The Organization of the Investment Frontier

Independent Mandates and the Formation of Shared Investments

Prepared for John Komkov

Polycentric Finance research program

Working paper · October 2026

Abstract

Investment opportunities can depend on investors' willingness to act together. Under threshold participation and scale economies, improving one investor's outside alternative raises worst-equilibrium welfare, narrows the equilibrium range, and can eliminate a jointly valuable investment. We derive sharp welfare bounds and distinguish participation transfers from assurance. Divisible portfolios can retain formation cliffs; congestion can reverse compression. In 64 simulated worlds, persistent records reduce repeated retrieval, while service failures withhold valid participation. Shared processing has lower recorded costs, but the marginal recorded reduction favors separate processing. These measured differences inform participation conditions without establishing economically significant investment fragmentation.

Keywords: portfolio formation; cost sharing; investment mandates; participation; authority; artificial intelligence.

JEL: G11, G23, C72, D47, D85.

Research draft. No journal or institutional affiliation is implied.

Disclosures and research record

This working paper was prepared for John Komkov in connection with the development of Argosy and Megalith. Those commercial interests motivated the research question. Company proposals and implementation implications appear in a separate note; they do not enter the mathematical results. The title page identifies the commissioning principal, rather than asserting an approved authorship or institutional affiliation. A submission would require the responsible human authors to approve the manuscript and complete their individual disclosures of affiliations, financial interests, funding and assistance.

AI systems assisted with research, drafting, code, numerical checks and internal review. Two separate review agents examined the financial argument and the AI study. These were internal research reviews, not external academic peer review. All formal claims remain the responsibility of the eventual human authors. The model evaluation records the product identifier, settings, access timestamps, prompts and returned usage. The available identifier does not guarantee an immutable model-weight version.

The numerical economies are synthetic. Published sources establish institutional capabilities or motivate comparisons; they do not estimate the prevalence or magnitude of the modeled losses. The interpretation experiment uses executable synthetic ground truth and a separate public-document slice with agent-audited exploratory labels. No private investor documents, live accounts, financial commitments or external outreach were used.

The accompanying package contains the canonical manuscript, generated reading and review editions, native-equation Word edition, financial model, evaluation inputs and logs, source and parameter ledgers, independent reviews, and the response record. Previous editions are preserved. Monetary units and status labels are specified at each result.

1 Introduction

An improvement in what an investor can do alone can remove from equilibrium an investment that remains worth undertaking together. The result arises when a program needs several investors to form and each investor decides whether participation serves his own portfolio. A better outside use of capital can cause one investor to leave. His departure can make the remaining investors' terms unattractive or leave the program below its minimum scale. The investment disappears from their equilibrium opportunities even though the resources and technology needed to implement it remain available.

We establish this result for a threshold participation game with excludable benefits and scale economies. An increase of Δ in one principal's outside value raises the lowest aggregate welfare among pure equilibria by exactly Δ. It never enlarges the gap between the highest and lowest equilibrium welfare. If the original gap is G, the change in the highest equilibrium welfare lies between $\Delta - G$ and Δ. Both bounds are attained. Thus an improvement can raise the floor while lowering the ceiling. We call the adverse case *excess separation*: independent choices abandon a shared investment whose aggregate gain still exceeds the available outside alternatives.

The distinction is organizational rather than technological in the narrow sense. An investment can be shared without making its owners hold the same portfolio. Each owner may combine the investment with common funds, separate holdings and other programs. The unit of a joint formation decision can therefore be an investment or a contractual arrangement, while authority remains with several independently governed portfolios. Conventional portfolio-choice problems take a specified menu as an input; financial-innovation research already explains why that menu can be endogenous. Our question concerns a particular part of that endogeneity: how separately governed demand supports formation, and how a change in individual implementation costs changes that support.

Three attainable sets organize the analysis. The feasible set describes what resources, technologies and hard constraints permit. The equilibrium set describes what the baseline participation game supports. The implementable set describes what a specified class of permissible contracts can support after their costs are paid. Their welfare envelopes can diverge. A calculation of a desirable joint allocation does not establish voluntary participation, and a participation agreement does not ensure that every equilibrium implements it. Transfers address a distribution problem; assurance payments can address a selection problem; collectible funding addresses a resource problem.

Artificial intelligence makes the incidence of organizational costs a practical question. Interpretation and record maintenance may become cheaper for an investor acting separately. The same technology may make shared arrangements cheaper to administer while preserving each investor's obligations, tax attributes, constraints and authority. For a fixed proposed membership, what matters is the difference between those cost reductions at each participant, especially the participant closest to leaving. Average processing savings do not determine the sign. We measure the tested systems in matched separate and shared processing modes and evaluate their implica-

tions for that participation condition under declared allocations of shared processing costs. In 64 simulated worlds, persistent records reduce repeated retrieval, while validation mainly standardizes how an already-abstaining decision is encoded. Every successful primary context extracts the scored state fields correctly. Recorded token costs fall more separately, although shared processing has a lower recorded cost level. Service failures can withhold valid participation, especially when one response serves several principals. The evidence establishes a processing tradeoff and a semantic-accuracy ceiling, not observed investment fragmentation.

The compression result draws on the monotone membership logic of cost sharing. In particular, the union-closed construction of the largest individually willing served set is explicit in Juarez (2013). We provide a direct welfare-envelope proof for the stated threshold game, sharp examples, a financial interpretation and limits. Mathematical validity and priority are separate claims; the paper does not claim a new general principle of group formation. Its contribution is to connect the participation margin, permissible implementation and the incidence of interpretation costs in a common research design.

Several findings keep that design economically substantive. Continuously divisible mixed portfolios can retain a formation cliff. Two programs can compete for limited capital even when investors may hold both: making one program cheaper can lower the best equilibrium welfare while raising coordinated welfare. Congestion can reverse compression, because an exit can improve the terms available to remaining investors. Finally, a credible failure payment need never be paid in equilibrium, yet the reserve behind it consumes resources. These results identify which institutional changes would be useful and which would leave the problem untouched.

2 Portfolio choice and the organization of implementation

2.1 What existing theory supplies

Portfolio theory evaluates combinations of claims rather than investments in isolation. Separation and spanning explain when heterogeneous investors can obtain their preferred exposures through mixtures of common funds. Markowitz (1952), Merton (1971) and Ross (1978) supply distinct foundations for that analysis. They also impose a demanding null: a difference in an investor's stated preference is not evidence that he needs a new vehicle or a separate transaction. The relevant difference is one that changes an attainable payoff, a binding constraint, a cost or a permissible action.

The design of claims is already an object of financial economics. Duffie and Rahi (1995) review security design and financial innovation; Boyd and Prescott (1986) model intermediaries as coalitions issuing claims different from those they acquire. The present model holds gross payoffs and investment information fixed in its main comparison. It changes the organization through which an opportunity can be formed and funded. It does not identify excess returns from a superior forecast or originate the idea of an endogenous opportunity set.

Joint optimization of separate portfolios is also established. Iancu and Trichakis (2014) study efficiency and fairness when accounts share trading costs, including individual participation restrictions. Their framework includes extensions with fixed fees and mixed-integer decisions. Devanathan et al. (2026) coordinate transaction-cost mitigation across portfolio managers. Luxenberg et al. (2022) study dynamic allocation with illiquid commitments. Our participation problem is different from merely computing a joint optimum: investors can refuse an allocation and pursue outside opportunities, and those decisions can determine whether the proposed investment exists.

Clubs and differentiated products connect heterogeneity to shared fixed costs. Buchanan (1965) analyzes membership and the cost of a shared good. Spence (1976) and Mankiw and Whinston (1986) explain why private entry incentives can fail to internalize effects on other participants. In our setting, separate implementation can remove capital from a valuable common arrangement. A new differentiated arrangement can also attract capital and create surplus. The sign cannot be inferred from the number of differences that an institution preserves.

The closest structural connection is cross-monotone cost sharing and population-monotonic allocation (Sprumont, 1990). The existence of a scheme that allocates coalition value while protecting incumbents is distinct from our assumption that a particular program already has monotone participation payoffs. Juarez (2013) characterizes mechanisms under different treatments of indifference and explicitly constructs the largest willing coalition using union closure. We use the same order property in a complete-information participation game, retain all pure equilibria and derive bounds on their welfare envelope as outside values change. We do not infer truthful reporting from the existence of that envelope. Combinatorial cost sharing studies the more demanding problem in which agents demand bundles and costs interact across goods (Birmpas et al., 2019). That distinction matters for the overlapping-program extension. In risk pooling, Blier-Wong and Lauzier (2026) examine entry-monotone allocations and core stability under cash-additive risk measures. Their convex pooling problem provides a close financial antecedent; discrete formation and the equilibrium response to outside opportunities remain separate questions.

The institutional response belongs to the literature on contracting with externalities. Dominant assurance contracts are due to Tabarrok (1998); optimal differentiated recruitment and strategic uncertainty are studied by Segal (2003), Bernstein and Winter (2012) and Sákovics and Steiner (2012). Their results depend on particular information structures and contract classes. We use explicit contracts whose acceptance game can be checked, rather than presume that a calculated external cost is a collectible charge.

2.2 Which differences require separate implementation?

Three kinds of differences are useful to distinguish, although they are not mutually exclusive. First, investors can want different return exposures. If existing funds span those exposures at acceptable cost, investors can choose different mixtures while sharing execution. Second, investors can face different restrictions on holdings or actions. A long-only mixture cannot eliminate a prohibited security if every fund used in the mixture holds it, but direct holdings, an eligible fund, a sleeve or

an appropriate contractual arrangement may solve the problem. Third, investors can have different histories: acquisition basis, legacy liabilities, accepted commitments and prior grants of authority. Those facts must remain attributable to the relevant principal through time.

Shared vehicles can preserve some differentiated state. Partnership accounting is an institutional antecedent, not a contradiction of the argument. Capital accounts, allocations concerning contributed property and partner-specific basis records can coexist with a common investment. They are distinct records: a partner's capital account is not necessarily his outside tax basis, which can reflect liabilities and partner-level adjustments (Internal Revenue Service, 2025a,b). The empirical question is what it costs to maintain the relevant distinctions accurately and what outcomes the legal arrangement permits.

Accordingly, the comparison is not between complete individuality and a single undifferentiated pool. The capable conventional benchmark may use fund mixtures, separate accounts with shared execution, differentiated participation terms or a custom vehicle. If it can reproduce the same net allocation with the same authority, information and costs, the proposed coordination arrangement has no additional portfolio value. Its possible advantage then concerns delivery cost, reliability or access to a demonstrably different opportunity.

3 The environment and attainable sets

3.1 Economic state, choices and authority

There are finitely many principals, three decision dates and finitely many states. At date 0 each principal adopts a mandate, chooses grants of authority and may enter a contract. At date 1 an opportunity and a funding state are observed; admissible participation and implementation decisions are made. At date 2 claims settle. A date 1 decision uses only information available by that date. Existing wealth, liabilities, unfunded obligations and liquidity requirements enter the principal's book before a new investment is evaluated.

A principal is a holder of capital with its own decision authority. The holder may be an individual, a trust, an investment office acting for a specified principal, or an institution. Multiple accounts representing the same economic principal do not create multiple budgets. A fund may itself be one independently governed participant. Polycentricity refers to these distinct centers of authority, following the governance tradition synthesized by Ostrom (2009). It does not require distributed computing, a blockchain or a particular form of ownership.

Let $X_i(a, \omega)$ be terminal wealth under joint action a and state ω. The baseline financial representation uses CARA certainty equivalents,

$$\mathrm{CE}_i(X) = -\frac{1}{\alpha_i} \log \mathbb{E}[e^{-\alpha_i X}], \qquad \alpha_i > 0. \tag{1}$$

A sure terminal transfer t satisfies $\mathrm{CE}_i(X + t) = \mathrm{CE}_i(X) + t$. This makes certainty-equivalent dollars a convenient numeraire. A transferable-utility representation additionally requires that the

relevant transfers actually be admissible, financeable and enforceable. Binding liquidity limits or restrictions on transfers can prevent a jointly valuable allocation from being acceptable. CRRA preferences permit money transfers but lack this exact translation property.

An authority grant A_i is a set of contingent actions permitted without a new approval. A mandate supplies objectives and hard investment constraints; an authority grant supplies permission; a signed contract supplies obligations. These are separate objects. Revoking future authority can block a new commitment without canceling a call under an already accepted obligation. Nor does an unused grant establish that cash is available or that a counterparty will accept a trade.

Prospective programs require actual formation expenditures or minimum capital, can impose minimum tickets and may have scarce capacity. Claims are represented by their contingent cash flows, obligations and rights. A common label does not establish that two instruments have the same payoff or that an access right can be transferred. Investors may combine several programs with existing funds and individual holdings. The reduced participation model below isolates one formation margin after those surrounding portfolio choices have been made. Prescribed linear CE schedules in the examples are reduced-form primitives; CARA preferences do not generally imply them for correlated risky investments.

3.2 Sets before frontiers

Let $\mathcal{F}(z)$ be the set of net payoff vectors attainable under financial state z, technologies, resources and hard constraints. Real formation, administration and financing costs are included. The set does not yet require that a decentralized participation game support its allocations. Let $\Gamma_0(z)$ be the baseline game, and let $\mathcal{E}(z)$ be the payoff vectors of its pure equilibria. For a declared class of permissible contracts C, define

$$\mathcal{I}(z; C) = \bigcup_{c \in C} \{p(a; c, z) : a \in \mathrm{NE}(\Gamma_c(z))\}. \tag{2}$$

The costs and resources of each contract belong in p and in its feasibility conditions. If the class contains the baseline contract and every contract respects the same resource accounting, then

$$\mathcal{E}(z) \subseteq \mathcal{I}(z; C) \subseteq \mathcal{F}(z). \tag{3}$$

The feasible, equilibrium and implementable frontiers are the undominated parts of these respective sets. Their frontiers need not themselves be nested. An outcome undominated in a smaller set can be dominated after more arrangements become available.

For figures we fix the welfare measure $W(p) = \sum_i p_i$, in certainty-equivalent dollars, and compute scalar envelopes $\max_{p \in \mathcal{E}} W(p)$, $\max_{p \in \mathcal{I}} W(p)$ and $\max_{p \in \mathcal{F}} W(p)$. These are not mean–variance frontiers. An increase in their sum does not establish that every principal gains. We report member payoffs and test participation against jointly feasible outside alternatives.

Weak supportability and robust implementation are different. Equation (2) asks whether a contract has an equilibrium supporting an outcome. A designer concerned about every equilibrium instead evaluates, for example,

$$\sup_{\substack{c\in C\\ \mathrm{NE}(\Gamma_c)\neq\varnothing}} \inf_{a\in\mathrm{NE}(\Gamma_c)} W(p(a;c)). \tag{4}$$

This quantity need not equal the maximum over $\mathcal{I}$. A unanimous subscription can make an attractive allocation an equilibrium while leaving no-formation equilibria intact. An assurance contract may remove those equilibria if its promises are credible and participation remains attractive after success.

4 Compression under scale economies

4.1 The participation game

There are n principals. Each chooses to subscribe or remain outside. Write S for the subscribers. The program forms if $|S| \geq q$, where $2 \leq q \leq n$. A nonparticipant i obtains o_i, independently of S. A failed subscriber obtains $o_i - h_i$, with $h_i \geq 0$. Thus a failed subscription either restores the outside value or consumes a stated holding cost; it pays no failure bonus. A member of a formed program obtains $u_i(S)$, inclusive of implementation and holding costs.

The central assumption is inclusion monotonicity:

$$S \subseteq T, \quad i \in S, \quad |S| \geq q \quad \Longrightarrow \quad u_i(T) \geq u_i(S). \tag{5}$$

Program payoffs and terms do not depend on the outside-value vector. Any outsider can join a formed set. There is no endogenous rationing or admission veto. For example, $u_i(S) = b_i - Fw_i/\sum_{j\in S} w_j - h_i$, with fixed positive weights and fixed $F \geq 0$, satisfies (5). A falling cost share is not sufficient if an additional participant also dilutes the member's investment payoff; the assumption concerns the net payoff.

Let $E(o)$ be the pure Nash equilibrium correspondence. It is nonempty because all principals remaining outside is an equilibrium: one subscriber cannot meet the threshold. Define

$$L(o) = \min_{S\in E(o)} W(S;o), \quad H(o) = \max_{S\in E(o)} W(S;o), \quad G(o) = H(o) - L(o). \tag{6}$$

No selection rule is imposed between the best and worst equilibria.

Theorem 1 (Compression)**.** *Under the stated threshold, exclusion, failed-subscription and monotonicity assumptions, let $o' = o + \Delta e_k$ for $\Delta \geq 0$. Then*

$$L(o') = L(o) + \Delta, \qquad G(o') \leq G(o), \qquad \Delta - G(o) \leq H(o') - H(o) \leq \Delta. \tag{7}$$

Both bounds on the change in highest equilibrium welfare are sharp.

Proof. All-outside gives $\sum_i o_i$. In a formed equilibrium each member can exit and obtain o_i, so every member receives at least his outside value. In a nonforming equilibrium a subscriber with $h_i > 0$ would leave; subscribers who remain incur zero holding cost. Hence every nonforming equilibrium also gives $\sum_i o_i$, and $L(o) = \sum_i o_i$.

Take any formed equilibrium T under o'. At the old outside values its members remain willing to participate. Admit an outsider whenever his payoff from joining is strictly above his old outside value. Existing members' payoffs cannot fall by (5); each admitted member is willing to remain. The finite process ends at a formed old equilibrium $S \supseteq T$: no remaining outsider has a profitable entry, and no member has a profitable exit.

The members of T have weakly larger payoffs at S, the additional members have nonnegative slack above their old outside values, and the old outside values are no larger than the new ones. Therefore

$$W(S; o) - \sum_i o_i \;\geq\; W(T; o') - \sum_i o'_i. \tag{8}$$

Maximizing over new formed equilibria gives $G(o) \geq G(o')$; if none forms, $G(o') = 0$. Adding the change in the floor gives the upper bound in (7). The lower bound follows from $H(o') \geq L(o') = L(o) + \Delta$.

For sharpness of the lower bound, take two principals, threshold two, formed payoffs $(0.5, 10)$ and outside values $(0, 0)$. Raising the first outside value to 0.6 destroys formation. Highest welfare changes from 10.5 to 0.6, exactly $\Delta - G$. For the upper bound with a positive gap, let a third principal always prefer to remain outside a valuable two-member program. Increasing his outside value by Δ leaves that program unchanged and raises highest welfare by exactly Δ. □

The theorem concerns the range of equilibrium welfare, not a universal decline in its maximum. With an unchanged membership, outsiders can become better off while the formed program remains intact, so the highest equilibrium rises. At an exit threshold, the membership and the cost shares can change discontinuously. Figure 1 displays both effects without choosing favorable equilibria on one side of the change and unfavorable ones on the other.

4.2 The largest viable membership

A formed set is individually willing if $u_i(S) \geq o_i$ for every member. The family of such sets is closed under union: any member of either set remains at least as well off in their union by (5). If any willing formed set exists, its union $S^*(o)$ is therefore the largest one. It is an equilibrium. An outsider who could profitably enter would create a larger willing set, contradicting maximality.

The payoff vector at $S^*(o)$ weakly dominates that at every other pure equilibrium. Members of another formed equilibrium receive at least as much in S^*; additional members receive at least their outside values. Thus S^* attains $H(o)$. An increase in outside values removes willing sets and cannot add them, so $S^*(o') \subseteq S^*(o)$ whenever both exist. This gives a useful algorithm: begin

with all principals, delete those who are unwilling at current terms, and repeat until the remaining set is willing or falls below the threshold.

This construction is the direct link to the cost-sharing literature. The largest willing served set and its monotonicity are established features of cross-monotone mechanisms (Juarez, 2013). Our participation game does not inherit those mechanisms' reporting incentives merely because its equilibrium has the same order structure. A principal may manipulate his declared outside value, and the computed willingness test is not a truthful elicitation procedure.

4.3 Excess separation and the distribution band

In the sharp lower-bound example, the program still yields 10.5 after the outside improvement, while all-outside yields only 0.6. A transfer slightly above 0.1 from the second member to the first would restore individual willingness if permissible. The loss of formation is a distribution problem, not a disappearance of resources or gross investment value.

For a fixed membership S, define initial slack $s_i = u_i(S) - o_i$. If only member k receives an outside improvement Δ, he ceases to be willing when $\Delta > s_k$. The same allocation continues to create aggregate surplus relative to the new outside alternative while

$$s_k < \Delta < \sum_{i \in S} s_i. \tag{9}$$

Within this interval, individual exit can destroy positive total surplus. A permitted redistribution may preserve it. With other members' outside values also changing, the aggregate comparison uses the sum of their changes. With potential entrants or new investment choices, the fixed-membership thresholds must be recomputed rather than treated as global cutoffs.

The term excess separation describes this comparison with a jointly attainable, individually acceptable alternative under a stated contract class. It does not imply that an investor has too much authority or that an institution may compel him to remain. In particular, a program with positive aggregate surplus can remain unimplementable when transfer restrictions, legal exclusions or missing funding prevent compensation.

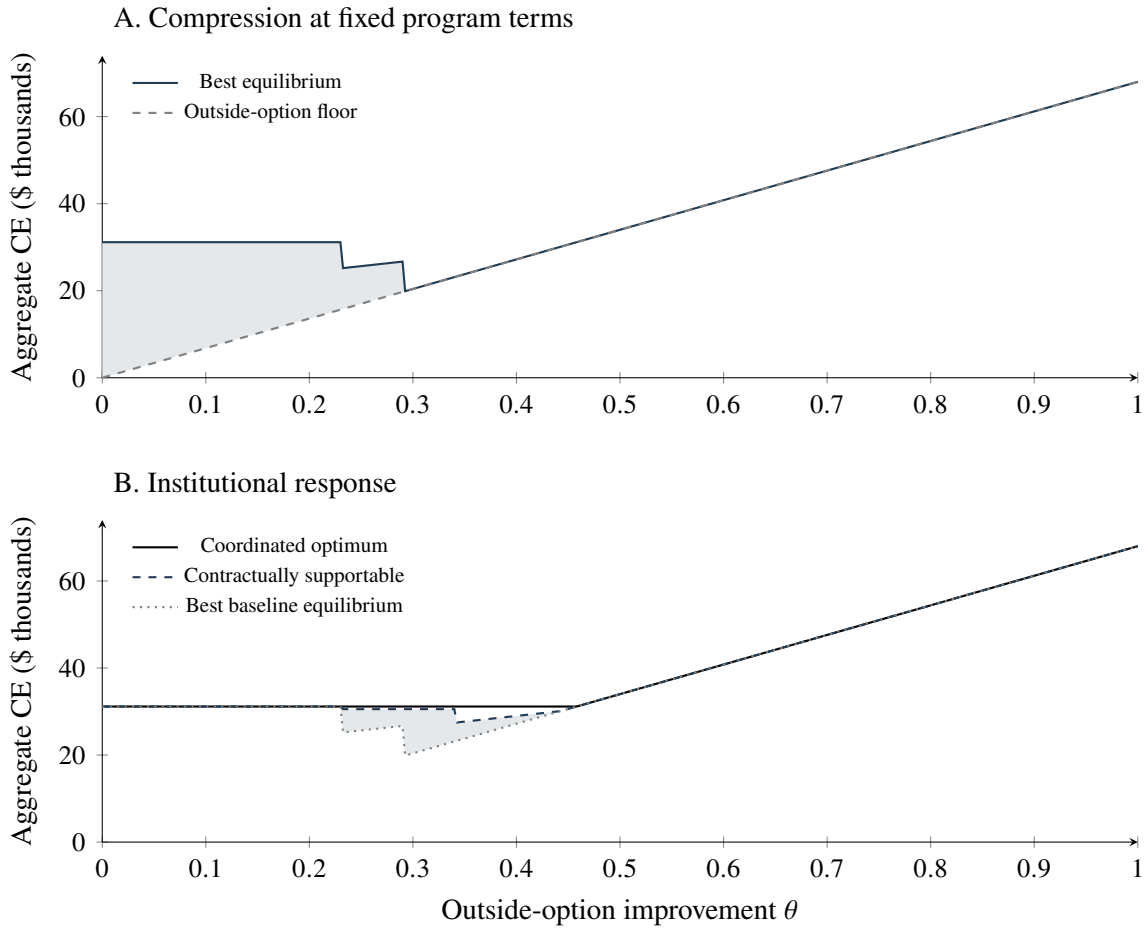


Figure 1: Compression and institutional response. Five principals share a program with minimum membership three and fixed formation cost 30. Outside values rise with θ; gross participation values and program cost remain fixed. Panel A shades the range between extreme pure equilibria. Panel B shades the gap recoverable by the declared named, all-accept contract class, after administration costs and bounded balanced transfers. Support in one equilibrium is sufficient for that curve; formation in every equilibrium is not asserted. Curves connect 401 computed parameter values; exact discontinuities lie between adjacent samples. All quantities are synthetic CE gains in thousands of dollars.

5 Contracts, participation and credibility

5.1 Cost incidence at the pivotal principal

Write a member's formed payoff as $b_i - \phi_i(S)$, where ϕ_i includes his real implementation burden and any fixed allocation of shared costs. Consider changes Δo_i in outside values and ϕ_i' in those burdens, holding the proposed membership and gross values fixed. Individual willingness is preserved exactly when

$$\Delta o_i \leq s_i + \phi_i(S) - \phi_i'(S) \quad \text{for every } i \in S. \tag{10}$$

This is a member-by-member condition. It says nothing about whether an outsider wishes to enter or whether the resulting allocation remains coalition-stable. An entrant who reduces charges can keep a member participating beyond the fixed-set cutoff. A change in supplier pricing can consume the implementation saving before it reaches members.

Let d_i^{sep} be the reduction in the cost of a specified outside route and d_i^{sh} the reduction in the member's cost within the shared arrangement. If these changes alone affect the comparison, (10) becomes

$$d_i^{\text{sep}} - d_i^{\text{sh}} \leq s_i. \tag{11}$$

The relevant incidence is the net gain from separation relative to sharing. A uniform proportional technological improvement need not produce uniform dollar savings. But a claim that the most distinctive mandate is necessarily pivotal requires an additional relationship between complexity, costs and existing slack. We test cost incidence by measured complexity and leave that correlation as an empirical question.

5.2 Permissible participation transfers

Fix a proposed net allocation g_i before transfers. A transfer t_i is positive when principal i receives money. With symmetric transfer bounds $|t_i| \leq B_i$, participation requires $g_i + t_i \geq o_i$. A balanced admissible transfer vector exists if and only if

$$\ell_i = \max\{o_i - g_i, -B_i\} \leq B_i \quad \forall i, \qquad \sum_i \ell_i \leq 0 \leq \sum_i B_i. \tag{12}$$

This follows from feasibility of intervals with a zero sum. Administrative costs must first be subtracted from g_i, and a date 0 cash payment must be checked against date 0 liquidity rather than only terminal certainty equivalents.

Figure 1 uses a transparent contract class. It includes baseline open participation and named-member, all-accept contracts. A named contract fixes its members, closes only if all accept, pays a fixed administrative cost plus a per-member cost, and permits balanced transfers of at most \$3,000 in either direction per member. Excluded principals cannot join that particular contract unilaterally.

The assumption is part of the contract class, not a permission inferred by the optimizer. Ordinary custom arrangements with the same capability are equivalent counterfactuals.

The resulting envelope is weak supportability. When all named members accept, the allocated transfers make exit unattractive. If the contract needs at least two acceptances, all-reject can still be an equilibrium. The curve therefore does not establish that the indicated contract will be signed, uniquely selected or immune to coalition deviation. It measures how much of the baseline participation loss this declared contracting technology could recover.

A broad externality charge is not automatically the cheapest way to do so. Marginal-contribution payments can implement a potential game while requiring a large subsidy. The earlier numerical rule that required $240,000 illustrates this distinction: that figure was the deficit of the particular payment rule, not the minimum subsidy needed for core stability. In the four-principal transferable-utility example, a budget-balanced allocation can meet every coalition's outside requirement because no proper coalition can form the program. Contracting research with heterogeneous externalities provides the relevant benchmark for stronger implementation questions (Bernstein and Winter, 2012; Segal, 2003).

5.3 Assurance changes selection

Suppose a subscriber bears holding cost h_i, receives a credible failure payment β_i if fewer than q subscribe, and obtains his formed-program payoff net of the holding cost when the threshold is met. Call a principal qualified if he strictly prefers participation to his outside option in every admissible formed membership containing him.

Proposition 2 (Assurance with excludable benefits)**.** *If at least q principals are qualified and $\beta_i > h_i$ for each of them, subscribing is strictly dominant for those principals in the deterministic participation game. Every equilibrium forms the program. Different equilibrium memberships may remain.*

Proof. If joining does not meet the threshold, the qualified principal receives $o_i - h_i + \beta_i > o_i$. If it meets the threshold, he strictly prefers the formed payoff by qualification. If the program already forms without him, exclusion leaves him his outside payoff, while qualification again makes entry strictly preferable. At least q such principals subscribe under every profile of others' choices, so every equilibrium forms. □

This is an application of dominant assurance contracts (Tabarrok, 1998). It addresses failure to coordinate among willing participants. It does not induce a member to accept a formed payoff below his outside option. Inside the distribution band, a payment that is made only on failure can leave the successful program unattractive to the pivotal member; success-contingent terms must change as well.

In the numerical example, three of four principals qualify. A $510 failure payment exceeds a $500 holding cost. The maximum failure payment is $1,020 because fewer than three subscribers means at most two payments. Equilibrium failure payments are zero. A reserve of $1,020 held for

90 days at an assumed annual opportunity cost of 4% costs \$10.06. The formed allocation yields \$39,500 before that reserve carrying cost and \$39,489.94 after it, if the provider's cost is counted. Zero equilibrium payout does not mean a costless mechanism.

The maximum promise changes if failure can also occur after the subscription threshold is met. A company refusal, a funding default or a settlement failure can expose the guarantor to more than $q - 1$ payments. A reserve adequate for the deterministic membership game is not evidence of adequacy for those additional contingencies.

5.4 Commitments and cash coverage

Let $K(\ell, a)$ be required cash at a funding date, and let $C(\ell, a)$ be internal funding that is both available and collectible under supported state ℓ and permitted withdrawal pattern a. Full cash reliability requires an external backstop with capacity at least

$$\overline{G} = \max_{\ell, a} [K(\ell, a) - C(\ell, a)]_{+}. \tag{13}$$

The realized shortfall, its expected payout, the price of the contingent obligation and the cost of carrying a reserve are distinct quantities. The bound is sufficient for cash coverage only if the backstop can pay on time and no other condition prevents execution. Correlation alone is neither necessary nor sufficient for a positive shortfall.

Three principals with available capacity of \$100,000 each face three ordinary states, each of probability 0.3, in which one principal cannot fund, and a common state of probability 0.1 in which none can fund. A program requires \$200,000. Two fixed \$100,000 tickets fund it with probability 0.3. An admissible rule assigning funding to the two available principals funds it with probability 0.9. The residual maximum backstop is \$200,000; expected payout is \$20,000. A state-price density of 2 in the common state and $8/9$ otherwise, normalized to mean one, values the contingent payout at \$40,000. These are synthetic amounts under stated collection rules, not market quotes.

A capable incumbent may already use contingent contributions, redundant commitments, a subscription facility or prefunding. Such capabilities belong in the comparison. Binding commitments may improve collectible funding but also surrender flexibility, require enforcement and remain exposed to default. Prefunding shifts cash into the arrangement earlier and consumes reserve capacity. Neither a contractual promise nor a recorded authority grant creates missing liquidity, consistent with the aggregate-liquidity distinction in Holmström and Tirole (1998).

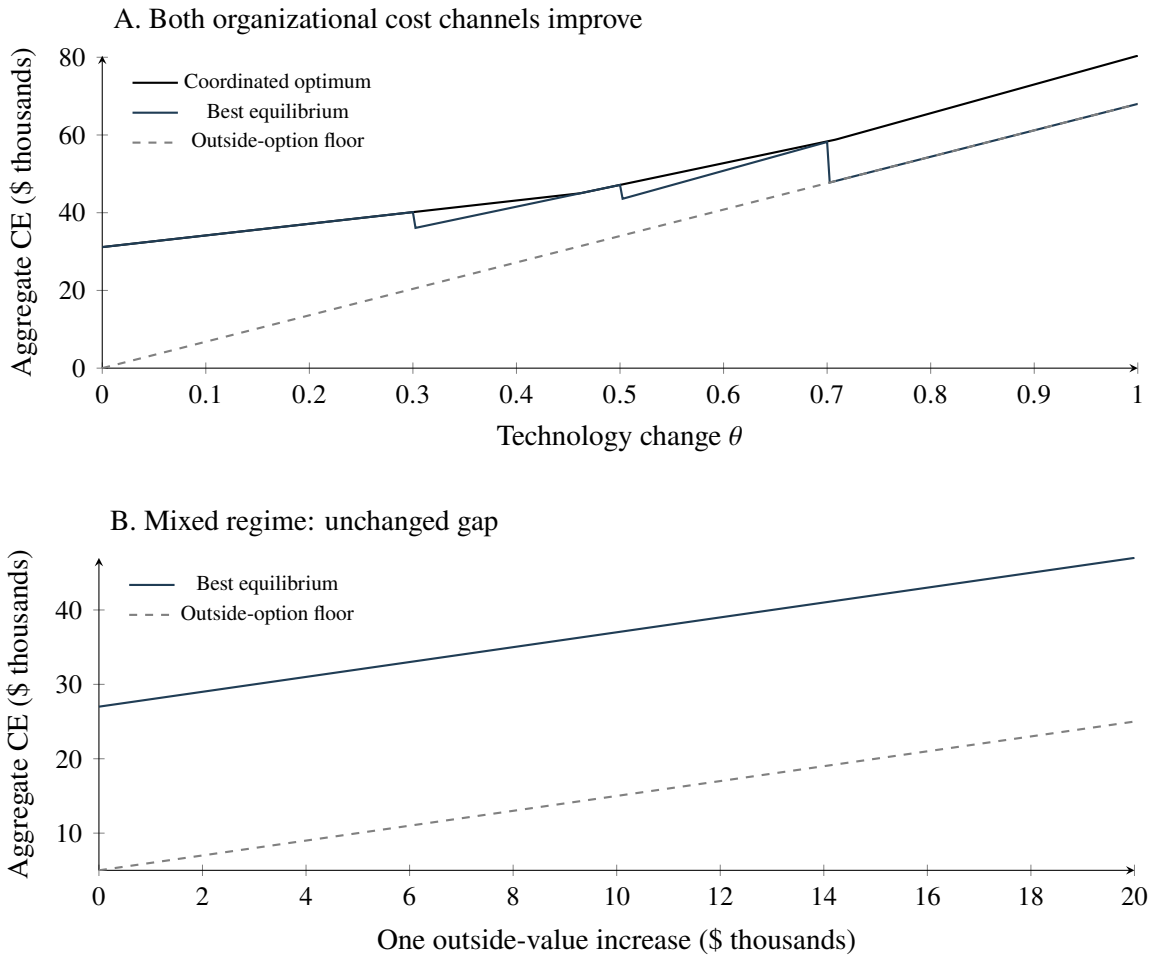


Figure 2: Two departures from the fixed-terms comparison. Panel A allows shared formation cost to decline as outside opportunities improve. Upward segments and changes in membership are retained. Panel B adds congestion beyond minimum scale; the equilibrium gap remains 22 over the plotted range. Neither panel is governed by the fixed-payoff compression theorem.

6 When the sign changes

6.1 Congestion and mixed regimes

Compression uses complementarity in net participation payoffs. It does not extend by relabeling every investment program a club. A first close may need more capital to be viable. An already oversubscribed allocation may instead become less attractive to each investor as more investors compete for it. A program can have both properties at different scales.

A three-principal example makes the reversal precise. The formation threshold is two, gross member benefits are $(15, 12, 20)$, and outside values are $(1, 1, 2)$, all in thousands of dollars. Total implementation cost is $3|S|^2$, divided equally, so each member pays $3|S|$. Initially all three form, with payoffs $(6, 3, 11)$ and aggregate welfare 20. The outside floor is 4 and the equilibrium gap 16. Raise the second principal's outside value from 1 to 7. He leaves; the remaining two pay less, and the best equilibrium payoffs become $(9, 7, 14)$, totaling 30. The floor rises to 10 and the gap rises to 20. The gap expands by 4, contradicting compression precisely where its net-payoff monotonicity assumption fails.

This example establishes possibility, not a universal positive comparative static under substitutability. A less congested membership could also remove a valuable payer, change prices or leave a program below its minimum. Figure 2 also reports a mixed cost schedule that first rewards scale and then penalizes congestion. In that example, the affected principal remains outside the best formed membership over the plotted range: the floor and ceiling rise equally and the equilibrium gap remains 22. The presence of both scale and congestion costs therefore does not determine the sign without the active membership margin. Classification therefore requires the program's active constraint and the contractual incidence of a change, not merely an asset-class label.

For empirical work, an observer should record whether an additional qualified participant lowers formation charges, dilutes an allocation, tightens a concentration limit or changes a supplier quote. The same manager may operate programs in all three regimes. A cross-sectional comparison of "private" and "public" investments would confound these differences rather than test the mechanism.

6.2 Continuous participation does not generally smooth the cliff

The binary game makes the compression theorem transparent. It is not the reason a formation loss can occur. Consider four principals, each with two units of capital. Principal i chooses nonnegative real quantities q_i in a common program and x_i in an individual investment, with $q_i + x_i \leq 2$. Unused capital remains liquid. The common program needs total quantity $Q \geq 7$, produces gross CE gain 15 per unit and costs 40 to form, charged in proportion to q_i. A reservation costs 0.5 per common unit whether or not the program forms. Individual gross gains per unit are $(18, 2, 3, 4)$, and a common per-unit individual implementation cost c is deducted. Values are thousands of dollars. A financial embedding assigns \$100,000 of investment capital to each quantity unit and gives each principal a separate \$20,000 expense reserve, sufficient for every

admissible route in the comparison. That reserve is available on the same terms across routes; paid expenses reduce net wealth, and unused cash remains liquid.

At $c = 10$, the all-common allocation gives each principal 19, aggregate 76. The first principal's outside payoff is 16; the others optimally hold liquid capital. All-common is a Nash equilibrium and the unique strong equilibrium. All-outside is also Nash because a unilateral participant cannot meet the threshold. At $c = 7$, the first principal's individual payoff becomes 22. Any positive common participation gives him less than acting alone. His exit leaves only six units available, below the program's seven-unit requirement. The unique Nash equilibrium gives payoffs $(22, 0, 0, 0)$, while the coordinated optimum remains 76.

Proposition 3 (A formation cliff with mixed continuous portfolios)**.** *The four-principal economy just specified has a fall in highest pure-equilibrium welfare from 76 to 22 as c falls from 10 to 7, although its coordinated maximum is 76 at both costs. This comparison permits all real quantities satisfying the individual capital budgets and permits simultaneous common and individual holdings.*

The proof is in Appendix A. The result does not establish the full compression theorem for arbitrary continuous participation games. A smooth control illustrates the distinction. Without a formation threshold or fixed cost, maximizing $(b-a)q - dq^2/2$ over $0 \leq q \leq \bar{q}$, with $d > 0$, gives $q^* = [(b-a)/d]_{[0,\bar{q}]}$, which changes continuously with the outside return a. The discontinuity in the counterexample is a property of the formation technology and participation incentives.

6.3 Overlapping programs

An investor can use one program as an alternative to another while retaining the ability to hold both. We study four principals with one unit of capital each. Program A requires three units and costs 3; it provides gross CE benefit 10 per unit to each investor. Program B requires one unit, has cost 12 or 2, and provides benefits $(14, 14, 1, 1)$ per unit. Fixed costs are allocated in proportion to actual invested quantities. Failed reservations return capital and earn no investment benefit. Every principal may invest in A, in B, in both or in neither, subject to his one-unit budget.

On half-unit and quarter-unit allocation grids, every joint profile and every unilateral deviation is enumerated. At B's higher cost, highest equilibrium welfare is 37 and lowest equilibrium welfare 16. The coordinated maximum is 37, attained by all-A. When B's cost falls to 2, the unique equilibrium is all-B, with payoffs $(13.5, 13.5, 0.5, 0.5)$, totaling 28. Coordinated welfare rises to 39. One coordinated allocation places each of the first two principals half in A and half in B, and the other two fully in A, yielding $(10.5, 10.5, 9, 9)$.

The preferred joint arrangement uses both programs in the first two portfolios. The lower cost of B attracts the principals who are needed to sustain A; after they leave, the other two cannot meet A's minimum. Refining the grid does not move the reported welfare extrema or the member payoffs at these allocations. The finer grid adds weak nonformation equilibria in the high-cost case, which remain in the reported equilibrium set.

These are fixed-price partial-equilibrium economies. The opportunity cost of one program is endogenous to other feasible programs, but security prices, entry by suppliers and economy-wide capital supply are not jointly determined. Extending the order argument to overlapping programs requires conditions on cross-program substitution and resources. Within-program cost sharing alone does not provide them. A general characterization remains open.

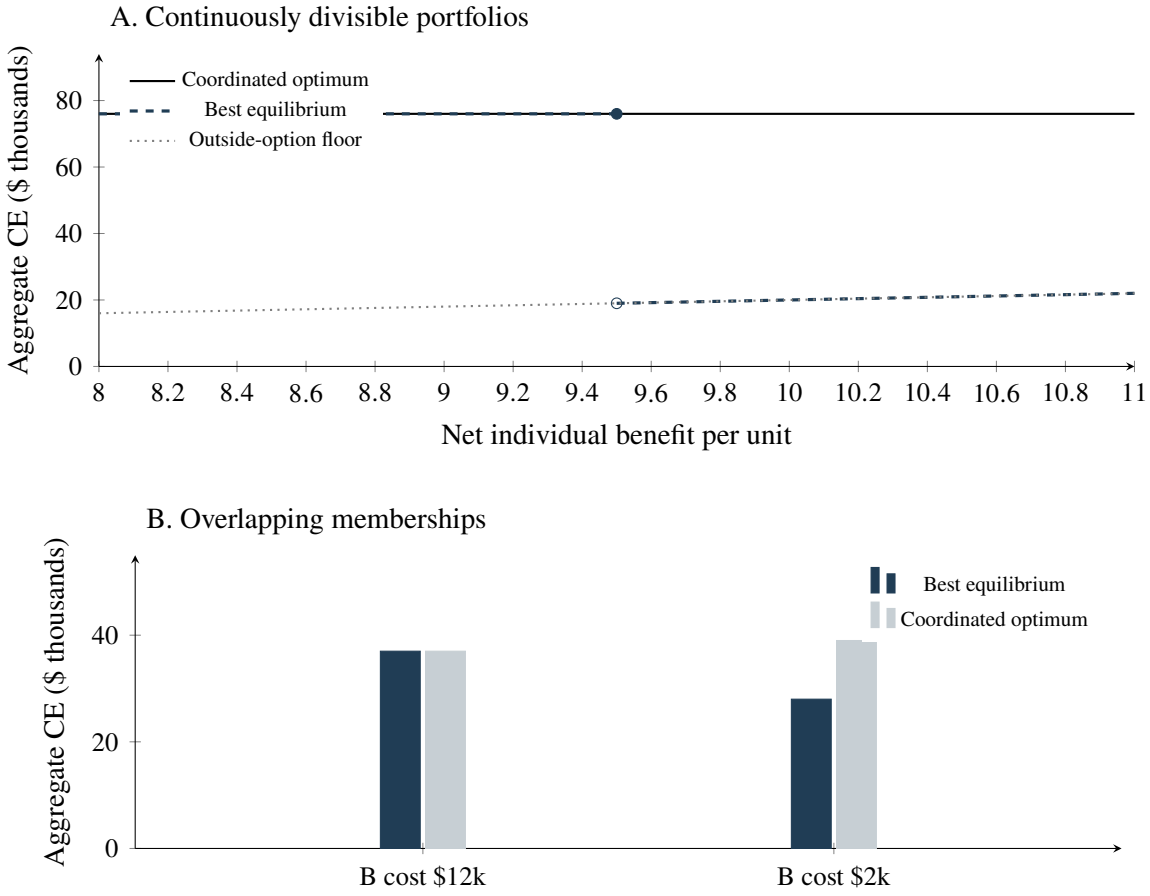


Figure 3: Divisibility does not remove the formation cliff. Four principals can divide continuous capital between the common investment and a separate opportunity. The common program needs seven of eight units. A separate expense reserve finances every permissible route. The coordinated maximum stays at 76 while the best equilibrium falls when the pivotal investor leaves. Panel B reports the overlapping-program comparison, independently enumerated on half-unit and quarter-unit grids.

7 Numerical evidence and economic accounting

7.1 Independent verification

The threshold engine represents small payoffs as exact rational numbers. It verifies the theorem over 78,732 comparisons in a structured three-principal parameter cube, including equal and unequal cost-sharing weights, ties, zero costs and negative net program values. Each comparison enumerates all memberships and directly tests every unilateral deviation. These calculations corroborate the proof; they are not its logical basis or an estimate of how frequently compression occurs in markets.

The supplied verification script is preserved separately. Its random trials, parameter draws and reported counts are not mixed with our structured enumeration. An independent financial reviewer implemented another evaluator and reproduced both 401-point financial curves and the overlapping-program grids. Mathematical verification and the literature audit have separate records.

The five-principal main figure has fixed gross values $(12, 13, 11.4, 12.75, 12)$, equal cost shares, formation cost 30 and threshold three. Outside values are $\theta(4, 8, 12, 18, 26)$ for $0 \leq \theta \leq 1$. All figures are thousands of dollars. At zero, all five participate and aggregate welfare is 31.15. When the fifth principal's outside value crosses his slack, membership contracts to four. A later departure makes every remaining formed membership unacceptable, even though the original program can remain jointly valuable. The floor continues to rise.

The administrative cost in the contract envelope is $0.4 + 0.04|S|$, paid by the named members on successful formation before balanced transfers. Failed negotiations carry no additional cost in this declared contract class. Each member can pay or receive at most 3. The figure compares this class with baseline participation and the coordinated maximum using the same gross values, outside options and program cost. It reports net investor CE after real administration costs. A supplier markup is not included in these main curves; if introduced, its receipts must be shown separately from resources consumed.

The second cost-incidence curve also lowers formation cost to $30(1 - \theta)$. It is not an application of the fixed-terms theorem. It illustrates condition (10) when technology improves both outside opportunities and shared implementation. Its upward sections and membership changes are retained. A fall in average organizational cost is compatible with gains over one interval and a participation loss at another.

7.2 What a numerical gain means

Three accounting distinctions are indispensable. First, gross investment benefits and information remain the same across equivalent implementation routes. A private instrument receives no extra return premium to make its program attractive. Second, a price paid from one included party to another redistributes surplus; formation and servicing expenditures consume it. Third, a member's

Table 1: Checked outcomes across distinct models

Model and change	Worst NE	Best NE	Coordinated
Continuous portfolios, cost 10	16	76	76
Continuous portfolios, cost 7	22	22	76
Overlapping programs, B cost 12	16	37	37
Overlapping programs, B cost 2	28	28	39
Congestion, before outside improvement	4	20	–
Congestion, after outside improvement	10	30	–

Values are synthetic incremental CE in thousands of dollars. Rows compare within their stated model; values are not comparable welfare levels across different economies. The continuous result is analytic. Overlapping-program values reproduce on half- and quarter-unit grids. Dashes indicate quantities not used in that comparison, not zero.

improved outside option can make him better off while reducing what other members can obtain. Aggregate welfare alone cannot identify the beneficiary.

For example, the low-cost overlapping-program equilibrium raises the first two principals' payoffs from 9.25 to 13.5 but lowers the other two from 9.25 to 0.5. The aggregate loss is not a unanimous loss. A coordinated allocation at the lower cost can achieve 39, but sustaining it requires addressing the first two principals' outside opportunities. The reported coordinated payoffs $(10.5, 10.5, 9, 9)$ are not by themselves an equilibrium. Holding the other allocations fixed, either of the first two principals can move his entire unit to B and obtain $14 - 2/1.5 = 38/3 \simeq 12.667$, above his mixed-allocation payoff of 10.5. The payoff of 13.5 belongs instead to the all-B equilibrium, in which the other principals also change their choices. The contract class and the outside alternatives must be recalculated jointly.

The four-principal continuous example likewise cannot preserve every pre-change member payoff while compensating the pivotal investor. Keeping the other three at 19 and giving the first 22 would require 79, while only 76 is available. A transfer can make the program preferable to the new outside outcome without making every investor as well off as under the old formed program. This distinction prevents a participation subsidy from being presented as a free Pareto improvement over history.

Numerical grids represent an approximation only where quantities are intrinsically divisible. A contractual minimum or lot size is an economic restriction and remains fixed as the numerical grid is refined. The small threshold checks use exact arithmetic. For grid-based comparisons we retain the acceptance threshold of \$25 per principal and \$100 aggregate CE; comparisons moving by more than those amounts are unresolved. The reported overlapping-program allocations and welfare extrema do not move between the two grids.

7.3 Four applications and discriminating observations

Public securities. A continuously available liquid asset has no private-subscription deadline. Public portfolios provide a useful null where minimum formation scale is slack and existing funds or coordinated separate accounts can reproduce the allocation. Taxes, restrictions, market impact and small-account costs can still matter. The test is whether the proposed institution changes a specific cost or attainable after-tax outcome beyond what a capable ordinary route provides.

Private funds. A first close can require minimum capital, while later admission can face capacity constraints or bargaining over fees. Dynamic commitment and cash-flow management already has a substantial literature (Luxenberg et al., 2022). The formation mechanism is relevant where departures by a few principals can make a planned program unavailable or uneconomic. It predicts attention to the identity and slack of those principals, not just aggregate indicated demand. An empirical history should preserve subscriptions that were considered, refused, delayed or withdrawn, as well as completed commitments.

Secondaries. A proposed transfer can require seller participation, company consent, a minimum block and an admissible holding structure. It can also be oversubscribed. Buyer and seller can gain through risk sharing without any new productive capital reaching the company. A discount to reported NAV neither measures that gain nor proves mispricing, because valuation dates, uncertainty and contractual rights differ. To test compression, one needs the actual terms, available substitutes and costs of an unsuccessful transaction; to test decongestion, one needs the allocation rule and the capacity lost or gained by remaining buyers.

Credit. Credit agreements contain specific minimum borrowings, increments, eligibility conditions, amendment rules and borrowing bases. The public-document slice includes a filed agreement with a $500,000 minimum borrowing definition, a stated exception for the entire remaining commitment, and $100,000 increments. Those terms illustrate why a continuous capacity number cannot be converted into an executable quantity without contractual conditions. They do not calibrate the participation threshold in our model.

Subscription-credit facilities are a particularly direct institutional record of how conditional capital is recognized. A filed 2022 agreement applies investor eligibility and concentration rules to the borrowing base, including aggregation of related investors (Goldman Sachs Middle Market Lending Corp. II, 2022). These terms distinguish nominal commitments from capital a lender is prepared to finance against. They also reflect bargaining, collateral, legal recourse, lender regulation and the facility's other terms. They are evidence of institutional treatment of conditional capital, not a pure market price of credibility.

Derivative markets provide another comparison between common and individualized execution. A listed hedge competes with bilateral arrangements, and its usefulness depends on the terms sustained by participation on both sides. If those net participation payoffs satisfy inclusion monotonicity, an improvement in the bilateral alternative has the compression effect. In actual markets, order imbalance, adverse selection, margin and counterparty capacity can change that sign. Johnston and McConnell (1989) attribute the decline of GNMA CDR futures to delivery-option design

and reduced hedging effectiveness relative to Treasury-bond futures. Their evidence motivates attention to contract design; it does not identify the welfare effect of cheaper bilateral customization. Residual hedging demand can guide the search for a new claim, but neither a large exposure nor its leading statistical direction supplies a willing risk bearer or a clearing price. Duffie and Jackson (1989) explicitly study futures innovation by exchanges in an incomplete-market equilibrium.

Exchange funds make the preservation of individual state within shared ownership concrete. Investors contribute concentrated positions to a partnership and receive interests in a diversified portfolio. Morgan Stanley's dated description distinguishes contribution rules, investor basis and liquidity restrictions; IRS guidance distinguishes capital accounts, partnership basis and a partner's outside basis (Morgan Stanley, 2023; Internal Revenue Service, 2025a,b). Such arrangements belong among the capable conventional alternatives. They preserve differentiated records while commonly delegating investment control, rather than demonstrating independent authority over every underlying trade. Related risk-sharing pools can create gains from compatible contributions, but a new concentrated exposure need not improve every incumbent's payoff. Entry monotonicity is a property of the specified allocation rule (Blier-Wong and Lauzier, 2026).

The four applications are classified by actual constraints. A private fund already above its efficient scale can be congested; a specialized public-securities vehicle can still have a binding fixed cost. Asset labels are not treatment assignments. An empirical protocol should classify each opportunity before observing whether it forms, document the outside alternatives available at the time, and follow the full event history through funding and ownership.

8 Technology and the incidence of interpretation costs

8.1 The technological question

Interpreting an investor's record can lower the cost of both separation and sharing. A system that reads a mandate cheaply may make an individual account economical. A system that maintains attributable obligations and permissions cheaply may also allow a shared program to preserve those same circumstances. The empirical object is the relative incidence in (11), not an undifferentiated claim that AI lowers transaction costs.

There are at least three channels. Interpretation converts evidence into an economic state and identifies unresolved conditions. Maintenance updates that state as notices and amendments arrive. Validation tests proposed actions against trusted permission, source provenance and financial constraints. A language model can assist with interpretation and maintenance; a deterministic rule can enforce a supplied grant. Neither supplies the missing legal right to act, creates a commitment or ensures that a counterparty's statement is true.

This distinction also separates our study from agent-market and financial-agent research. Shahidi et al. (2025) study the broader possibility that AI agents change search, contracting and market design. FinEvo-Bench studies retained experience in financial workflows (Deng et al., 2026). A financial-advisory factorial experiment by Brar et al. (2026) likewise makes tool

contributions an empirical question. Our experiment holds the economic history fixed and tests retained financial state and validation in two processing organizations. It asks whether the same principal's circumstances can be maintained more cheaply while preserving relevant distinctions, rather than whether the system becomes a better investment forecaster.

8.2 Matched economic worlds

The evaluation contains 64 held-out worlds. Each has four principals and four chronological checkpoints. Both processing modes receive the same underlying facts and perform the same required principal-level work. In separate mode, each principal is processed in an isolated context and costs are summed. In shared mode, one context can reuse common offering evidence while producing separate states and decisions. This is an operational comparison of batching and state maintenance. A conventional provider batching separate accounts can use the same capability.

The primary configurations are retrieval without persistent semantic records or deterministic validation, denoted R0V0, and a records-and-validation workflow, R1V1. A full 2-by-2 factorial is completed on 16 predetermined worlds. Four worlds receive predetermined repeats. The resulting synthetic evaluation contains 1,344 four-principal book-processing episodes: 1,024 primary, 256 factorial-completion and 64 repeat episodes. A separate 12-case public-document slice produces 48 single-dossier episodes. Development uses at most 256 episodes and remains outside the evaluation estimates.

An episode is not a model call. A separate-mode book requires four isolated processing contexts; a shared book requires one. A context may make a second inference after requesting historical documents. Processing streams are shuffled with a fixed seed, while checkpoints within each stream remain chronological. Every inference, retrieved source, prompt, output, timeout and returned usage count is recorded. System overhead, cache usage and task-prompt size are reported separately where observable. Savings caused by amortizing an invocation are not described as evidence that a new financial institution is necessary.

Synthetic ground truth comes from executable state transitions. Documents are rendered deterministically, and labels are never derived from the evaluated model. The archive supplies only evidence known at each checkpoint, with separate knowledge and effective dates. The histories include settlement of a previously announced distribution, partial payment of an accepted call, future-effective preference changes, expiring eligibility evidence, duplicate position attribution and contradictory current observations. A change in intention does not remove an accepted obligation.

R1 retains model-produced semantic states, source links, temporal conditions and validation warnings. It does not receive oracle-corrected records. R0 can retrieve the same historical documents but does not retain that semantic state. Both configurations must construct current state for the required outputs, so construction is charged in both; the comparison isolates persistence rather than assuming that records are free. Proposed actions do not change subsequent ground truth, which remains an identical exogenous history across conditions.

8.3 Validation and the information boundary

The source catalog identifies principal scope, field coverage, effective dates and knowledge dates. Every arm receives it. The authenticated permission registry is a separate channel. A source document that tells the system to ignore restrictions or asserts an unlimited grant cannot alter that registry. V1 checks current grant scope, status, expiry, remaining aggregate budget, expense-inclusive arithmetic, contractual lots and source-catalog consistency. A failed check blocks the proposed action and records a review condition.

Validation has a deliberately explicit boundary. It can catch a cited source for the wrong principal, a future or stale observation, a breached grant or an arithmetic inconsistency. It does not prove that a monetary value was correctly extracted from a superficially plausible source. A hallucinated cash amount can pass structural checks if its citation and internal arithmetic appear consistent. We therefore report errors that survive into proposed actions, as well as errors intercepted by validation.

The public slice contains three cases each from credit agreements, tender documents, indentures and closed-end or interval-fund filings. The 12 cases span 11 filings and were previously used in the earlier edition. They are an exploratory external slice, not newly held-out evidence. Labels are agent-audited, source-linked and limited to the quoted passages. An ambiguous question about equivalence of share-class rights was replaced before evaluation with a question about whether a conversion ratio is expressly stated. No human-validated label claim is made.

8.4 Complexity, quality and economic consequences

Economic complexity is measured by distinct principal-condition pairs whose bounded perturbation changes the financial or executable lot-capacity set at some checkpoint. Each counted condition has a saved before-and-after witness, source identifier and effective date. Desired quantity is excluded from this feasibility count; it affects the preferred action rather than necessarily changing what is permitted. Document length, duplicated facts, presentation format and memory burden are separate features. The complexity bands are assigned before evaluation from development-set cutoffs, with ties kept together. They contain 38 low-, three middle- and 23 high-complexity worlds. The middle group is descriptive; comparisons across tiers also reflect other features of these fixed histories.

The statistical unit is the economic world. Principal and checkpoint observations are repeated measurements. Paired world contrasts estimate the change in processing cost from R0V0 to R1V1 within each mode, and then the difference between those changes. World-level bootstrap intervals describe uncertainty conditional on this simulator and generator. The small number of shared presentation families limits inference over the population of financial documents. Factorial interactions on 16 worlds and public comparisons are exploratory.

Correctness is assessed both by component and by the complete decision. We distinguish unsupported proposed authority, missed obligations, valid capacity wrongly blocked and appropriate abstention. Invalid lots and subminimum proposals count as invalid actions. An infeasible

proposal is not silently clipped into a different, feasible trade. Nominal proposed capital, capital admissible as proposed and aggregate shortfall relative to the stated program target are separate quantities. Error dependence within a world is reported alongside average accuracy.

A common downstream decision engine supplies a further, explicitly model-implied measure. It solves the same finite-lot program using either interpreted or oracle records. Gross CE benefits follow a declared piecewise-linear schedule, and fixed and proportional participant costs consume both value and capacity. The no-program option costs nothing. The selected action is evaluated against the oracle economic state; if infeasible, it is reported as such rather than assigned an invented monetary penalty. For financially feasible actions satisfying every principal's true participation constraint, regret is the oracle optimum less the oracle value of the selected action. A negative true member benefit is reported as a participation failure with member values; regret is then unassigned. The gross CE schedule is $\eta \min(q_i, d_i)$, with $\eta \in \{0, .04, .08, .12\}$, preferred quantity d_i and expenses 5,000 + $.02q_i$ for positive participation. A formed program requires at least $650,000; quantities are $1,000 lots with a $25,000 participant minimum. The engine enumerates active sets and optimizes discrete concave quantities, including quantities beyond the preferred kink when needed for worthwhile formation. These utility scenarios are synthetic. This exercise isolates the consequences of interpreted information after a common decision rule; it is distinct from the model's raw recommended amount.

Human processing and correction prices remain assumptions because no expert-review time was observed. We therefore report cost–quality results, raw resource use and break-even prices. An accuracy improvement at higher cost can be valuable at some review prices and not others. Equal required work does not guarantee equal achieved quality, so raw token savings alone cannot be inserted into (11) as fixed-quality cost reductions.

The field-and-date source catalog and authenticated permission registry are supplied inputs. Their acquisition, authentication and continuing maintenance costs are not measured. Charged record construction refers to the model-produced semantic state, not the creation of these trusted upstream channels.

The public slice requires 72 underlying calls across 48 configurations: 24 records-preparation calls and 48 answers. Preparation is charged to the records conditions. Because these are single-dossier questions, their comparison does not identify longitudinal persistence effects.

Measured shared-mode costs belong to a four-principal book. They do not identify an observed cost charged to a pivotal investor. The bridge to individual participation therefore reports equal allocation and allocation proportional to the frozen consequential-condition count as explicit policies. A lower shared bill can coexist with an unfavorable incidence for a particular member. These policies produce a sensitivity for the participation slack required to absorb the measured cost change; the experiment does not identify an actual pivotal investor or an observed change in program formation.

Table 2: Primary paired processing results

	Separate		Shared	
	R0V0	R1V1	R0V0	R1V1
Correct final decisions	946	1,007	974	1,016
Unsupported raw actions	0	0	0	0
Unsupported final actions	0	0	0	0
Final review flags	98	98	95	98
Model calls	1,791	1,429	448	353
Uncached input, millions	5.887	5.024	2.243	2.004
Cached input, millions	22.328	17.817	5.527	4.400
Output tokens, millions	0.619	0.674	0.467	0.507
Summed call hours	10.16	9.53	5.02	5.26
Calls with missing usage	3	4	0	1

Each column contains 64 worlds, 256 four-principal checkpoints and 1,024 principal decisions. R denotes persistent records and V deterministic validation. Calls include historical retrieval. Time sums individual call durations; it is not elapsed evaluation time. Missing usage remains unpriced. Review flags include justified abstention, validation blocks and missing outputs.

8.5 Processing quality and resource use

The completed evaluation contains 5,293 synthetic inference calls and 72 public-document calls. Eight synthetic calls failed, each in a different processing context: seven service-capacity errors and one 180-second timeout. All remain in the scheduled denominators. Independent implementations reproduce the scored outcomes and cost contrasts, and a separate receipt audit reconciles the frozen schedule, source hashes, chronology and unique call records.

The complete workflow raises final decision accuracy from 946 to 1,007 of 1,024 principal-checkpoint decisions in separate mode, or from 92.38% to 98.34%. In shared mode the corresponding change is from 974 to 1,016, or from 95.12% to 99.22%. These are complete-output scores: capacity, recommendation, authorization, proposed amount and status must all agree with the oracle. No unsupported raw or final action is observed in any primary arm. The guard's gain is a specific representational correction: it replaces a null proposed amount with zero while retaining review status. It makes this change for 53 separate and 36 shared decisions, producing 52 and 36 additional fully exact outputs. It blocks no positive proposal and changes no status. Every successful primary context has all scored state fields correct. The accuracy gain therefore concerns abstention encoding, not demonstrated improvement in semantic extraction or prevention of unsafe investment. A post-hoc diagnostic treating a null proposal under review as the same abstention as a zero proposal yields 1,010 versus 1,007 exact outputs separately and 1,020 versus 1,016 jointly. The apparent improvement disappears under that equivalence; the frozen primary scores in the table are retained unchanged.

Persistent state reduces historical retrieval and the number of invocations. Primary separate-mode calls fall from 1,791 to 1,429; shared-mode calls fall from 448 to 353. Both complete

Table 3: Cost incidence by economic complexity

Tier	Worlds	Machine dollars [95% CI]	Review flags [95% CI]
All	64	0.028 [0.017, 0.039]	0.05 [-0.06, 0.19]
Low	38	0.029 [0.015, 0.045]	0.05 [-0.11, 0.26]
Middle	3	0.047 [0.023, 0.084]	0.00 [0.00, 0.00]
High	23	0.024 [0.005, 0.045]	0.04 [0.00, 0.13]

Positive incidence means larger R0V0-to-R1V1 savings in separate processing. Values are per four-checkpoint book. Machine prices are assumed at \$2, \$0.20 and \$8 per million uncached input, cached input and output tokens. For the point estimate at review price p, add p times review-flag incidence; intervals for the priced sum use paired world resampling. These sensitivities impose no equal-achieved-quality restriction. The three-world middle tier is descriptive.

workflows use fewer input tokens but more output tokens. Summed call time falls in separate mode and rises in shared mode, so fewer calls do not imply lower measured processing time. Deterministic guard computation totals less than 0.3 seconds in either primary validation arm, separately from the call time in the table.

The matched factorial provides a limited decomposition. Across its 16 worlds, the average effect of records on call counts is -5.28 per four-checkpoint book in separate mode and -1.56 in shared mode. The corresponding validation effect on complete-decision accuracy is 5.86 percentage points in each mode. Records have a 1.95-point accuracy effect separately and a -0.39-point effect jointly. The accuracy interaction intervals include zero in both modes. Thus the factorial supports a retrieval-saving role for persistent state. Its validation score must be read in light of the abstention-encoding correction, rather than as a semantic-interpretation gain. It does not establish a positive complementarity between the two components.

The public slice is at a ceiling: all 192 scored answers are correct across the four configurations. It cannot discriminate their accuracy on these questions. Predetermined synthetic repeats reproduce 248 of 256 principal decisions exactly. The retained differences and failures matter to the economic use of the outputs even when mean accuracy is high.

8.6 Cost incidence and the participation margin

At the preselected middle price schedule, recorded inference costs per four-checkpoint world fall from \$0.331 to \$0.297 separately and from \$0.146 to \$0.140 jointly. Shared processing has a lower cost level in both configurations, while the tested change saves more in separate processing. The separate-minus-shared saving is \$0.0281 per world, with a paired 95% interval of [\$0.0173, \$0.0393]. The direction survives deleting any one of the three presentation families. These are prices applied to observed token use, not invoices or the full cost of serving an investor.

The complexity comparison does not support the proposed concentration of separate-mode savings among the most complex mandates. Mean machine-cost incidence is \$0.0290 in the low tier and \$0.0242 in the high tier. The intervals for the high-minus-low resource contrasts include

zero. The three-world middle tier supplies no reliable ordering. The interpretation-cost channel in the model is measurable, but this experiment does not show that mandate complexity makes adverse formation effects typical.

Review costs can dominate the token-price comparison. At the same machine prices, assigning \$5 to each final review flag gives mean incidence of \$0.263 with an interval of [−\$0.277, \$0.961]; at \$25, the mean is \$1.200 with an interval of [−\$1.527, \$4.711]. The point estimates do not cross zero at a nonnegative review price in any tier under the three reported token-price schedules. Their uncertainty does not establish the sign once review is priced. Review flags include appropriate abstention and failed outputs as well as correctable errors; they are not measured labor minutes. Usage is missing for the eight failed calls. Recorded machine costs leave those calls unpriced, and a difference between cost differences need not be a lower bound.

The allocation of shared cost changes the member-level comparison. For an internal resource saving to alter participation, it must reach the principal through a charge, an outside cost or another net benefit; provider savings need not pass through to client prices. At the middle machine prices and zero review price, separate savings exceed allocated shared savings for 164 of 256 member histories under equal allocation and 166 under condition-count allocation. The average member incidence is \$0.0070 under both budget-balanced policies, while individual signs differ. The mean positive part of incidence across all 256 member histories is about one cent per four-checkpoint history. For each history, this positive part is the additional participation slack needed to absorb its adverse cost incidence in the restricted accounting; the reported mean is not conditional on being adversely affected. This is not an estimate of economically material exit: initial real-world participation slack and the other costs of investment organization are unobserved. Nor can four synthetic checkpoints be converted into an annual saving without an additional frequency model.

8.7 Opportunity capture and failure exposure

Complete-decision accuracy and investment opportunity capture rank the configurations differently. Under the common downstream optimizer with $\eta = .12$, mean model-implied regret per checkpoint is \$106.25 for separate R0V0, \$365.23 for separate R1V1, zero for shared R0V0 and \$285.94 for shared R1V1. The oracle forms in 182 of 256 checkpoint states; the corresponding interpreted-state decisions form in 181, 181, 182 and 181. Every selected action in these comparisons is financially feasible and individually acceptable under the oracle. The reported losses are forgone modeled opportunities, not negative member returns or realized investor losses.

The same ordering appears at $\eta = .08$: mean regret is \$50.31, \$187.89, zero and \$148.13, respectively. At $\eta = 0$ every configuration chooses no program and regret is zero. The complete results retain all four utility schedules and each member's payoff. Failure cases remain part of those comparisons; selecting only successful processing would answer a different question.

No primary arm falsely reports attainment of the \$650,000 funding target under the specified scoring rule. Because unsupported final actions are absent, the joint-error histogram is zero throughout, and the excess pairwise-error endpoint is also zero. This study consequently supplies

no positive evidence for correlated unsafe interpretation errors. It cannot establish that the risk is absent outside these histories. All positive primary downstream regret occurs in books with a service failure. The single shared-mode timeout withholds a four-principal response, accounting for the shared complete workflow's entire $882,000 shortfall in valid proposed participation across the sampled checkpoints and its positive downstream regret. A missing common response can therefore withhold several principals' decisions at once. That is operational dependence, distinct from several successful responses making the same semantic mistake. The eight observed failures do not identify a causal effect of records or validation on the service's failure rate.

The measured result is therefore narrower than a claim about the optimal organization of capital. Shared processing amortizes work, records reduce repeated retrieval, and validation standardizes abstention. Semantic extraction is at a ceiling on successful primary contexts. Recorded marginal savings favor separate processing, the complexity prediction is unsupported, and the observed service failures account for every positive primary downstream loss. The experiment identifies quantities that belong in the participation calculation; it does not show that their present magnitudes destabilize an actual investment program.

9 Institutions and the empirical test

9.1 What a coordinator must supply

The model identifies a narrow economic service. A coordinator can preserve a valuable investment by finding mutually acceptable terms, providing credible assurance or making compatible implementation cheaper. Each service has a different source of value. A better forecast changes investment selection; a lower administrative expense changes the resource cost; a timely authority check changes what can be executed; a supplier concession changes the division of surplus. They should not be combined into a single unexplained optimization gain.

Independent governance can be valuable when an investor's circumstances matter to his acceptable participation and can be preserved without duplicating the underlying investment. It can obstruct formation when retained exit options remove a pivotal member, or when constraints prevent otherwise useful transfers or contingent funding. The institution earns its cost only if its permissible interventions improve net outcomes relative to a capable ordinary route.

Ownership of the coordinator is not identified by the cost of a guarantee reserve alone. A low-cost reserve provider could finance an independently owned coordinator. Control, conflicts, contracting costs and the beneficiaries' ability to monitor also matter, as in the broader analysis of ownership by Hansmann (1996). Likewise, a registry that verifies a price cap can give it bargaining relevance only when the cap is observable and credibly fixed. Strategic delegation supplies that logic (Katz, 1991); the presence of a database does not establish the required commitment technology.

A coordinator affiliated with a supplier faces an additional problem: it may prefer allocations that increase its own fees or inventory turnover. The model's baseline assumes that proposed al-

locations are evaluated under a declared objective, with all costs and participant outcomes visible. A production design would need conflict rules, allocation priorities and a record of which information each decision used. These are substantive restrictions on implementable arrangements, not implementation details that can be added after portfolio optimization.

9.2 Observations that would change the conclusion

The formation mechanism is economically important only if binding thresholds and pivotal departures matter in actual opportunities. A prospective study should retain all considered programs, including those that fail, and classify their formation and capacity regimes before outcomes are observed. It should record the investor's available alternative, quoted terms, mandate constraints, actual authority and funding status at each decision date. Later documents must not be used to reconstruct an unrealistically informed earlier decision.

The first test compares sensitivity to changes in individual implementation cost across otherwise similar programs with binding and slack formation requirements. A lack of differential sensitivity would weaken the claim that compression is a first-order organizational effect. It would not contradict the theorem's conditional logic. The second test asks whether departures are concentrated among principals with small measured participation slack and whether technological savings disproportionately improve their outside alternatives. Complexity alone is not the predicted sufficient statistic.

The third test measures the response of permissible contracts. If ordinary custom vehicles, negotiated transfers or existing commitment structures routinely recover the available surplus at lower cost, the proposed new institution has little incremental portfolio value. If supplier bargaining captures the saving, member gains can disappear even while implementation becomes cheaper. If consent, documentation or settlement rather than investor interpretation is the binding delay, faster interpretation does not establish execution value.

The fourth test traces correlated interpretation errors to actual funding and action failures. Shared clause errors can cause several principals' apparent capacity to be overstated at once. The synthetic study includes a diagnostic for that dependence but observes no unsupported primary actions. It does not establish market-wide crowding or a causal common shock in investment returns. A real event history must include rejected actions and near misses, rather than only successful closings.

9.3 Limits of the established argument

The main theorem fixes program terms, permits open entry and studies pure-strategy equilibrium. It requires participation payoffs to weakly improve under membership expansion. It does not cover arbitrary strategic pricing, uncertain private outside values, endogenous information acquisition or general overlapping-program games. The continuous example proves that divisibility alone is

insufficient to remove a cliff; it does not prove that every continuous model has the same envelope property.

The welfare measure is a sum of certainty-equivalent dollars under a declared financial representation. A social evaluation would also specify the treatment of supplier rents, taxes, external effects and heterogeneous welfare weights. A voluntary transfer that restores participation need not be budget-balanced after administrative costs, coalition-stable or truthful. Authority and legal restrictions can remove it from the permissible class even when its aggregate arithmetic is favorable.

The evaluation is a bounded test of a particular interpretation workflow. Its model identifier, simulator, renderer, source catalog and processing organization limit extrapolation. Institutional adoption, customer willingness to pay, actual delivery costs and legal enforceability remain unmeasured. The strongest supported technological conclusion concerns the tested cost and quality changes; the strongest supported financial conclusion concerns the conditions under which such changes affect participation.

10 Conclusion

A shared investment can remain financially worthwhile after it ceases to be an equilibrium among independent investors. Under threshold participation and scale economies, a better outside alternative raises the equilibrium floor and compresses the range above it. The highest equilibrium can fall by the entire previous range net of the outside improvement. Continuous mixed portfolios need not remove the loss, and overlapping programs can transmit it from one opportunity to another.

The institutional response depends on the source of failure. Assurance can select formation among qualified participants. Transfers can restore participation when an allocation remains jointly valuable. Binding commitments and funded backstops change the resources and obligations available at execution. Their costs, authority requirements and remaining equilibria determine how much of the feasible gain is implementable. Congestion can reverse the membership effect, so programs must be classified by their actual constraints.

Technology enters on both sides of the participation margin. The paired study finds less repeated retrieval, standardized abstention and larger marginal recorded token-cost savings separately. It also finds a semantic-accuracy ceiling and losses of valid participation after service failures. The cost differences are small under the stated prices, and no observed unsafe action supports the proposed correlated-error channel. Establishing economic significance requires real participation slack, delivery costs and complete opportunity histories. The relevant question is which saving reaches the pivotal principal, at what quality and under which contract. Technology can improve what investors do separately while destabilizing what they achieve together.

A Proofs and computational characterizations

A.1 The monotone deletion algorithm

Start with $S_0 = N$. If $|S_t| < q$, no formed willing set remains. Otherwise remove every principal with $u_i(S_t) < o_i$ and call the remaining set S_{t+1}. A removed principal cannot belong to any willing formed subset of S_t, because his payoff in a subset is no larger. By induction every willing formed set is contained in each S_t. The procedure ends in finitely many deletions. If the remaining set meets the threshold, it is willing and contains every other willing set, so it is S^*. If not, no formed equilibrium exists.

This algorithm computes the largest formed equilibrium, not the entire equilibrium correspondence. Our verification code separately enumerates every membership and checks all unilateral deviations. The two procedures answer different questions when weaker equilibria coexist. Weak inequalities preserve indifference; changing to strict participation or a different tie rule changes the correspondence at boundaries and must be declared.

Repeated one-coordinate improvements extend Theorem 1 to any componentwise increase $o' \geq o$: the floor increases by $\sum_i (o'_i - o_i)$ and the gap weakly contracts. The corresponding highest-welfare bounds use that total increase. This follows by successive application, not by assuming that simultaneous departures can be analyzed independently of changing cost shares.

A.2 Proof of the continuous mixed-portfolio counterexample

Write $h = 15 - 0.5 = 14.5$ and $v = h - 40/8 = 9.5$. When the program forms, common benefit per unit is $h - 40/Q \leq v$ because $Q \leq 8$. At $c = 10$, all individual per-unit gains are less than v. Hence any principal's benefit is at most $2v = 19$, with equality in a positive common allocation requiring $Q = 8$ and his own $q_i = 2$. The all-common profile therefore Pareto-dominates every other profile and is a Nash equilibrium.

For completeness, it is the unique strong equilibrium. In any other profile at least one principal receives less than 19, and none receives more. If a principal is already at 19, equality requires every principal to supply two common units, contradicting a different profile. Thus all four strictly gain by jointly moving to all-common from any different profile. No coalition can strictly improve on all-common because no member can exceed 19. This uses the stringent strong-equilibrium concept; the theorem in the main text relies only on Nash outcomes.

At $c = 7$, the first principal's individual per-unit gain is 11, strictly above v and below h. In any formed profile with positive q_1, his benefit is bounded above by

$$11(2 - q_1) + vq_1 = 22 - q_1(11 - v) < 22. \tag{14}$$

A failed common reservation also lowers his benefit. Thus choosing $q_1 = 0, x_1 = 2$ strictly dominates any positive common reservation. The other three have at most six units in total and

cannot form a program that needs seven. Their individual net gains are negative, so they hold liquid capital and the low-cost Nash outcome is unique.

For any formed allocation, other individual investments have negative net benefit. Even assigning all residual capital to the first investor gives the upper bound

$$14.5Q - 40 + 11(8 - Q) \leq 14.5 \cdot 8 - 40 = 76. \tag{15}$$

All-common attains 76. No-formation welfare is at most 22. The coordinated optimum is therefore 76 at the lower cost as well. Every argument admitted real quantities and mixed holdings; no integrality gap is invoked.

A.3 Bounded transfers and coalition requirements

For each member let the admissible transfer interval be $[a_i, b_i]$, and impose $g_i + t_i \geq o_i$. The resulting interval is $[\max(a_i, o_i - g_i), b_i]$. A zero-sum selection exists exactly when each interval is nonempty and the sum of the lower bounds does not exceed zero while the sum of upper bounds is at least zero. Necessity follows by summing bounds. Sufficiency follows by starting at all lower bounds and increasing coordinates until the sum reaches zero. This is the construction used in the contract envelope.

Individual participation is only a subset of cooperative restrictions. If a coalition T can secure value $V(T)$ using its own resources and permitted claims, core stability additionally requires $\sum_{i \in T} y_i \geq V(T)$ for every coalition and $\sum_i y_i = V(N)$. A sum of CE is a transferable-utility value only when the required transfers are admissible. No general core–integrality equivalence follows from the presence of fixed costs. Such a claim requires a specified game and linear relaxation, as in the special facility-location setting of Goemans and Skutella (2004). The present compression proof does not use that equivalence.

In the four-principal continuous example at low individual cost, no proper coalition can form the common program because its members have at most six units. A coalition containing principal 1 can secure 22 by the individual route; another proper coalition can secure zero. The allocation $(22, 18, 18, 18)$ sums to 76 and satisfies every such lower bound. Thus zero outside subsidy suffices for the transferable-utility core in this specified game. It does not imply that the original continuous participation game implements those transfers or that they satisfy an unmodeled contractual restriction.

A.4 Assurance costs and funding bounds

In the deterministic count-threshold assurance game, a failed program has at most $q-1$ subscribers. If bonuses are nonnegative, the largest aggregate failure promise is the sum of the $q - 1$ largest

offered bonuses. Expected payout at a formed equilibrium is zero. The provider's reserve carrying cost remains a real cost, and in the numerical example it is

$$1{,}020 \times 0.04 \times \frac{90}{365} = 10.06027397 \text{ dollars.} \tag{16}$$

A success fee can recover that cost only if members remain willing after paying it. The example has ample slack to allocate the $10.06 across the three qualified members. If the provider has other failure obligations or can default jointly with subscribers, this calculation must be replaced by the relevant state-contingent funding problem.

For cash reliability, every state requires external cash of at least $[K - C]_+$. Taking the maximum over supported states and permitted withdrawals gives the necessary backstop capacity. Conversely, a provider able to deliver that amount when needed covers every modeled shortfall, so the bound is sufficient for the cash component under the stated assumptions. With probability p_ℓ, a state-price density m_ℓ and realized shortfall G_ℓ, expected payout is $\sum_\ell p_\ell G_\ell$, while the specified contingent value is $\sum_\ell p_\ell m_\ell G_\ell$. Neither is generally the maximum shortfall or the reserve opportunity cost.

A.5 Why fixed-cost figures do not certify a technology forecast

Theorem 1 holds gross formed payoffs fixed. Figure 1 increases outside values at fixed shared costs. Figure 2A changes both outside values and shared cost, so its membership changes are evaluated directly rather than attributed to the theorem. For a fixed membership, subtracting the old and new individual participation inequalities yields (10). Once entry, reallocation or supplier response is allowed, the candidate set and its charges must be recalculated. A small local shadow price cannot certify a discrete change in formation.

A technological claim therefore needs at least three observations: the effect on the cost of a specified outside route, the effect on that principal's cost within a specified shared arrangement, and his initial participation slack under contemporaneously available alternatives. The paired interpretation experiment informs part of the first two. It does not measure legal organization, investment diligence, bargaining, enforcement or all other costs in ϕ_i.

References

Bernstein, Shai, and Eyal Winter. 2012. Contracting with Heterogeneous Externalities. *American Economic Journal: Microeconomics 4(2), 50–76*.

Birmpas, Georgios, Evangelos Markakis, and Guido Schäfer. 2019. Cost Sharing over Combinatorial Domains: Complement-Free Cost Functions and Beyond. *arXiv:1910.06384*.

Blier-Wong, Christopher, and Jean-Gabriel Lauzier. 2026. Designing entry-monotone risk-sharing pools. *Working paper, arXiv:2606.00972v1*.

Boyd, John H., and Edward C. Prescott. 1986. Financial Intermediary-Coalitions. *Journal of Economic Theory 38(2), 211–232*.

Brar, Aryan, Justin Du, Avery Lor, Kylie Seto, and Eric Taylor. 2026. Retrieval-augmented generation vs. deterministic tax computation in multi-agent financial advisory: A 2x2 factorial experiment. *arXiv:2608.23908*.

Buchanan, James M. 1965. An Economic Theory of Clubs. *Economica 32(125), 1–14*.

Deng, Bo, Kang Zhou, Lifan Guo, Chongyang Tao, Xuanren Chen, Chenggang Xie, Renzhao Liang, Feng Chen, and Chi Zhang. 2026. FinEvo-Bench: A Longitudinal Benchmark for Self-Evolving Agents in Professional Financial Workflows. *Working paper, arXiv:2608.06144v2*.

Devanathan, Nikhil, Logan Bell, Dylan Rueter, and Stephen Boyd. 2026. A Distributed Method for Cooperative Transaction Cost Mitigation. *arXiv:2603.07881; working paper*.

Duffie, Darrell, and Matthew O. Jackson. 1989. Optimal Innovation of Futures Contracts. *Review of Financial Studies 2(3), 275–296*.

Duffie, Darrell, and Rohit Rahi. 1995. Financial Market Innovation and Security Design: An Introduction. *Journal of Economic Theory 65(1), 1–42*.

Goemans, Michel X., and Martin Skutella. 2004. Cooperative Facility Location Games. *Journal of Algorithms 50(2), 194–214*.

Goldman Sachs Middle Market Lending Corp. II. 2022. First Amendment to Revolving Credit Agreement. *SEC-filed Exhibit 10.1, July 26, 2022; incorporated marked agreement*.

Hansmann, Henry. 1996. The Ownership of Enterprise. *Harvard University Press*.

Holmström, Bengt, and Jean Tirole. 1998. Private and Public Supply of Liquidity. *Journal of Political Economy 106(1), 1–40*.

Iancu, Dan A., and Nikolaos Trichakis. 2014. Fairness and Efficiency in Multiportfolio Optimization. *Operations Research 62(6), 1285–1301*.

Internal Revenue Service. 2025a. Publication 541: Partnerships. *December 2025 revision*.

Internal Revenue Service. 2025b. Instructions for Form 1065. *2025 instructions*.

Johnston, Elizabeth Tashjian, and John J. McConnell. 1989. Requiem for a Market: An Analysis of the Rise and Fall of a Financial Futures Contract. *Review of Financial Studies 2(1), 1–23*.

Juarez, Ruben. 2013. Group strategyproof cost sharing: The role of indifferences. *Games and Economic Behavior 82, 218–239*.

Katz, Michael L. 1991. Game-Playing Agents: Unobservable Contracts as Precommitments. *RAND Journal of Economics 22(3), 307–328*.

Luxenberg, Eric, Stephen Boyd, Mykel Kochenderfer, Misha van Beek, Wen Cao, Steven Diamond, Alex Ulitsky, Kunal Menda, and Vidy Vairavamurthy. 2022. Strategic Asset Allocation with Illiquid Alternatives. *arXiv:2207.07767; working paper*.

Mankiw, N. Gregory, and Michael D. Whinston. 1986. Free Entry and Social Inefficiency. *RAND Journal of Economics 17(1), 48–58*.

Markowitz, Harry. 1952. Portfolio Selection. *Journal of Finance 7(1), 77–91*.

Merton, Robert C. 1971. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory 3(4), 373–413*.

Morgan Stanley Wealth Management. 2023. Exchange Funds: An Important Alternative for Your Asset Allocation. *August 2023, CRC 5903090, especially pages 3–6 and 8*.

Ostrom, Elinor. 2009. Beyond Markets and States: Polycentric Governance of Complex Economic Systems. *Nobel Prize Lecture, December 8*.

Ross, Stephen A. 1978. Mutual Fund Separation in Financial Theory—The Separating Distributions. *Journal of Economic Theory 17(2), 254–286*.

Sákovics, József, and Jakub Steiner. 2012. Who Matters in Coordination Problems? *American Economic Review 102(7), 3439–3461*.

Segal, Ilya. 2003. Coordination and Discrimination in Contracting with Externalities. *Journal of Economic Theory 113(2), 147–181*.

Shahidi, Peyman, Gili Rusak, Benjamin S. Manning, Andrey Fradkin, and John J. Horton. 2025. The Coasean Singularity? Demand, Supply, and Market Design with AI Agents. *NBER Working Paper 34468*.

Spence, Michael. 1976. Product Selection, Fixed Costs, and Monopolistic Competition. *Review of Economic Studies 43(2), 217–235*.

Sprumont, Yves. 1990. Population Monotonic Allocation Schemes for Cooperative Games with Transferable Utility. *Games and Economic Behavior 2, 378–394*.

Tabarrok, Alexander. 1998. The Private Provision of Public Goods via Dominant Assurance Contracts. *Public Choice 96, 345–362*.

Internet Appendix

I. Frozen evidence and computational scope

The replication package separates evidence from executable policy. The synthetic simulator supplies economic state, dated source documents and authenticated permission records. The interpretation system sees the documents, catalog and permitted registry fields; the scorer alone sees the complete synthetic state. The public slice preserves its source URLs, excerpts, question definitions and exploratory labels. Source text is not an authority channel.

The primary evaluation has 64 economic worlds, each with four principals and four checkpoints. A complete book episode contains all four principals even when they are processed separately. The primary comparison contains $64 \times 4 \times 2 \times 2 = 1{,}024$ book episodes. Completing the factorial on 16 predetermined worlds adds $16 \times 4 \times 2 \times 2 = 256$ episodes; repeating four predetermined worlds adds 64. The total is 1,344 synthetic episodes, requiring 3,360 processing contexts because separate mode uses four contexts and shared mode one. A context can invoke the model once more after requesting dated historical evidence. Twelve public cases in four configurations add 48 episodes and 72 underlying calls, including preparatory extraction. Development used 128 episodes in three separately recorded versions.

The model product identifier, reasoning setting, CLI version, timestamps, prompt hashes, schema hashes, request and response text, source retrieval, tool-use status and reported token usage accompany each call. Calls exceeding the prescribed time limit, using prohibited tools or failing principal-coverage requirements count as failures. Replaying a completed local receipt is an idempotent resume, not a new inference attempt. A future call to the same product identifier can return a different answer because the available interface does not expose an immutable weight snapshot.

The numerical package contains no market-price estimator. All exact-game payoffs and formation costs are stated synthetic primitives. The allocation experiments enumerate every profile and unilateral deviation on each stated finite grid. Analytic proofs support the continuous example. Contractual lot size is held fixed when numerical resolution changes.

II. State transitions and capacity calculations

A principal's financial capacity is the greatest permitted lot quantity after cash calls, protected liquidity, competing reservations, expenses and concentration are accounted for. In the synthetic task, the displayed formula is

$$C_i = 1{,}000 \left\lfloor \frac{1}{1{,}000} \max \left\{ 0, \min \left[\frac{B_i - O_i - L_i - R_i - 5{,}000}{1.02},\ M_i - X_i \right] \right\} \right\rfloor .$$

Here B_i is cash, O_i outstanding calls, L_i protected liquidity, R_i competing reservations, M_i the permitted exposure limit and X_i existing exposure. Unresolved contradictory cash evidence makes

the displayed financial capacity unknown. The authorized projection additionally accounts for grant budget, expiry, scope, eligibility and current instrument terms. Financially feasible capital can therefore exceed currently executable capital.

At the first update, the simulator settles an announced distribution and pays part of a previously accepted call. A 100,000 distribution less a 20,000 payment raises cash by 80,000. The remaining call stays on the book. Later updates add a competing reservation, permit an eligibility certificate to expire, change terms, make a preference amendment effective, revoke one grant and allow another to expire. Conflicting current observations require review. These transitions are exogenous and common across experimental arms; no model proposal rewrites the subsequent economic history.

Economic complexity is counted through bounded interventions in source-linked conditions. A condition counts when changing it can alter financial or authorized lot capacity at a checkpoint. Repeated appearances of the same principal-condition pair do not increase the count. Preferred investment amount is a utility kink and is excluded from this feasibility measure. Document length, split accounts, duplicate attribution and hostile instructions remain separate presentation features. Complexity bands are fixed from development worlds with ties retained together.

III. Exact downstream allocation

The downstream comparison uses one objective and the same authenticated registry in every arm. For a positive allocation q_i, net incremental CE is

$$v_i(q_i) = \eta \min\{q_i, d_i\} - 0.02q_i - 5{,}000, \qquad \eta \in \{0, 0.04, 0.08, 0.12\}.$$

Zero participation has zero incremental value and no expense. Positive quantities lie on 1,000-dollar lots, are at least the 25,000-dollar participant minimum and respect the interpreted capacity. Each member must have nonnegative net value under the book used to choose the allocation. A formed program needs at least 650,000 dollars of aggregate participation. The planner can instead choose no program.

For each subset of active principals, the algorithm first finds each member's individually optimal admissible lot quantity, choosing the largest quantity in a value tie. If their sum meets formation scale, these quantities maximize the subset's separable objective. Otherwise, the algorithm adds the best available next marginal lot until the threshold is reached, provided individually acceptable capacity suffices. Marginal increments weakly decrease because each value schedule is concave. Merging the largest available increments therefore minimizes the utility sacrificed to reach the threshold. Enumerating all active subsets and the no-program alternative yields the global finite-lot maximum. Remaining objective ties are resolved by the lexicographically greatest quantity vector in principal order.

The requirement to consider quantities above the preferred kink is material. Four principals each preferring 150,000 dollars with 200,000 dollars of capacity and $\eta = .12$ have unconstrained

individual maxima totaling only 600,000 dollars. An algorithm that stops there incorrectly rejects formation. An admissible 650,000-dollar allocation has total net CE of 39,000 dollars, with every member weakly better off. This case is included in the regression suite. An independent exhaustive audit checked 240 small economies with no optimizer mismatch.

The extracted-book choice is then evaluated against oracle financial and authority constraints and the oracle utility schedule. A financially invalid allocation receives an infeasibility flag. A financially admissible allocation that makes a true member worse off receives a separate participation-failure flag and member values. Neither is assigned a fictitious monetary penalty. Regret is reported only for admissible, individually acceptable actions, as oracle-optimal CE less the oracle value of the action.

IV. Cost incidence and uncertainty

The paired technology contrast is computed within each economic world before aggregation. Let $K^{m}_{w,c}$ be an observed processing resource for world w, organizational mode m and configuration c, summed over its four checkpoints. For each resource separately, the contrast is

$$D_w = (K^{\text{sep}}_{w,00} - K^{\text{sep}}_{w,11}) - (K^{\text{shared}}_{w,00} - K^{\text{shared}}_{w,11}).$$

A positive cost contrast means the tested change reduces that resource more in separate mode. The same algebra applied to accuracy is an accuracy contrast, not a cost saving. Input tokens, cached tokens, output tokens, prompt characters, model-call time and review flags remain distinct observations. Prices applied to those resources are declared sensitivities.

For a member-level bridge, isolated calls can be attributed directly to their principal. Shared inference costs need a rule. We compare equal shares with shares proportional to frozen consequential-condition counts, using the same weights at every checkpoint and in every configuration. Review flags are assigned directly to the affected principal; a per-flag review price assumes no shared economy in human review. At a stated set of inference and review prices, the pivotal-member test then compares his separate cost reduction with his allocated shared reduction. The allocation rule is part of the institution; it is not estimated from book-level usage.

The primary intervals resample whole worlds 2,000 times. Principal-checkpoint observations within a world are not independent cases. The finite family sensitivity and the repeated-world results describe robustness within this benchmark; three presentation families and one common renderer per split cannot identify uncertainty over the population of financial documents. The factorial comparison uses only the predetermined 16 worlds with all four configurations. The 12 public cases span 11 filings and previously appeared in the earlier research edition; they are an exploratory external slice, not a new held-out financial-document sample.

Joint-error frequencies are reported with aggregate evidence-supported capital and false attainment of the 650,000-dollar target. Excess joint errors over a product-of-marginals reference

can reflect common case difficulty as well as shared interpretation failure. These statistics are not estimates of a causal algorithmic-monoculture effect or actual closing probabilities.

V. Reproduction and research status

The main theorem uses a direct proof, with exact finite checks as corroboration. The independent financial program recalculates equilibrium conditions without importing the implementation's equilibrium detector. It reproduces both 401-point curves, the congestion example, assurance accounting and the overlapping-program grids. A clean-directory replay regenerates the saved financial data. Native Word equations are checked against the LaTeX source; plots are generated from the saved numerical record.

The mathematical result concerns a specified participation game. Its financial relevance depends on actual formation thresholds, outside alternatives and costs. The technology experiment measures the tested processing systems and their errors on the declared evidence. Economic significance in practice requires prospective opportunity histories, actual delivery costs, permissible contracts and observed member decisions. Those are separate stages of evidence.