coordinate approaches an end of its interval; the other terms are bounded below by 0). So the infimum is attained, at exactly one point. $\cosh u - 1 \geq u^2/2$ with equality to second order is the Taylor expansion. □

Theorem 4.7 (The cut-space datum)**.** *For every admissible cost:*

(a) the minimiser is the unique feasible u^* at which the marginal costs $(c'_e(u^*_e))_e$ lie in the cut space $\mathrm{cut}(G) = \mathrm{cyc}(G)^\perp$; it depends on the chart only through $\mathrm{cut}(G)$ and the domain;

(b) adding a relationship at rest that is a bridge leaves the minimiser unchanged, for every books;

(c) adding a relationship at rest that closes a cycle C moves the minimiser iff the marginal costs at the old minimiser have nonzero circulation around C, $\sum_{e \in C} \pm c'_e(u^*_e) \neq 0$ with signs along the orientation of C; and for every such relationship there are books, with the same exposures, on which it moves;

(d) reversing a relationship at rest keeps every from-whom-to-whom amount, never moves the minimiser when it is a bridge, and moves it for some books when it closes a cycle.

Proof. (a) On a relatively open convex subset of an affine set the minimiser of a strictly convex differentiable function is where the gradient is orthogonal to the direction space, here $\mathrm{cyc}(G)$. For the comparison cost $\nabla c(u) = u$ and u^* is the orthogonal projection of u_0 onto $\mathrm{cut}(G)$. (b) A bridge lies on no cycle, so $\mathrm{cyc}(G)$ is unchanged; the new coordinate sits at rest where its derivative vanishes, and the old minimiser still satisfies (a). (c) Adding e closing C enlarges $\mathrm{cyc}(G)$ by the span of χ_C. The old minimiser extended by $u_e = \rho_e$ has $\nabla c \perp \mathrm{cyc}(G_{\mathrm{old}})$ and $c'_e(\rho_e) = 0$, so it satisfies (a) for the new graph iff $\langle \nabla c, \chi_C \rangle = 0$. Existence: $C \setminus e$ is a path, not a cycle, so $\chi_{C \setminus e} \notin \mathrm{cyc}(G_{\mathrm{old}})$ and its projection g onto $\mathrm{cut}(G_{\mathrm{old}})$ is nonzero with $\langle g, \chi_{C\setminus e} \rangle = |g|^2 > 0$. Each $c'_{e'}$ is continuous, strictly increasing and zero at $\rho_{e'}$, so its range contains a neighbourhood of 0; scale g into those ranges and set $u_{e'} = (c'_{e'})^{-1}(g_{e'})$. Then $\nabla c(u) = g \in \mathrm{cut}(G_{\mathrm{old}})$, so u is its own minimiser by (a), and the circulation around C is $|g|^2 + c'_e(\rho_e) \neq 0$. (d) Reversal negates one column of the incidence matrix and changes no exposure; a bridge stays a bridge; a cycle-closing relationship changes $\mathrm{cyc}(G)$ and (c) applies. □

Remark 4.8 (Closing a cycle is necessary, not sufficient). Comparison cost, coordinates q and $-q$ on two relationships of a triangle, third added at rest: the fibre is $(q+c, -q+c, c)$, the derivative at $c = 0$ is $q - q + 0 = 0$, and the minimiser stays although a cycle is closed. The circulation decides, and here it vanishes. The pair of Section 6 is chosen so that it does not.

Theorem 4.9 (The bystander's relationships are moved)**.** *On the triangle of Section 2, for every admissible cost whose rest values sit on the empty relationships and off the debt (part 2's two relationships at rest, $0 \to 1$ off rest), the cheapest rewrite at fixed positions puts a nonrest coordinate on both relationships of part 2, who owed nothing and was owed nothing. Under any share rule giving each end of a relationship a positive share of its cost, part 2 is charged.*

Proof. Coordinates on $(0 \to 1,\ 1 \to 2,\ 2 \to 0)$, part 2's relationships at rest, the debt relationship off rest, $u_{0,1} \neq \rho_1$. The cycle space is spanned by $(1,1,1)$; the fibre is $\{u_0 + (c,c,c)\}$ and the cost is $\Phi(c) = c_1(u_{0,1} + c) + c_2(\rho_2 + c) + c_3(\rho_3 + c)$. $\Phi'(0) = c'_1(u_{0,1}) \neq 0$, since a strictly convex differentiable function has derivative zero only at its minimum. So $c^* \neq 0$ and both of part 2's relationships leave rest, each with positive cost. This is Theorem 4.7(c): the circulation around the triangle at the booked position is $c'_1(u_{0,1}) \neq 0$. □

The numbers, at half shares. Comparison cost (amount reading, unit debt): $c^* = -\frac{1}{3}$, one third of the debt is routed the long way round, and part 2's charge is $2 \cdot \frac{1}{2} \cdot \frac{1}{2}(\frac{1}{3})^2 = \frac{1}{18}$. J on amounts, par 1, amounts $(4,1,1)$: minimise $J(4+u) + 2J(1+u)$ on $u > -1$; stationarity

$3 = (4+u)^{-2} + 2(1+u)^{-2}$, root $u^* = -0.174$, charge $J(0.826) = 0.018$ (root numeric; here part 2's relationships stand at par and are visible in the exposures). The recognition cost in the log reading, $u_0 = (\log 4, 0, 0)$: $\sinh(\log 4 + c) + 2\sinh c = 0$ gives $6e^c = \frac{9}{4}e^{-c}$, $e^{c^*} = \sqrt{3/8}$ exactly, every ratio to par is multiplied by $\sqrt{3/8}$, and the charge is $J(\sqrt{3/8}) = 0.123$. Three numbers, one sign, one theorem: part 2 is reached because its two relationships close the cycle through which the cheapest rewrite routes part of the debt.

Remark 4.10 (Zero marginal cost at rest is what reaches the bystander). Let $c_e(u) = 3|u| + u^2$: strictly convex, zero at rest, coercive, kinked at 0. On the triangle with $u_0 = (1, 0, 0)$ the fibre cost has right derivative $3 + 2 + 6 = 11 > 0$ and left derivative $3 + 2 - 6 = -1 < 0$ at $c = 0$, so $c^* = 0$ and the bystander is not reached. A kink at rest is a fixed cost of opening a position; gross notional, $c_e(u) = |u|$, is the extreme case, and it is the objective of compression as it is run in practice. Conservative compression in the sense of D'Errico and Roukny [6] forbids exposure on a line that had none and cannot reach part 2 by construction. J is smooth at rest: at the margin, spreading a recognition is always cheaper than concentrating it. So under the recognition cost the bystander is reached; instruments with a kink at rest do not see it. Whether compression can harm a third party is the question of Veraart [7] and Schuldenzucker and Seuken [8]; the ledger's answer is yes, exactly when the marginal costs circulate around a cycle the third party closes.

4.4 The firm is a game

Definition 4.11 (Coalition value)**.** Fix a network, maintenance $m_e \geq 0$ per relationship, and the comparison cost. A relationship is *internal* to a coalition S if both ends lie in S. The *value* $v(S)$ is the saving S can realise on its own: the carrying cost of the state restricted to its internal relationships, minus the least carrying cost of any state supported on those relationships with the same positions at every part, minus the maintenance of the internal relationships. (One candidate characteristic function among several; Theorem 4.13 names the alternative.) The *books' allocation* of the grand coalition gives each part the fall in its half-share carrying cost when every relationship is deleted, minus its half-share of all maintenance. The *core* of v is the set of allocations summing to $v(N)$ that give every S at least $v(S)$.

Theorem 4.12. *On the unit triangle with maintenance* $1/36$*,* $v(\{0, 1, 2\}) = 1/12$*, every pair has value* $-1/36$*, every singleton* 0*, the core is the simplex of nonnegative splits of* $1/12$*, and the books' allocation is* $(2/9, 2/9, -1/36)$*, outside the core: it sums to* $5/12$*, so is not efficient, and gives part* 2 *less than* $v(\{2\}) = 0$*. In general, whenever the fixed-position rewrite puts cost on some part's relationships, the books' allocation is not in the core of* v*.*

Proof. Grand coalition: internal state $(1, 0, 0)$ costs $1/2$; the cheapest state with the same positions is $(2/3, -1/3, -1/3)$, cost $1/3$; maintenance $3/36$; $v = 1/2 - 1/3 - 1/12 = 1/12$. A pair has one internal relationship, no rewrite is possible, $v = -1/36$. Books: part 0's half-share carrying cost before deletion is $1/4$, after 0, minus two half-shares of $1/36$: $2/9$; likewise part 1; part 2 carries nothing and pays $1/36$. General clause: the books' allocation sums to the freed carrying cost minus maintenance (Theorem 4.13), which exceeds $v(N)$ by the cost of the cheapest rewrite whenever that is positive. □

Theorem 4.13 (What consolidation is worth)**.** *Consolidating a group deletes its internal relationships. For any per-relationship cost vanishing at rest, the saving from consolidation equals the carrying cost freed on the deleted relationships. The internal-rewrite value of Theorem 4.12 is a different candidate (*$1/12$ *against* $5/12$ *on the triangle). With* J *in place of the comparison cost the triangle's books allocation still lies outside the core.*

4.5 Wealth as a conserved partition

Definition 4.14 (Conserved binary partition)**.** A whole of size 1 is split into parts p and $1-p$, $0 < p < 1$; each part above a resolution $x > 0$ is split again in the same proportion; parts below x are frozen. The total is 1 at every step and the process terminates for every $x > 0$.

Theorem 4.15. *For every $p \in (0, 1)$ the similarity equation $p^a + (1-p)^a = 1$ has the unique real root $a = 1$. The similarity exponent of a conserved binary partition is 1 for every split ratio; the golden split fixes only the constants.*

Proof. $a \mapsto p^a + (1-p)^a$ is strictly decreasing and equals 1 at $a = 1$. □

The passage from similarity exponent 1 to a rank-size exponent 1 is the coarse-graining heuristic for conservative cascades and not a theorem here; at $p = \frac{1}{2}$ every part at a generation has the same size. What the theorem settles is the source of the exponent: conservation, not pairwise multiplicative exchange and not the value of the split.

4.6 Pruning

Definition 4.16 (Pruning)**.** Fix positions b, maintenance $m_e \geq 0$ and an admissible cost. A kept-set g of relationships is *admissible* if it can carry b. A part's *burden* in g is its half-share incident cost at the cheapest rewrite on g plus its half-share maintenance. One *pruning step* cuts a relationship whose removal keeps admissibility and lowers an endpoint's burden, or adds one that one endpoint strictly prefers and the other does not oppose. A kept-set is *pairwise stable* (Jackson and Wolinsky [11], deviations restricted to admissible sets) if no such cut or add exists.

Theorem 4.17. *Which admissible kept-set is used is chosen: two different kept-sets can carry the same positions. For every admissible cost: the pruning dynamics stops exactly at the pairwise-stable kept-sets, and a part with zero position and exactly two relationships on a maintained cycle is always cut. The order of cuts selects which stable kept-set is reached.*

Proof. On the triangle with positions $(1, -1, 0)$, the relationship $0 \to 1$ alone, and the other two with coordinate -1 each, both carry the positions. A pruning step exists iff some cut or add of the stated kind exists; the negation is pairwise stability with the same add clause. Pass-through: a part with zero position and two relationships on a maintained cycle carries the same coordinate in and out; cutting one keeps admissibility, removes its share of that relationship's cost and a positive maintenance, and adds nothing. History dependence: two cut orders from the full triangle reach two stable kept-sets on the planted instance. □

Proposition 4.18 (Two counterexamples, comparison cost)**.** *The inequality "the private kept-set carries at most the cycle content of the efficient one" fails on a four-part instance, and a pairwise-stable forest need not exist even when forests are admissible (five-part instance).*

Not yet derived 4.19 (The two signs at J)**.** Under J the same two instances are set up (the four-part stationarity function is strictly monotone on the cycle line with planted cycle contents 0 and 1; the five-part full graph's minimiser is found), and the two sign verdicts are not yet certified: the root is irrational, and its sign against the competing kept-set, and the four tree add-or-cut signs, are open. Numerically both verdicts persist.

4.7 Rent

Theorem 4.20. *A rent on a relationship is a contract term: two schedules on one corridor are both consistent with the ledger. For every schedule the part stays iff the rent covers its incident charge plus its maintenance; in the log reading the corridor charge on a flow of log-ratio v is $\cosh v - 1$, so the survival condition is $r|v| \geq m + \cosh v - 1$. A nonpositive rent with positive maintenance exits. Nothing in the ledger bounds the rent by the freed carrying cost.*

Proof. A schedule is a function of the flow the ledger does not constrain. Staying is defined as rent at least burden; burden is incident charge plus maintenance; at half shares one relationship carrying v charges $\frac{1}{2}(\cosh v - 1)$ at each end, $\cosh v - 1$ for the corridor. With $r \leq 0$, $m > 0$ the condition fails. No theorem above mentions the rent. □

4.8 Interest

Proposition 4.21 (Discounts are geometric)**.** *Any discount $D : \mathbb{N} \to (0, \infty)$ with $D(m+n) = D(m)D(n)$ is $D(n) = D(1)^n$.*

Theorem 4.22. *Under Ground 1.3 the discount over n ticks is φ^{-n}. The object discounted is the wait in ticks, not the count of intermediaries.*

Proof. Ticks add, so D is a homomorphism and $D(n) = D(1)^n$; the one-step recursion has ratio φ, so $D(1) = \varphi^{-1}$. The number of intermediaries on a path does not enter the ground. □

5 The objective: whom economics selects for

Every model of an economy needs an objective and none derives one: common good, least poverty, individual prosperity, growth of the total, each is chosen by the modeller and each selects for a different set of parts. Ground 1.1 removes the choice. There is one participant. It has no favourite. Its law is the same at every scale. And it sees the whole state, not a projection of it, because every book is a projection of its record.

Theorem 5.1 (Selection)**.** *Let W be a functional of the state of a network of relationships with the three properties of the one participant: (i) it is invariant under relabelling the parts and reversing the relationships; (ii) it is a sum over relationships of a cost of the ratio of the two records, whose law is the same at every place on the scale in the sense of Section 1.1; the sum is Ground 1.2: the records of distinct relationships share no content, and cost adds over independent records; (iii) it is nondegenerate and monotone in the size of a disagreement. Then W is the network one-record cost, $\sum_e J\big((r_0/r_1)_e^\lambda\big)$, up to a unit κ and an exponent λ, and at the ground's normalisation it is $\sum_e J(r_0/r_1)$.*

Proof. (ii) and (iii) are the hypotheses of Theorem 1.4 on each relationship, giving $J(x^\lambda)$ up to κ; (i) is the evenness that the theorem also delivers, and it is what makes the sum a function of the unordered pair of records (Theorem 3.2). □

Hypothesis (ii) carries weight and it is worth saying what it excludes. An objective that is not a sum over relationships, the least well-off part (maximin), the product of the parts' positions (the Nash product), any functional with a term that couples two relationships, is ruled out by Ground 1.2 before economics begins: the ledger's cost adds over independent records. That is a fact about the ground, not a modelling choice made here; the reader who rejects it rejects the cost of a recognition, not this theorem.

Theorem 5.2 (Objectives stated in what parts hold are projections)**.** *Every objective that is a function of the parts' positions or of the from-whom-to-whom exposures matrix is an exposure functional in the sense of Definition 6.1. By Theorem 6.2 it takes the same value on two systems that differ on who is reached by a rewrite and who consents to it. It cannot see harm.*

So the objective of economics is settled by the ground. The one participant selects for itself, which is to say for the total disagreement of the whole, and the moral content of that objective is not added to it: an objective blind to harm is not the whole's objective because the whole is the one to whom the harm is done. The parts' objectives (their own prosperity, the poverty of some, the growth of a total) are what the one participant cannot see, and Theorem 3.4 says those are exactly the chosen variables.

6 Exposure functionals: what amounts cannot see

Balance-sheet reasoning works from positions; network finance works from the from-whom-to-whom matrix of exposures. Both read amounts. We say exactly what they measure and name the datum they cannot record.

Definition 6.1 (The exposures projection)**.** The *exposures matrix* of a chart and books assigns to each ordered pair of parts the total booked claim of the first on the second (summing over relationships joining the pair); positions are row sums minus column sums. A functional of (chart, books) is an *exposure functional* if it depends on the data only through the exposures matrix.

A relationship carrying zero balance contributes nothing to the matrix: an open line with nothing outstanding is invisible to it.

The pair. Four parts. In both systems part 0 is owed one unit by part 1 and nothing else is outstanding. In A the standing relationships are $0 \to 1$, $1 \to 0$, $2 \to 3$, $3 \to 2$: two two-cycles. In B they are $0 \to 1$, $1 \to 2$, $2 \to 3$, $3 \to 0$: one four-cycle. The exposures matrices are identical (a single 1 at $(0, 1)$) and so are the positions $(1, -1, 0, 0)$. Under the comparison cost the fixed-position rewrite of A is $(\frac{1}{2}, -\frac{1}{2}, 0, 0)$: the debt is split across the antiparallel pair and parts 2, 3 are untouched. The rewrite of B is $(\frac{3}{4}, -\frac{1}{4}, -\frac{1}{4}, -\frac{1}{4})$: a quarter of the debt is routed through 2 and 3. Part 2's half-share charge is $\frac{1}{32}$ in B and 0 in A; it consents (its charge does not rise) in A and not in B. For every admissible cost with rest on the empty relationships and off the debt, the two rewrites differ: in B the circulation around the four-cycle at the booked state is $c'_{01}(u_{01}) \neq 0$ (Theorem 4.7(c)), in A part 2's two-cycle is disconnected from the debt and stays at rest. A cost whose rest on $0 \to 1$ is the booked figure makes both systems optimal as booked, which is why the hypothesis is stated.

Theorem 6.2 (Exposure functionals are cycle-blind)**.** *Every exposure functional takes the same value on any two (chart, books) with the same exposures matrix. On the pair every exposure functional agrees, and for every admissible cost with rest on the empty relationships and off the debt the cheapest rewrite, the set of parts it moves, the half-share bystander charge and the consent verdict all differ. None of the four is an exposure functional.*

Proof. The pair has the same matrix; the four quantities differ on it (computed above for the comparison cost; for every admissible cost by Theorem 4.7). A functional that differs on two inputs with the same projection does not factor through it. □

Theorem 6.3 (Where amounts are enough)**.** *On a fixed chart the exposures matrix determines every position and the fibre, hence the rewrite. On a chart whose standing graph is a forest the rewrite is determined by the matrix outright.*

An instrument that reads amounts is right about everything amounts determine: positions, the fibre, and on a forest the rewrite itself. It is silent about who is reached when a cycle of obligations is rewritten and whether a bystander consents; the datum it cannot record is the set of standing relationships that carry no balance, and Theorem 4.7 says that datum matters through the circulation around the cycles it closes. A model that also observes the standing graph is not blind to it, and the theorem says what that observation buys.

7 Ethics is the same subject

The ethics of recognition defines, on the ledger of bonds between parts, the harm an action does to a part (the cost it externalises onto that part's bonds), the part's consent (its burden does not rise), and moral debt (the cost one part's action leaves on another). These are not imported from outside; they are the one participant's accounting of its separations, and the question is only whether their operators coincide with the economic ones.

Definition 7.1 (Translation)**.** A chart, books and a fixed-position rewrite determine an action on the bond ledger whose agents are the parts and whose bonds are the relationships: the bond change on each relationship is the rewrite on it.

Theorem 7.2 (One operator, two names)**.** *Under the translation:*

(a) the harm to a part equals twice the change in its half-share incident charge (the ethics posts the full cost of a bond at each end; the ledger posts half); on the triangle in the log reading the part who owed nothing suffers harm $2J(\sqrt{3/8}) > 0$*;*

(b) a part consents iff its incident charge does not rise; a part consents to every fixed-position rewrite iff every relationship incident to it is a bridge, so a zero-balance relationship can change the consent verdict only by closing a cycle, and on the pair of Section 6 it does;

(c) harm to a part on none of the touched relationships occurs only when the touched relationships close a cycle through it, never on a forest;

(d) the moral debt created by closing a gap equals the gap in amounts, however the two revaluations split it.

Theorem 7.3 (Where the vocabularies diverge)**.** *The economic bystander (a part who owed nothing and is charged) is not the part the ethics calls internalised; the two notions differ on the triangle. The ethics carries objects the translation does not use (a value functional, patienthood, energy, a tick); the ledger carries objects the ethics does not (the unit, the fibre, accounts, journal entries).*

Harm, consent and debt; bystander charge, non-increase and counterpart gap: one operator, two names, exact up to the factor of two in how a bond's cost is posted. Economics is the ethics of recognition, quantified, and the objective of Section 5 is the moral objective because the one to whom harm is done is the one who keeps the record.

8 What the parts choose

The modelling choices that any derivation of economics must make are listed here, each with the theorem stating what holds for every value of it. The pattern is the one Section 3 predicts: the choices do not become theorems, they become free variables each carrying a theorem about what no value of it can change.

Choice	Status	What holds for every value
Cost of a disagreement	Ground 1.2	Derived from the one participant's two properties (Theorem 1.4); unit κ and exponent λ change no ordering, minimiser or sign.
Carrying cost per tick	Ground 9.1	Theorem 9.6: a standing gap costs $NJ(r)$ over N ticks and neither part reads it on its own books. The per-tick charge is inherited; the cost's form is free within the admissible class.
Shares of a cost	Chosen	Theorem 4.5: the total is forced.
The agreement move	Chosen	Theorem 4.1: the split is free, the sum is the gap.
Walk-away figures	Chosen	Theorem 4.2: the ledger supplies two figures; the outside option is free.
Point in the interval	Chosen (known)	Remark 4.3.
Characteristic function	One candidate	Theorem 4.12, 4.13: books outside the core under both.
Kept relationships	Chosen	Theorem 4.17: dynamics stops exactly at pairwise-stable kept-sets, for every cost.
Rent schedule	Chosen	Theorem 4.20.
Split rate of wealth	Bounded, not fixed	Theorem 4.15 fixes the exponent; Theorem 9.7 bounds the rate by one event per tick.

9 Genesis: how one becomes two

Every result above holds on given books with given buffers. This section derives the form of those givens from the ground and names what it does not derive. Two facts of the ground are used, stated here as the premises they are.

Ground 9.1 (Per record, per tick)**.** Cost adds over independent records (Ground 1.2): accumulating a record that shares no content with what is already held, and that costs something, strictly increases the cost. At most one unit moves per tick, and a book records a unit while the unit is in its cell.

Theorem 9.2 (Two records)**.** *Two sides' records of one shared recognition, sharing no content and each costing something, do not merge: their join is neither one of them. Separation produces two records, not one shared datum.*

Proof. If the join of the two records were the first record back again, its cost would equal the cost of the first; by Ground 9.1 the cost of the join is strictly larger. The same for the second. □

Theorem 9.3 (Never simultaneous)**.** *A crossing of one unit from the cell of book α to the cell of book β is recorded by α at tick t and by β at tick $t + 1$. At no tick do both books record it. Read as books on the one-claim chart, the crossing is the one-sided state at t (source posted, destination not) and the opposite one-sided state at $t + 1$; it is never matched.*

Proof. At one tick a unit is in one cell, so two books that both record it are the same book. With $\alpha \neq \beta$, β does not record at t and α does not at $t+1$. On the one-claim chart the source's asset carries x and the destination's liability 0 at t, and 0 and $-x$ at $t + 1$; neither sums to zero when $x \neq 0$. Both states balance inside each part. □

Theorem 3.7 said, negatively, that the one record of physics does not hold two books. Theorem 9.3 says the positive thing: a matched pair of books is a relation between *kept* records with different stamps, never a state of the one record. The two-level ledger is the memory of the one record; parts must keep records to match at all, and the one-sided state is the state every recognition passes through. This is where the two records come from.

Theorem 9.4 (The buffer is the residue of separation)**.** *For a part whose books balance, the total on its undesignated accounts is minus its external net. A part with a nonzero net holds a nonzero balance on some undesignated account. When every part balances, the undesignated totals of all parts sum to zero exactly when the counterpart gaps cancel in aggregate.*

Proof. The designated accounts of a part correspond one to one with the relationships incident to it, an asset account per relationship leaving and a liability account per relationship entering, no account both. So the designated total is the net; balance makes the rest its negative. Summing over parts, the nets sum to the sum of counterpart gaps. □

The buffer that absorbs a revaluation is not chosen and not supplied from outside. It is what separation leaves in each part: the part's record of being still one with the rest. The buffers add to zero exactly when the whole still adds up.

Theorem 9.5 (No buffer, no first claim)**.** *A part every account of which is designated cannot post: its books are unchanged along every legal history.*

Proof. A journal entry avoids designated accounts, so on such a part it is zero. A cross posting needs an undesignated counter account at each end, so such a part is neither end. Induct along the history. □

Theorem 9.6 (Carrying cost)**.** *Assume, with Ground 9.1, that a standing record is charged at every tick it stands. Then a gap at ratio $r \neq 1$ costs $NJ(r)$ over N ticks, which is positive and exceeds every bound as N grows; and each book compared with itself reads $J(1) = 0$.*

Proof. The sum of N copies of $J(r)$; $J(r) > 0$ for $r \neq 1$; $J(1) = 0$. □

The carrying cost is a one-record functional neither part sees on its own books. Interest (Section 4) is the parts' agreement about who books a cost the whole already carries.

Theorem 9.7 (Rate bound)**.** *At most one posting occurs in the whole ledger per tick; over N ticks at most N. The split rate of Theorem 4.15 is at most one event per tick.*

Proof. Ground 9.1: at most one unit moves per tick. □

Not yet derived 9.8 (The amount and the rate)**.** The ground supplies the crossing and its two stamps, not the magnitude of the unit that crosses. The amount of the first claim is not derived here; it is the same wall as the emission-side absolute of physics. The split rate is bounded by one event per tick and not fixed below that bound.

10 What each result stands on

The ground has three parts and the calculus has one primitive. A reader who accepts separation but not the ground can keep every result in the first two rows; the third row is where the ground enters.

Stands on	Results
Separation only (two records, balanced books, no cost)	Reachability 2.5, fold 2.6, unit 2.7, matched 3.7, gap conservation 4.1, two figures 4.2, exponent one 4.15, amounts are blind 6.2 and right 6.3, buffer 9.4, frozen 9.5.
Separation plus any admissible cost (differentiable, strictly convex, at rest)	Swap 3.2, forced and chosen 3.3, 3.4, signed gap 3.5, shares 4.5, exact form 4.6, cut space 4.7, bystander 4.9, firm 4.12, freed 4.13, pruning 4.17, rent 4.20, dictionary 7.2, divergence 7.3. Numbers quoted are at J.
The ground	J from no favourite and same law 1.4 (Ground 1.1); strain 3.6 and the objective 5.1, 5.2 (Ground 1.1, additivity of Ground 1.2); interest 4.22 (Ground 1.3); two records 9.2, never simultaneous 9.3, carrying cost 9.6, rate 9.7 (Ground 9.1).

Economics stands on Recognition Science as arithmetic stands on the first distinction: not from nothing, from the ground. What this document adds to the ground is one primitive, separation, and the theorems in the first two rows are what separation forces with no further input.

11 Related formalisms

The calculus of Section 2.1 is double-entry bookkeeping made algebraic: the part-level balance is the accounting identity and the cross posting is the quadruple entry of a bilateral claim; Ijiri [1] and Ellerman [2] treat double entry algebraically with one set of books and so have no counterpart gap. The fixed-position rewrite is a flow on a graph with a separable convex cost, after Rockafellar [3]; its cut-space dependence is the decomposition of edge space into cut and cycle spaces [4]. Network clearing after Eisenberg and Noe [5] takes the exposures matrix as input; Section 6 states what that input cannot identify. Portfolio compression after D'Errico and Roukny [6] is the fixed-position rewrite of a cycle of obligations in practice, and whether it can harm a third party is the question of Veraart [7] and Schuldenzucker and Seuken [8]; the bystander theorem and Remark 4.10 answer it for the ledger. The two-figure interval is the assignment market of Böhm-Bawerk [9] and Shapley and Shubik [10]. Pruning and pairwise stability follow Jackson and Wolinsky [11]; the firm game is a cooperative game with the core of Gillies [12]. The partition result is the conserved route to Zipf's law, as against Gabaix's proportional random growth [13]. The functional equation of Section 1.1 is d'Alembert's, treated in Aczél [14].

References

[1] Y. Ijiri, *Theory of Accounting Measurement*, American Accounting Association, 1975.

[2] D. Ellerman, The mathematics of double entry bookkeeping, *Mathematics Magazine* 58 (1985) 226–233.

[3] R. T. Rockafellar, *Network Flows and Monotropic Optimization*, Wiley, 1984.

[4] N. Biggs, *Algebraic Graph Theory*, 2nd ed., Cambridge University Press, 1993.

[5] L. Eisenberg, T. H. Noe, Systemic risk in financial systems, *Management Science* 47 (2001) 236–249.

[6] M. D'Errico, T. Roukny, Compressing over-the-counter markets, *Operations Research* 69 (2021) 1660–1679.

[7] L. A. M. Veraart, When does portfolio compression reduce systemic risk?, *Mathematical Finance* 32 (2022) 727–778.

[8] S. Schuldenzucker, S. Seuken, Portfolio compression in financial networks: incentives and systemic risk, working paper, 2020.

[9] E. von Böhm-Bawerk, *Positive Theorie des Kapitales*, Wagner, Innsbruck, 1889; Book IV.

[10] L. S. Shapley, M. Shubik, The assignment game I: the core, *International Journal of Game Theory* 1 (1971) 111–130.

[11] M. O. Jackson, A. Wolinsky, A strategic model of social and economic networks, *Journal of Economic Theory* 71 (1996) 44–74.

[12] D. B. Gillies, Solutions to general non-zero-sum games, in *Contributions to the Theory of Games IV*, Princeton, 1959.

[13] X. Gabaix, Zipf's law for cities: an explanation, *Quarterly Journal of Economics* 114 (1999) 739–767.

[14] J. Aczél, *Lectures on Functional Equations and Their Applications*, Academic Press, 1966.