

The Organization of the Investment Frontier

Independent Mandates and the Formation of Shared Investments

John Komkov¹

Working paper October 2026

Abstract

An improved outside option can eliminate a jointly valuable investment. Under scale economies, it raises worst-equilibrium welfare and compresses the equilibrium range. Two characterizations extend the argument. With private outside values, universal subscription requires strict participation and funded assurance for all but one indispensable investor. With continuous concave portfolios, acceptable capital determines formation, and the best individually acceptable allocation is already Nash. A paired AI experiment measures interpretation-cost incidence: shared processing has lower recorded inference costs, while the improvement saves more separately. The findings distinguish participation, selection and funding, and identify the institutions through which cheaper implementation changes investment opportunities.

Keywords: investment formation; portfolio choice; cost sharing; conditional commitments; artificial intelligence.

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

¹John Komkov has commercial interests in Megalith and Argosy. AI tools assisted with research, drafting, code and analysis.

1 Introduction

An improvement in what an investor can do alone can remove from equilibrium an investment that remains worth undertaking together. A shared program needs enough independently governed capital to form. An investor whose separate alternative improves may leave even when the others would gain enough from his continued participation to compensate him. His decision changes their opportunities as well as his own. The investment frontier depends on how investors are organized.

Direct investment and alternative vehicles make this question concrete. Fang, Ivashina and Lerner (2015) study 390 direct private equity investments made between 1991 and 2011 by seven institutions. Lerner et al. (2022) examine alternative investment vehicles alongside conventional private equity funds. These arrangements allow investors to combine delegated and direct exposure rather than choose one organizational form for their entire portfolio. The evidence documents the alternatives whose coexistence motivates the model. It does not establish that an improvement in direct investment caused a shared program to fail. That causal mechanism is the object of the theory and the proposed empirical test.

The first result is a compression theorem. Consider a threshold participation game with excludable benefits, independent outside options and fixed program terms under which additional members weakly improve existing members' payoffs. Improving one investor's outside-option value by Δ raises the lowest pure-equilibrium welfare by exactly Δ and never enlarges the range of equilibrium welfare. If the previous range is G , the highest equilibrium changes by an amount between $\Delta - G$ and Δ . Both bounds are attained. The floor can rise while the ceiling falls. The adverse case is *excess separation*: individual choices abandon a shared investment that still creates aggregate surplus.

The next two results identify how far this logic survives richer financial choices. With private outside-option values, information alone cannot eliminate nonparticipation when no investor can form the program alone and a failed subscription brings no reward. For a closed, indispensable cohort, an exact Bayesian characterization distinguishes participation from selection. Every type must strictly prefer full subscription, and all but at most one investor must receive a credible strategic-failure payment exceeding his subscription expense. These conditions are necessary and sufficient for full subscription in every Bayesian equilibrium, including mixed equilibria. A binding commitment by itself supplies neither condition. Risk aversion and a common funding failure can make success-only compensation inadequate at any finite amount.

With continuously divisible portfolios, investors choose genuinely interior allocations between the shared program and an outside investment. A second characterization identifies every pure equilibrium for a class with concave incremental values, fixed formation charges and indispensable investors. A formed equilibrium exists exactly when every investor has a nonnegative participation maximum and their largest acceptable contributions can meet the threshold. A fall in those acceptable contributions can destroy formation while every investor's individual payoff schedule still has a positive peak. The same model gives a useful limit: its best individually acceptable

coordinated allocation is already a Nash equilibrium. Better portfolio computation alone cannot improve that equilibrium. The institution must resolve selection, redistribute participation gains, change terms or secure funding.

Three attainable sets separate these tasks. The feasible set describes what resources, technologies and hard constraints permit. The equilibrium set describes what voluntary choices support under the baseline rules. The implementable set describes what a specified class of permissible contracts can support, net of its costs. Their undominated boundaries need not coincide. A contract may support an attractive outcome in one equilibrium without guaranteeing it in every equilibrium. Even universal subscription does not guarantee that the cash or the required consent will arrive.

Artificial intelligence enters through the incidence of organizational cost. Machine interpretation and persistent records can make separate implementation cheaper. They can also reduce the cost of a shared arrangement that preserves each investor's circumstances. The participation condition depends on which saving reaches the investor closest to leaving. A completed paired experiment in 64 synthetic economic worlds measures one component of that comparison. Records reduce repeated retrieval. Shared processing has lower recorded token costs, but the tested improvement saves more in separate processing. Successful primary contexts are at the scored semantic-extraction ceiling; most apparent exact-output improvement comes from normalizing abstention. All positive primary model-implied decision regret occurs after service failures. These findings locate a measurable cost channel without establishing actual investment fragmentation.

The theory draws on monotone games, cost sharing and assurance contracts. [Juarez \(2013\)](#) provides the closest willing-coalition construction behind compression. [Tabarrok \(1998\)](#) already treats assurance with private values, and [Halac, Kremer and Winter \(2020\)](#) studies capital assembly with heterogeneous investments and success- and failure-contingent payments. The results here give a welfare bound and two restricted characterizations suited to investment participation. They do not propose a general replacement for these theories. Their common contribution is to separate four questions that a portfolio optimum leaves open: who wishes to participate, which equilibrium is selected, what a permissible contract can change, and which resources remain available when called.

The unit of coordination can consequently be a shared investment even when the surrounding portfolios and their authority remain distinct. This is the sense of polycentric finance used here. Conventional portfolio choice remains essential, and the literature on financial innovation already makes claims endogenous. The additional organizational margin is whether separately governed demand can sustain a particular investment. Technology can improve what investors do separately while destabilizing what they achieve together.

2 Environment and attainable sets

2.1 Economic state, choices and authority

There are finitely many investors, three decision dates and finitely many states. At date 0 each investor adopts a mandate, chooses grants of authority and may enter a contract. Date 1 separates subscription from funding. The opportunity and public contract are observed before subscription; any later funding condition is learned before implementation. The Bayesian extension explicitly places subscription before the independent funding event. 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 investor's book before a new investment is evaluated.

An investor 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 investor, 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,

$$CE_i(X) = -\frac{1}{\alpha_i} \log \mathbb{E}[e^{-\alpha_i X}], \quad \alpha_i > 0. \quad (1)$$

A sure terminal transfer t satisfies $CE_i(X + t) = CE_i(X) + t$. 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.

2.2 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

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

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

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

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 investor gains. We report member payoffs and test participation against jointly feasible outside alternatives.

Weak supportability and robust implementation are different. The implementable-set definition asks whether a contract has an equilibrium supporting an outcome. A designer concerned about every equilibrium instead evaluates, for example,

$$\sup_{\substack{c \in C \\ \text{NE}(\Gamma_c) \neq \emptyset}} \inf_{a \in \text{NE}(\Gamma_c)} W(p(a; c)). \quad (4)$$

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.

For the Bayesian extension, a payoff vector is a vector of interim expected utilities or the corresponding explicitly constructed certainty equivalents, indexed by own type. The equilibrium concept then becomes Bayesian Nash equilibrium, including mixed strategies. Comparisons use the same prior and funding technology. A guarantee about signatures is stated separately from the probability of successful investment funding.

2.3 Portfolio choice and the closest organizational results

Separation and spanning explain when investors can obtain heterogeneous exposures through mixtures of common funds (Markowitz, 1952; Merton, 1971; Ross, 1978). A different preference need not require another vehicle. Desired exposure, restrictions on holdings and holder-specific histories are different reasons for differentiation. A long-only mixture cannot remove an excluded security when every constituent fund holds it; tax attributes and accepted obligations require attributable records. Which of these facts requires separate execution depends on the available legal and operational arrangement.

Security design already treats available claims as endogenous (Duffie and Rahi, 1995; Boyd and Prescott, 1986). Joint portfolio optimization already allocates shared trading costs across separate accounts (Iancu and Trichakis, 2014; Devanathan et al., 2026), and illiquid-commitment research already treats liquidity and pacing dynamically (Luxenberg et al., 2022). The formation question here concerns investors' ability to refuse a proposed joint solution. Shared fixed costs connect that choice to clubs and entry (Buchanan, 1965; Spence, 1976; Mankiw and Whinston, 1986). Neither a new product label nor a new computational architecture establishes a new economic opportunity.

Population-monotonic allocation and cross-monotone cost sharing supply the membership order used below (Sprumont, 1990; Juarez, 2013). The compression proof uses a union-closed family of willing memberships and derives its welfare envelope as outside values change. It does not inherit truthful reporting or coalition stability. In larger combinatorial environments, those are additional mechanism-design requirements (Birmas et al., 2019).

Monotone-game theory distinguishes a different order property. Topkis (1979), Milgrom and Roberts (1990) and Vives (1990) study complementarity of actions and equilibrium responses; Milgrom and Shannon (1994) provides the ordinal comparative-statics conditions. An increase in a member's payoff when others join a *formed* investment does not by itself establish increasing differences across the complete subscription game. Failed subscriptions must also be compared. The next section makes that boundary explicit rather than treating a threshold as either sufficient for, or incompatible with, supermodularity.

3 Compression

3.1 The participation game

There are n investors. 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:

$$S \subseteq T, \quad i \in S, \quad |S| \geq q \quad \implies \quad u_i(T) \geq u_i(S). \quad (5)$$

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 investors remaining outside is an equilibrium: one subscriber cannot meet the threshold. Define

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

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*

$$L(o') = L(o) + \Delta, \quad G(o') \leq G(o), \quad \Delta - G(o) \leq H(o') - H(o) \leq \Delta. \quad (7)$$

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

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

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 investors, 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 investor 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 . $\square$

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.

3.2 The complete game and increasing differences

Let $D_i(S)$ be the difference between subscribing and staying outside when the other subscribers are S . Before the threshold can be met, $D_i = -h_i$. Once the investor's entry would form the program, $D_i = u_i(S \cup \{i\}) - o_i$. Under inclusion-monotone formed payoffs, the complete game has increasing differences exactly when

$$u_i(T) \geq o_i - h_i \quad \text{for every minimal formed membership } T \text{ containing } i. \quad (9)$$

If the first formed membership gives less than the failed-subscriber payoff, the incremental return to subscribing falls at the threshold. Otherwise it rises or stays constant and then follows formed-payoff monotonicity. Appendix A proves this equivalence. Standard strategic-complementarity results apply in the latter region. Compression applies to the stated welfare envelope even when this transition condition fails. Thus ordering payoffs inside formed memberships, ordering full-game incentives, and bounding welfare are distinct claims.

3.3 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 investors, 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. An investor may manipulate his declared outside value, and the computed willingness test is not a truthful elicitation procedure.

3.4 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

$$s_k < \Delta < \sum_{i \in S} s_i. \quad (10)$$

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 investors 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.

4 Participation under technological change

4.1 Cost incidence at the pivotal investor

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

$$\Delta o_i \leq s_i + \phi_i(S) - \phi'_i(S) \quad \text{for every } i \in S. \quad (11)$$

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, (11) becomes

$$d_i^{\text{sep}} - d_i^{\text{sh}} \leq s_i. \quad (12)$$

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.

4.2 Proportional savings and endogenous membership

Let initial separate and allocated shared processing costs be c_i^{sep} and c_i^{sh} . A proportional reduction θ in both gives net incidence $\theta(c_i^{\text{sep}} - c_i^{\text{sh}})$. For a fixed membership with initial slack s_i , continued willingness requires

$$\theta(c_i^{\text{sep}} - c_i^{\text{sh}}) \leq s_i \quad \text{for every member } i. \quad (13)$$

Among members with positive cost differences, the first fixed-set departure occurs at the smallest ratio $s_i/(c_i^{\text{sep}} - c_i^{\text{sh}})$. A common percentage saving can therefore have a concentrated participation effect. A complex mandate is pivotal only if its cost incidence is large relative to its slack. Complexity alone is not a theorem about departure.

The ratio describes a fixed proposed membership. Under open entry, a new investor may reduce charges and preserve a participant who would otherwise leave. Figure 2 recomputes the full equilibrium correspondence as both outside and shared costs change; it does not extend a fixed-set cutoff beyond its assumptions. Its participation panel reports the largest willing set at each point. Entry is permitted, although this illustrated path starts with all five investors and exhibits only

departures. A separate saved entry case starts with gross values (12, 12, 12, 5), outside values (1, 1, 1, 0), threshold three and expense $24(1 - \theta)$. The fourth investor initially stays outside; he joins at $\theta = 1/6$, when his charge falls to 5. The largest willing membership then expands from three to four. This countercase shows why fixed-set exit ratios are not a general entry model.

The figure's synthetic five-investor economy has gross formed values (12, 13, 11.4, 12.75, 12) and outside values $\theta(4, 8, 12, 18, 26)$, in thousands of dollars. It needs three members. Figure 1 holds the common cost at 30; Figure 2 reduces it to $30(1 - \theta)$. A capable ordinary provider receives the same information, costs and implementation opportunities. No forecast advantage or private-return premium enters either comparison.

4.3 Supportability and compensation

For a fixed proposed net allocation g_i and transfer bounds $\underline{T}_i \leq T_i \leq \bar{T}_i$, a balanced participation transfer exists exactly when

$$\ell_i = \max\{o_i - g_i, \underline{T}_i\} \leq \bar{T}_i \quad \forall i, \quad \sum_i \ell_i \leq 0 \leq \sum_i \bar{T}_i. \quad (14)$$

This is feasibility of intervals with a zero sum. Real administrative expenses are first subtracted from the proposed allocation. A dated cash payment also requires liquidity at its payment date, not only terminal surplus.

Figure 1 permits named-member contracts with unanimous acceptance, transfers bounded by \$3,000 per member and an administrative charge of \$400 plus \$40 per member. The plotted recoverable gap is net of those costs. Excluded investors cannot enter that particular closed contract unilaterally. The class includes baseline open participation, so supportability cannot be worse than the baseline. These contracts support an all-accept equilibrium; they do not by themselves eliminate all-reject. The next section supplies the stronger implementation test.

A computed externality is not automatically an admissible charge. In particular, the ability to evaluate aggregate gains does not authorize restricting an investor's outside route. Fixed transfers, their budget and their acceptance game must be specified. Coalition stability is another condition: a mutually acceptable allocation can still be blocked by a subgroup with a better feasible arrangement. Neither a pure Nash equilibrium nor budget balance establishes that stronger property.




Figure 2: Cost incidence and membership. Common expense falls from 30 to $30(1 - \theta)$ while outside values rise to $\theta(4, 8, 12, 18, 26)$. Panel A enumerates every pure equilibrium at each point. Panel B marks membership in a largest best formed equilibrium, retaining ties. Entry is permitted in the calculation; this particular path exhibits only departures. All values and membership points come from the saved five-investor economy.

5 Information, institutions and credible capital

5.1 The information-only boundary

Uncertainty about other investors' willingness is different from uncertainty about whether a formed investment pays. Consider finitely many investors with private types drawn from a common prior. A type specifies an outside-option value and may also specify risk preferences. Investors observe their own types and any signals permitted by the information structure. If no investor can meet the formation requirement alone and signing an unsuccessful subscription gives no positive net reward, universal refusal is a Bayesian equilibrium. Given that everyone else refuses at every type, signing cannot form the program and cannot improve the investor's payoff. This argument applies at every information set. Greater signal precision, public disclosure of willingness or a recorded expression of interest does not remove this equilibrium.

Global-games selection changes additional primitives. In [Sákovics and Steiner \(2012\)](#), a noisy fundamental affects the investment threshold and the support includes regions of dominance. The limiting selection and the identity of important participants follow from that structure. [Morris and Shin \(2003\)](#) makes the role of dominance regions explicit. The present fixed-threshold financial game retains states in which one investor is physically unable to form the program regardless of his information. A uniqueness claim therefore requires a contract or another specified change to incentives; it does not follow from adding noise to private outside values.

5.2 A declared class of conditional contracts

Fix a cohort C of $m \geq 2$ investors, positive tickets q_i and a minimum K . The cohort is minimal winning:

$$\sum_{i \in C} q_i \geq K, \quad \sum_{j \in C \setminus \{i\}} q_j < K \quad (i \in C). \quad (15)$$

Each investor signs or refuses a public contract. Signing binds him to the specified ticket if the contractual conditions are met. No private type is reported. Types are finite, every own type has positive marginal probability, and their common prior may be correlated. Conditional on an investor's own type, his payoff from full subscription does not depend on the other investors' private types. Thus private outside opportunities are allowed, but interdependent information about the investment's return is not.

All signatures trigger a funding attempt. A separately verifiable condition Z succeeds with probability $p > 0$, independently of the type profile. For a signer, $s_i(t_i)$ is his success-state gain over his type-specific outside portfolio before a fixed success transfer T_i . Signing consumes expense $h_i \geq 0$ in every state. If the cohort fails to sign, the signer receives $\beta_i \geq 0$. If everyone signs but Z fails, he receives $\gamma_i \geq 0$. On success he receives $s_i(t_i) + T_i$. Non-signers retain their outside option. Strategic failure and external funding failure are observable contractual events, not labels assigned from private types.

Let $D_i(t_i)$ denote the expected-utility gain from signing when everyone else signs, integrating over Z . Let $e_i(t_i)$ denote the utility gain from the sure strategic-failure receipt $\beta_i - h_i$. Strictly increasing utility implies that e_i has the sign of $\beta_i - h_i$. At an arbitrary belief the gain from signing is a probability-weighted average of D_i and e_i . This reduction is why the timing, independent funding event and private-values restrictions matter.

Theorem 2 (Bayesian subscription under funded assurance). *In the stated closed-cohort game, all investors sign at every positive-probability type in every Bayesian Nash equilibrium if and only if*

$$D_i(t_i) > 0 \quad \text{for every } i, t_i, \quad \#\{i : \beta_i \leq h_i\} \leq 1. \quad (16)$$

The statement includes mixed equilibria. It guarantees subscription; the specified funding attempt still succeeds with probability p .

For sufficiency, the investors with $\beta_i > h_i$ strictly prefer signing whether the other signatures arrive or not. At least $m - 1$ investors therefore sign at every type. The remaining investor then strictly prefers signing by $D_i > 0$. For necessity, a negative D_i defeats universal signing; two weakly assured investors can refuse together without either being able to complete the cohort alone. The zero case is more subtle. A type with $D_i = 0$ can mix, and sufficiently infrequent refusal preserves the remaining investors' willingness. Appendix A gives the complete argument, including correlated types.

Strictness is economically consequential. A contract can have a unique fully subscribed *pure* equilibrium while also having a mixed equilibrium with incomplete subscription. In a two-investor boundary example, full-sign utility gains are $(0, 1)$ and strategic-failure gains are $(1, -1)$. All-sign is the only pure equilibrium. The first investor signing with probability $3/4$ and the second always signing is also an equilibrium. Reporting only pure equilibria would incorrectly certify universal implementation.

Binding commitments prevent accepted tickets from being revised unilaterally; they do not force investors to accept. Ordinary subscriptions and publicly recorded binding subscriptions retain universal refusal. Formation-contingent transfers can make full subscription attractive without removing refusal. Assurance can remove refusal without making an unattractive successful investment acceptable. The theorem states precisely when both margins have been addressed within this class.

5.3 Risk preferences, participation prices and transfer limits

A useful financial specialization uses CARA preferences over an incremental payment lottery that is independent of residual background wealth. Let $a_i > 0$ be fixed across investor i 's own types. His conditional certainty-equivalent gain from signing when everyone else signs is

$$C_i(t_i) = -\frac{1}{a_i} \log \left[p e^{-a_i(s_i(t_i) + T_i)} + (1 - p) e^{-a_i \gamma_i} \right] - h_i. \quad (17)$$

Its sign is the sign of $D_i(t_i)$. A CE already integrating a correlated background risk cannot simply be substituted into this lottery with physical probabilities. When funding and background wealth covary, the full terminal-wealth distribution must be evaluated instead.

Write $s_i^{\min} = \min_{t_i} s_i(t_i)$. If $e^{-a_i h_i} - (1-p)e^{-a_i \gamma_i} > 0$, strict participation at every type is equivalent to

$$T_i > r_i \equiv -s_i^{\min} - \frac{1}{a_i} \log \left(\frac{e^{-a_i h_i} - (1-p)e^{-a_i \gamma_i}}{p} \right). \quad (18)$$

If the numerator is nonpositive, no finite success-only transfer suffices. A large payment in the successful state cannot undo the utility cost of sufficiently likely uncompensated failure. For example, with $p = .1$, $a_i = 1$, $h_i = 1$ and $\gamma_i = 0$, the largest possible CE gain is $-1 - \log(.9) < 0$. The risk-neutral limit of the floor is $[h_i - (1-p)\gamma_i]/p - s_i^{\min}$. With private risk aversion, compute the floor separately by type and take its maximum.

Suppose transfers satisfy $\underline{T}_i \leq T_i \leq \bar{T}_i$ and $\sum_i T_i = 0$. Strict participation is feasible exactly when the following finite linear program has a strictly positive optimum:

$$\max_{T, \eta} \eta \quad \text{subject to} \quad T_i \geq r_i + \eta, \quad \underline{T}_i \leq T_i \leq \bar{T}_i, \quad \sum_i T_i = 0. \quad (19)$$

Here η is slack in success-state transfer dollars, not a common CE gain. Collection of negative transfers, settlement timing and contractual eligibility are additional financial requirements. Sure-money translation under CARA does not remove them. A fixed fee for administering the contract must also be paid before its proposed welfare improvement is evaluated.

5.4 The resources behind assurance

For $\gamma_i = 0$, the smallest strategic-failure reserve has a simple characterization. Conditional on strict full-sign willingness, its infimum within the declared nonnegative-bonus class is

$$R_{\text{sel}}^{\text{inf}} = \sum_{i \in C} h_i - \max_{i \in C} h_i. \quad (20)$$

Leave a highest-expense investor unassured, give every other investor $h_i + \varepsilon$, and let ε approach zero. Those $m - 1$ payments are simultaneously due if the remaining investor refuses. Necessity follows because at least $m - 1$ investors must be strictly assured. The infimum generally is not attained: exact equality restores a refusal equilibrium. This is an infimum over the stated closed cohort and payment class, not over all sequential mechanisms or possible cohorts.

The reserve is a maximum credible promise. Its equilibrium payout is zero when universal subscription is implemented, yet its carrying cost need not be zero. If external failure also pays γ_i , a supported all-sign failure state requires $\sum_i \gamma_i$. Taking the maximum of that amount and the strategic reserve is valid only when the states are mutually exclusive and the same funded reserve can legally be reused at the relevant dates. Capital needed to buy the investment is separate.

For required cash $K(\ell)$ and collectible internal funding $B(\ell)$ at settlement, a cash backstop needs maximum capacity $\max_{\ell}[K(\ell) - B(\ell)]_+$. Expected payout averages the state shortfalls; an economic price values them using the provider's state prices; carrying cost prices capital held ready. A cash backstop cannot repair missing issuer consent or ineligible ownership. The distinction follows the aggregate-resource logic of [Holmström and Tirole \(1998\)](#). Accepted obligations also survive a later preference change or revocation of permission for new transactions.

5.5 A finite private-information economy

The four-investor example uses tickets of \$80,000, \$100,000, \$120,000 and \$100,000. The minimum is \$350,000, so the proposed \$400,000 purchase requires all four. Values in this paragraph are thousands of dollars. Gross success values are (14, 13, 16, 12), and each investor has two equally likely outside values: (6, 10), (6, 10), (7, 11) and (4, 12.2). Subscription expenses are (.5, .8, .6, 1) and the common funding condition succeeds with probability .9. The high type of investor 4 dislikes success at the initial terms.

Success transfers $(-.6, -.5, -.5, 1.6)$ balance and make every type strictly willing under both risk neutrality and CARA with $a_i = .02$ per thousand dollars. Strategic-failure payments (.51, .81, .61, 0) then implement universal subscription. Their maximum reserve is \$1,930, compared with the \$1,900 infimum; carrying it for 90 days at an assumed annual 4% costs \$19.04. Expected strategic bonuses are zero. Common funding failure remains at 10%; in that state none of the planned purchase capital is collectible. A backstop guaranteeing completion of the stated \$400,000 purchase would need \$400,000, with an expected gross funding draw of \$40,000 in this two-state specification. The draw is neither a net credit loss nor a quoted guarantee price. These are different obligations from funding the smaller minimum or insuring failed subscription.

A transfer cap of \$500 per investor cannot support every type; a cap of \$2,000 can. Pure Bayesian enumeration checks ordinary subscriptions, binding subscriptions, transfers alone, assurance alone, the combined contract and a contract with two unassured investors. The all-equilibrium conclusion rests on Theorem 2, not on the absence of another pure equilibrium in a finite search.

5.6 What is inherited from contracting theory

[Tabarrok \(1998\)](#) already establishes assurance with excludable goods, private values and subscription costs. [Halac, Kremer and Winter \(2020\)](#) develops a substantially richer capital-assembly problem with heterogeneous investment sizes, success and failure payments, issuer capital and unique implementation. Its iterative-dominance argument is a close antecedent. [Segal \(2003\)](#) and [Bernstein and Winter \(2012\)](#) study differentiated contracting with externalities and the cost of recruitment. Bernstein and Winter also discuss success rewards paired with failure reimbursement. The current closed-cohort result does not replace these mechanisms or their optimal-contracting analyses.

Table 1: Institutional diagnosis: different failures require different promises

Failure	Relevant instrument	Exact boundary
Selection	Funded strategic-failure bonus	Full-sign utility must be strictly positive. All but at most one indispensable investor need a bonus exceeding expense.
Participation	Success-contingent transfer	The worst-type transfer floor must be finite and compatible with transfer limits and the budget.
Cash funding	Prefunding or collectible backstop	Sufficient cash must arrive in every covered state. Maximum capacity, expected payout and price differ.
Permission or consent	Current grant and satisfied legal condition	Compensation does not enlarge authority; cash does not cure missing consent.
Quantity revision	Binding accepted ticket	A new preference cannot rewrite acceptance; expiry can block a new contract.

The Bayesian characterization concerns a closed minimal cohort and observable failure categories. An assurance reserve pays contingent bonuses; it is not the capital required to purchase the investment. Support in an all-accept equilibrium is weaker than implementation in every equilibrium.

Its narrower role is an exact classification with finite private outside values, ordinary mixed equilibria, separately verifiable strategic and funding failures, strict participation, bounded transfers and a funded-reserve infimum. The ability to distinguish failure causes is itself a contractual assumption. If it cannot be verified, the payment class must change. The theorem does not identify an optimal ownership structure for the coordinator, establish truthful elicitation of hidden types or prescribe the cheapest mechanism outside this class.

6 Continuous portfolios and the scope of compression

6.1 Interior allocations and indispensable capital

An investor who improves his outside investment need not withdraw his entire shared allocation. To distinguish adjustment from formation, let investor i choose a continuous contribution $q_i \in [0, b_i]$. The program forms if $\sum_i q_i \geq K$. Capacities satisfy $\sum_i b_i \geq K$ and $\sum_{j \neq i} b_j < K$ for every i : every investor is indispensable, even though his contribution is divisible. Capacities are net of protected intermediate liquidity and existing obligations. Fixed formation-cost shares are paid on success. They do not change with other investors' chosen quantities.

The formed payoff is $O_i + f_i(q_i)$, where O_i is the optimized outside-option value. The incremental function f_i is continuous and strictly concave, with a unique maximizer $d_i \in (0, b_i)$ and $f_i(0) < 0$. The negative intercept represents a fixed participation expense, charged only if a

positive participation forms. A failed reservation refunds capital but consumes $\rho_i q_i$, with $\rho_i > 0$, giving payoff $O_i - \rho_i q_i$. There is no aggregate upper capacity or congestion. Preferences over the formed portfolio are separable across investors after the fixed expense shares have been assigned.

These assumptions admit genuine mixing. One representation is

$$f_i(q; t) = P_i(q) + v_i(B_i - q; t) - v_i(B_i; t) - F_i - \rho_i q. \quad (21)$$

Here P_i is the value of the shared position, v_i is the value of an outside investment he can retain alongside it, B_i is available investment capital and F_i is his fixed expense share. Concavity of the portfolio components produces concavity of f_i ; a strictly concave component gives strictness. A separate liquid reserve is retained before determining B_i and b_i . The representation is an admissible reduced-form portfolio valuation, rather than a claim that arbitrary correlated risks decompose this way.

Theorem 3 (Continuous formation and individual acceptability). *In this indispensable-cohort game, all-zero is the only failed pure equilibrium. A formed profile q is a pure equilibrium exactly in either of these cases:*

1. $\sum_i q_i > K$, $q_i = d_i$ and $f_i(d_i) \geq 0$ for every i ;
2. $\sum_i q_i = K$, $q_i \geq d_i$ and $f_i(q_i) \geq 0$ for every i .

If all peaks $f_i(d_i)$ are nonnegative, define

$$R_i = \max\{q \in [d_i, b_i] : f_i(q) \geq 0\}. \quad (22)$$

A formed equilibrium exists if and only if all peaks are nonnegative and $\sum_i R_i \geq K$. Moreover, every individually acceptable formed allocation is weakly Pareto-dominated by a formed equilibrium. Consequently the best pure-equilibrium surplus above the outside floor is

$$G = \max \left\{ 0, \max_{\substack{0 \leq q_i \leq b_i, \sum_i q_i \geq K \\ f_i(q_i) \geq 0 \forall i}} \sum_i f_i(q_i) \right\}, \quad (23)$$

with the inner value absent when its feasible set is empty.

The proof uses two properties. In a failed profile, a positive reservation can be canceled to avoid its expense. At a formed profile with slack, a quantity away from d_i can move toward it without destroying formation. Exactly at the threshold, a quantity below d_i can be increased profitably, whereas any reduction destroys formation. A contribution above d_i can therefore be sustained by the shared threshold, provided its payoff remains individually acceptable.

The welfare equality is a substantive limit. Starting from an acceptable allocation, first raise quantities below their individual maxima. If the sum of the maxima reaches K , the vector d improves everyone weakly. Otherwise reduce only quantities above their maxima until the threshold

binds. These reductions improve their owners' payoffs and yield an equilibrium. The best acceptable allocation is therefore already supported by the baseline game. The coordinator can improve a selected bad equilibrium, but cannot claim an additional portfolio-optimization gain over the best one under these same terms.

The equilibrium set need not possess a greatest quantity vector. With two investors, capacities one, $K = 1.6$ and $d_1 = d_2 = .6$, suitable strictly concave positive participation values support $(.6, 1)$ and $(1, .6)$. Their componentwise maximum $(1, 1)$ is not an equilibrium: a small reduction moves an investor toward his preferred quantity while formation remains slack. Membership monotonicity in the binary game does not establish a lattice of continuous contributions.

6.2 Two technological changes

An improvement available only upon complete exit raises O_i by δ_i and subtracts δ_i from every positive-participation incremental value. A mixable outside improvement enters both terms involving v_i in (21). They are different experiments. If v_i has increasing differences in retained capital and technology, the larger all-outside capital position benefits at least as much as any $B_i - q$. Thus $f_i(q; t') \leq f_i(q; t)$ pointwise when technology improves. In a differentiable interior model, $v_{wt} \geq 0$ gives $f_{qt} = -v_{wt} \leq 0$ and a nonincreasing preferred contribution d_i . This is a standard monotone-comparative-statics argument (Milgrom and Shannon, 1994); concavity alone does not supply the sign.

Corollary 4 (Continuous compression). *Keep the cohort, capacities, threshold and failed-reservation rule fixed. Suppose $O'_i \geq O_i$ and $f'_i(q) \leq f_i(q)$ pointwise, with both economies satisfying Theorem 3. Let $\Delta = \sum_i (O'_i - O_i)$. Then*

$$L' = L + \Delta, \quad G' \leq G, \quad \Delta - G \leq H' - H \leq \Delta. \quad (24)$$

The new individually acceptable set is contained in the old one and its objective is weakly lower at every common allocation. Equation (23) therefore gives $G' \leq G$. This proof does not require a greatest equilibrium quantity vector. It also identifies the boundary: declining per-unit charges, changing capacities or a more attractive complementary hedge do not by themselves give the required pointwise order. A hedge that is especially useful while holding the shared investment can increase f_i , and the equilibrium gap can expand even if the outside value also rises.

6.3 A cliff with interior mixed portfolios

A four-investor example gives each investor two units of available risky capital and a separate half-unit protected liquid reserve. One capital unit is \$100,000. Each investor also has a separate \$20,000 expense account available on identical terms to every route; fixed formation and reservation expenses are paid from that account, and unused cash remains liquid. Shared capaci-

ties are (1, 1.1, .9, 1.2) and the program needs 3.5 units. Payoffs below are thousands of dollars. Incremental values are

$$f_i(q_i; \theta) = A_i(\theta)q_i - 10q_i^2 - F_i, \quad (25)$$

where $A(\theta) = (17 - \theta, 18.4, 15.6, 20)$ and $F = (6, 2, 2, 2)$. These values arise from (21) with a quadratic outside investment, a concave shared investment and reservation expense $.2q_i$. Only investor 1's outside technology improves. His outside value rises by 2θ ; the other outside values remain 12. The investment changes no legal constraint or authority.

At $\theta = 0$, preferred contributions are (.85, .92, .78, 1), totaling 3.55. Every choice is interior, every portfolio retains the outside investment and liquid reserve, and the threshold is slack. At $\theta = 1.5$, the best equilibrium instead binds the threshold, with contributions (.78125, .92625, .78625, 1.00625). At $\theta > 17 - \sqrt{240} \approx 1.5081$, investor 1's best participation gain becomes negative, so no formed equilibrium remains. Immediately beyond that threshold the program still yields more than \$18,000 of aggregate incremental value at the unconstrained coordinated allocation. Partial adjustment precedes the cliff but does not eliminate it.

A distinct three-investor case isolates a capital shortfall without a negative peak. Let $b_i = 2$, $K = 4.7$, $d = (.5, 1.5, 1.5)$ and $f_i(q) = a_i - (q - d_i)^2$ with peaks $a = (.2, 2, 2)$. Initially the largest acceptable quantities sum to 4.9472. An exit-only improvement of .17 for investor 1 reduces that sum to 4.6732, below the threshold, although all peaks remain positive. The feasible allocation (.7, 2, 2) still has increments (−.01, 1.75, 1.75), totaling 3.49. Compensating a \$10 loss can preserve \$3,490 of aggregate surplus in this synthetic economy. Neither a smoother optimizer nor a better equilibrium-selection device supplies that compensation by itself.

6.4 Combining private values and continuous allocation

At $\theta = 1.6$ in the four-investor case, freeze the coordinated quantities as a contractual offer:

$$q = (.7775, .9275, .7875, 1.0075).$$

Its success increments before transfers are, approximately,

$$s = (-.0716, 6.4634, 4.0834, 7.9994).$$

For this combined synthetic construction these are fixed success-state incremental payments, independent of residual wealth and the funding event. This assumption permits the CARA lottery calculation; arbitrary portfolio certainty equivalents cannot be mixed that way. Add two private outside-value increments per investor, (0, .2) for the first and (0, .4) for each other investor. These are additional exit-only type differences. Success transfers of .4 to the first investor, financed equally by the others, restore strict full-sign willingness when the common funding probability is .9, signing costs .01 each and CARA risk aversion is .02 per thousand.

Strategic-failure bonuses of .011 to the first three investors and zero to the fourth secure all signatures by Theorem 2. The worst-type CE gains, net of signing costs, are approximately (.1056, 5.2944, 3.1735, 6.6573). The reserve for strategic bonuses is \$33; the investment itself requires \$350,000 and formation expenses of \$12,000 are already deducted in the portfolio values. Common funding failure remains at 10%.

This construction connects the branches through a fixed, acceptable portfolio offer. After acceptance, quantities are contractual tickets, not freely revised strategic choices. The result does not characterize a simultaneous Bayesian game in which agents both choose continuous quantities and report hidden values. It shows that a continuous formation loss can be repaired by a contract whose information and funding conditions are independently checked.

6.5 Congestion, entry and overlapping opportunities

Scale economies concern the contractual effect of an entrant on incumbents' net payoffs. In a three-investor congestion example, membership charges rise with the number of participants. Raising one investor's outside value by 6 increases the best equilibrium from 20 to 30 and the equilibrium gap from 16 to 20. Exit relieves congestion. This is a counterexample to an unrestricted compression claim, rather than a sign theorem for all substitute participants.

Entry is excluded from the indispensable-cohort characterization and can change it materially. An entrant may replace a departing investor, improve cost shares or supply enough acceptable capital to keep formation feasible. The fixed-membership pivotal condition must then be recomputed. A capable conventional custom vehicle with the same admissible allocations and contracts reproduces the modeled outcome.

The preserved overlapping-program experiment permits four investors to divide their capital between two programs. Reducing program B's formation cost from 12 to 2 changes the best equilibrium welfare from 37 to 28 while raising coordinated welfare from 37 to 39. The improved coordinated allocation uses both programs in two investors' portfolios. Half-unit and quarter-unit grids reproduce these extrema. This is a fixed-price numerical counterexample, not a general-equilibrium result: capital constraints connect the programs, but prices and supplier entry are not determined jointly.

The new continuous characterization is analytic. Verification separately checks 193,116 quantity profiles against analytic unilateral deviations and 400 pointwise payoff deteriorations. An independently written evaluator checks 538,002 profiles in 810 additional quadratic games. Refining the canonical allocation grid to 1/160 of a capital unit leaves maximum investor CE error below \$0.18 and aggregate error below \$0.04 relative to the analytic solution, within the \$25 and \$100 acceptance thresholds. Economic minimums and fixed expenses are kept distinct from these numerical grids.

7 AI cost incidence

7.1 Matched organizations and a frozen experiment

The technological question is whether the same investor becomes cheaper to serve separately or within a shared arrangement that preserves his individual state. The completed experiment compares those modes using identical economic histories. It measures interpretation and record maintenance, one component of the organizational cost in the participation condition. It does not measure investment forecasting or the full cost of a legal vehicle.

There are 64 held-out economic worlds, each with four investors and four chronological checkpoints. Separate processing isolates each investor’s context and sums its costs. Shared processing reuses common offering evidence while maintaining separate states and decisions. Both perform equivalent required work. The primary comparison is retrieval without persistent semantic records or deterministic validation (R0V0) against the complete workflow (R1V1). A full factorial on 16 predetermined worlds separates records and validation. Four predetermined worlds are repeated. The evaluation contains 1,344 synthetic book-processing episodes: 1,024 primary, 256 factorial-completion and 64 repeat episodes. Twelve public-document cases add 48 episodes.

An episode can require several inference calls. The archived synthetic evaluation has 5,293 calls; the public slice has 72, including charged record preparation. Development is separate. Model-produced records incur construction, update and retrieval costs; they are not corrected by an oracle. Each checkpoint supplies only then-available evidence. Synthetic truth comes from executable state transitions and deterministic documents. The public cases are an exploratory, previously used slice with source-linked, agent-audited labels. They are not newly held-out or human-validated cases.

The source catalog and authenticated permission registry are common trusted inputs. Validation tests provenance scope, dates, grant expiry, remaining authority, expenses and contractual lots. Documents cannot create permission by asserting it. Validation does not establish that every plausible cited amount was extracted correctly. Acquisition and authentication of the trusted inputs are unmeasured. Human review prices remain assumptions because no expert review time was observed.

Economic complexity counts distinct, source-linked conditions whose perturbation changes feasible participation. Presentation difficulty and document length are separate. Frozen tiers contain 38 low-, three middle- and 23 high-complexity worlds. The economic world is the statistical unit; its investors and checkpoints are repeated observations. Paired world bootstrap intervals condition on this simulator and its few presentation families. The three-world middle tier is descriptive.

7.2 Quality, repeated work and operational failure

Table 2 preserves the frozen complete-output scores. A decision is exact only when capacity, recommendation, authorization, proposed amount and status all match. Every successful primary

Table 2: Primary processing results

	Separate		Shared	
	R0V0	R1V1	R0V0	R1V1
Exact final decisions / 1,024	946	1,007	974	1,016
Exact raw decisions / 1,024	946	955	974	980
Inference calls	1,791	1,429	448	353
Failed contexts	3	4	0	1
Recorded dollars / world	0.3311	0.2970	0.1457	0.1398
Mean regret / checkpoint (\$)	106.25	365.23	0.00	285.94

Each arm contains 64 four-checkpoint worlds. Dollar costs price returned usage at \$2, \$0.20 and \$8 per million uncached input, cached input and output tokens. Failed-call usage is missing. Regret uses the common synthetic downstream objective with $\eta = .12$; it is not a realized return or an observed service benefit.

processing context has all scored semantic state fields correct. No unsupported raw or final action appears in a primary arm.

Most of the exact-output improvement is a narrow encoding correction. Validation converts a null proposed amount to zero while retaining review status. It makes 53 such corrections separately and 36 jointly, adding 52 and 36 fully exact outputs. It blocks no positive proposal and changes no status. A post-hoc equivalence rule that treats null-under-review and zero-under-review as the same abstention yields 1,010 versus 1,007 exact decisions separately and 1,020 versus 1,016 jointly. The apparent improvement disappears under that representation; the original scores remain unchanged.

Persistent records reduce repeated retrieval. Primary calls fall by 362 separately and 95 jointly. On the factorial subset, the records main effect is -5.28 calls per four-checkpoint book separately and -1.56 jointly. The validation main effect on complete-output accuracy is 5.86 percentage points in each mode, reflecting the same encoding correction. Accuracy interaction intervals include zero. Fewer calls do not necessarily imply lower summed call time; shared-mode call time rises in the complete configuration.

All 192 public-slice answers are correct, another accuracy ceiling. Predetermined repeats reproduce 248 of 256 investor decisions exactly. Eight synthetic calls fail: seven service-capacity errors and one 180-second timeout. All remain in scheduled denominators, with no outcome-dependent retry. A shared failure can withhold several investors' decisions in one response. The observed timeout is such an operational event; it is not evidence that successful interpretations share a common semantic error or that financial markets have acquired a new common factor.

7.3 Which saving reaches the investor?

At the token prices in Table 2, the separate-minus-shared reduction in recorded cost is \$0.0281 per four-checkpoint world, with a paired 95% interval [\$0.0173, \$0.0393]. Shared processing is cheaper in level; the tested technology change saves more separately. Leaving any one presen-

tation family out preserves that direction. These amounts are applied prices on recorded usage, not invoices. Missing usage for failed calls makes each recorded bill incomplete; the difference between two savings need not be a lower bound.

The complexity hypothesis is unsupported. Mean incidence is \$0.0290 in the low tier, \$0.0470 in the three-world middle tier and \$0.0242 in the high tier. High-minus-low intervals include zero. The experiment does not show that the most complex mandates receive the largest advantage from separation.

Review can dominate the token bill. Pricing each final review flag at \$5 produces mean incidence of \$0.263 with interval $[-\$0.277, \$0.961]$; pricing it at \$25 gives \$1.200 with interval $[-\$1.527, \$4.711]$. Flags include appropriate abstention and missing outputs as well as correctable errors. These sensitivities neither measure labor nor impose equal achieved quality.

The theoretical bridge is member-specific pass-through. Under equal allocation of shared processing costs, separate savings exceed allocated shared savings for 164 of 256 investor histories; allocation by frozen condition counts gives 166. Both balanced rules produce mean incidence of \$0.0070 per investor history, but individual signs differ. For a fixed contract, a positive incidence $d_i^{\text{sep}} - d_i^{\text{sh}}$ consumes that amount of participation slack. In the continuous model, reductions in costs incurred only upon exit shift f_i downward; costs that also change the retained outside portfolio must instead enter both terms in (21). The experiment observes neither an actual investor's initial slack nor a provider's pass-through policy. Its cent-scale contrasts cannot be converted into annual economic gains without additional frequency, quality and service-cost evidence.

7.4 From interpreted state to financial consequences

A common downstream engine selects allocations using either interpreted or oracle state. It accounts for fixed and proportional expenses, a formation minimum and contractual lots. It then evaluates the selected action against the oracle. Regret is assigned only to financially feasible, individually acceptable choices; unsupported or loss-imposing actions are separately reported. The utility schedule is synthetic, and its parameters are described in the Internet Appendix and all member outcomes are retained in the replication record.

At $\eta = .12$, the oracle forms a program in 182 of 256 checkpoints; interpreted-state decisions form in 181, 181, 182 and 181 across the four primary arms. Every selected action is feasible and individually acceptable. All positive primary regret occurs in books with a service failure. The single shared-mode timeout accounts for all \$882,000 of withheld valid participation in shared R1V1 and its positive regret. No primary arm falsely attains the \$650,000 funding target, so the prespecified joint unsupported-action endpoint is zero. This is an absence of such errors in the tested histories, not a measured reduction in economy-wide correlated risk.

The study therefore supplies cost incidence, a semantic ceiling and an operational failure record. Its strongest positive result is less repeated retrieval. Its main caution is that amortizing work can concentrate the consequences of a missing response. Both quantities belong in the economics of implementation; neither establishes which institutional form investors should

adopt. The archived model identifier, access timestamps, prompts, evidence and usage permit exact rescoring of the recorded experiment. They do not identify immutable proprietary model weights or guarantee identical future responses.

8 Empirical implications and conclusion

8.1 What would identify the mechanism?

The theory predicts responses to contractual effects, not to asset-class labels. A private fund seeking its first close can be limited by minimum scale; the same manager's mature co-investment may be congested. A secondary purchase may depend on enough eligible sellers and compatible buyers, or its supply may already be rationed. A credit facility may require a minimum committed amount and then encounter concentration limits. A broad public-securities portfolio often has slack formation requirements, but a specialized mandate can still incur meaningful fixed costs.

A prospective test should classify these constraints before outcomes are known and retain every considered opportunity, including refusals and failed closings. For each decision date, the record needs the then-available outside route, the investor's requested allocation, net participation terms, authority, existing obligations and funding state. A later offer or document cannot be used to improve an earlier information set. Program size alone cannot identify whether entry helps incumbents: the relevant observation is how one more compatible participant changes their net terms.

Three contrasts would establish economic importance. First, changes in the cost of separate implementation should matter more for programs with binding formation requirements than for comparable opportunities whose minimum scale is slack, after accounting for changes in shared implementation. Second, departures should be associated with small participation slack and adverse member-level cost incidence. Complexity is a candidate determinant of that incidence, not a substitute for measuring it. Third, the use of transfers and assurance should depend on whether the failure is participation or selection. A selection device cannot make a negative success value acceptable, and compensation cannot buy a missing legal permission.

These predictions require variation not supplied by the synthetic economies. Outside opportunities can improve because expected returns change, access becomes available, fees fall or liquidity improves. Some changes also alter the shared program's payoffs. A credible design must separate those channels. The complete event history should distinguish investor refusal, diligence, documentation, consent, funding and settlement. Otherwise a program that fails for operational reasons can be misclassified as a participation failure.

8.2 Institutions that already preserve differences

The conventional counterfactual is capable. Investors can combine funds, direct investments and custom vehicles, and existing intermediaries can batch records or negotiate conditional commit-

ments. Partnership capital accounts, contributed-property allocations and partner-specific basis records already preserve differentiated state within shared ownership (Internal Revenue Service, 2025a,b). A capital account is not necessarily the investor's outside basis. Exchange funds offer a concrete example of shared investment with attributable contribution histories and restrictions (Morgan Stanley, 2023). These capabilities do not imply that every investor retains independent authority over each underlying trade.

Subscription-credit agreements likewise recognize heterogeneous capital. The filed agreement discussed in the evidence ledger differentiates investor eligibility, advance rates and concentration limits (Goldman Sachs Middle Market Lending Corp. II, 2022). Such terms reveal how contractual capital enters a borrowing base. Their prices combine credit, collateral, legal, maturity and liquidity effects; they are not a pure market price of credibility. They nevertheless provide a practical source of restrictions for a formation model.

Related market examples require the same discipline. Johnston and McConnell (1989) documents the decline of a futures contract in relation to design and competing hedges. It motivates endogenous participation, but does not identify cheaper bilateral customization as a cause of compression. Risk-sharing pools can benefit from compatible entrants, while unsuitable new exposures or scarce capacity can reverse the incidence (Blier-Wong and Lauzier, 2026). Public versus private, pooled versus separate, and digital versus manual are therefore insufficient classifications.

8.3 The limits that matter

The binary compression theorem fixes program terms and permits open entry under monotone formed payoffs. The Bayesian result instead fixes a closed minimal cohort, private values, independent verifiable funding risk and observable failure categories. The continuous result fixes indispensable investors, capacities, separable concave payoffs and expense shares. These are separate environments. The combined construction freezes a continuous allocation into a binding offer; it is not a general Bayesian portfolio equilibrium theorem.

The welfare measure sums CE dollars. Admissible sure transfers move CE dollars one-for-one, while real expenses reduce their sum. Success-contingent transfers in the Bayesian model instead balance cash in the success state; their ex-ante CE changes need not sum to zero because they reallocate risk. Supplier rents, tax effects and externalities require their own accounting if the objective is social welfare. CARA supplies sure-money translation, but implementability still requires admissible, collectible transfers. Under CRRA the same financial transfers can be made, while the CE arithmetic changes. None of the results permits an optimizer to enlarge authority or erase a prior obligation.

A coordinator earns its cost by delivering something the capable incumbent cannot provide as cheaply: acceptable terms, better equilibrium selection, dependable funding or compatible execution. Ownership, conflicts and incentives then matter. A low reserve carrying cost does not identify who should own the coordinator; the broader costs of contracting and control enter that

question (Hansmann, 1996). An affiliated supplier can capture savings or steer scarce capacity. Those effects must be measured against investor outcomes and specified allocation rules.

8.4 Conclusion

The investment frontier is partly an organizational outcome. Under scale economies, a better outside option raises the equilibrium floor and compresses the welfare range. It can remove a jointly valuable investment from equilibrium. Private information preserves a nonparticipation equilibrium unless incentives change. Within a declared contract class, strict full-sign willingness and funded assurance for all but one indispensable investor characterize universal subscription. Physical funding remains a separate condition.

Continuous mixed portfolios preserve the same economic tension under a pointwise order on incremental participation values. Their exact equilibrium characterization also limits the scope for coordination: the best individually acceptable allocation is already Nash. The institution adds value by selecting a better outcome, changing participation terms or securing execution, not by relabeling that existing optimum.

The AI experiment finds less repeated retrieval and lower recorded costs for shared processing, but greater marginal savings from the tested technology in separate processing. Successful semantic extraction is at a ceiling; service failures account for positive primary regret. Economic significance requires observed participation slack, actual cost pass-through and complete opportunity histories. Technology can improve what investors do separately while destabilizing what they achieve together. Institutions determine which investors receive the savings.

A Proofs and computational characterizations

A.1 Increasing differences at formation

Fix investor i and order the other subscribers by inclusion. His payoff difference between subscribing and remaining outside is $-h_i$ when his subscription still fails and $u_i(S \cup \{i\}) - o_i$ when it forms. Within the failed region the difference is constant. Within the formed region it is non-decreasing by net-payoff monotonicity. Every transition between the regions passes through a minimal formed set containing i . Consequently increasing differences holds if and only if every such set satisfies $u_i(S) \geq o_i - h_i$. A violation supplies two ordered profiles at which the difference falls. Satisfaction makes each transition nondecreasing. The finite binary action lattice then meets the relevant strategic-complementarity condition. A threshold alone does not decide the issue.

A.2 Bayesian subscription: necessity including ties

Sufficiency follows by strict dominance for the investors with $\beta_i > h_i$ and then strict optimality for the possible remaining investor. For necessity, write $e_i(t_i)$ for the strategic-failure utility increment. If any $D_i(t_i) < 0$, universal signing is not a Bayesian equilibrium. A mixed equilibrium exists in the finite Bayesian game, so universal signing cannot hold in every equilibrium. If at least two investors have $e_i \leq 0$, let all such investors refuse at every type and all others sign. Each refuser still faces another refuser; signing cannot complete the cohort and gives no positive gain. Every signer has positive e_i . The profile is an equilibrium.

It remains to consider all $D_i(t_i) \geq 0$, at most one weakly assured investor, and a zero D at some positive-probability type. Call the possibly weakly assured investor j . If j has a zero- D type, let those types refuse and its positive- D types sign; all assured investors sign. Investor j is optimizing because everyone else signs. Each assured investor's signing gain is a convex combination of a nonnegative D_i and a strictly positive e_i , so signing remains optimal.

Otherwise every type of j has strictly positive D_j . Choose a zero- D type of an assured investor i , let that type refuse with probability $\varepsilon > 0$ and let all other types and investors sign. Investor i is indifferent because everyone else signs. Every other assured investor remains willing for the same mixture reason. Conditional on any type of j , the refusal probability is at most ε . Finitely many strictly positive D_j values imply a sufficiently small common ε for which every j type continues to sign. If there is no weakly assured investor, that last restriction is unnecessary. The resulting mixed equilibrium has positive failure probability. Correlation of types changes the conditional probabilities but not the bound by ε . This completes necessity and the theorem.

A.3 Transfer floors and reserve infimum

Strict positivity of (17) is equivalent to

$$pe^{-a_i(s_i(t_i)+T_i)} < e^{-a_i h_i} - (1-p)e^{-a_i \gamma_i}. \quad (26)$$

The left side is strictly positive at every finite transfer. A nonpositive right side therefore rules out strict willingness. With a positive right side, taking logarithms and maximizing the required transfer across types gives (18). Closed finite transfer bounds make the feasible domain of (19) compact in T . A vector strictly above every floor has a positive minimum slack and conversely; the LP optimum is positive exactly in that case.

For the reserve bound, Theorem 2 requires at least $m - 1$ strict strategic assurances. Those investors may all sign while the remaining investor refuses. The liability in that supported action profile exceeds the sum of their h_i . The least such sum is $\sum_i h_i - \max_i h_i$. Leaving a highest-expense investor unassured and paying $h_i + \varepsilon$ to the rest approaches the bound. No profile of strategic failure can contain more positively insured signers, so that amount suffices. The theorem requires strictness, explaining the infimum. This proof concerns the payment reserve, not the financing of tickets, expenses or collectible success transfers.

A.4 Continuous equilibrium and the acceptable frontier

Indispensability implies that every formed profile has $q_i > 0$ and that exit by any investor destroys formation. Exit therefore gives O_i , so formed equilibrium requires $f_i(q_i) \geq 0$. If the program fails and some $q_i > 0$, reducing it to zero saves $\rho_i q_i > 0$; hence the only failed equilibrium is all-zero.

At a formed profile with $\sum_i q_i > K$, any investor away from his unique maximizer can move a small amount toward it without losing formation. Thus $q = d$ is necessary. It is sufficient when all peaks are nonnegative: no formed deviation improves on a peak, while a failed deviation gives at most O_i . At a threshold profile, $q_i < d_i$ permits a profitable increase. If $q_i \geq d_i$, every increase weakly lowers the formed value and every decrease destroys formation. Individual acceptability then suffices. This proves the equilibrium characterization, including zero-gain ties.

If any peak is negative, no acceptable formed profile exists. Otherwise the acceptable upper-branch interval is $[d_i, R_i]$. If $\sum_i d_i \geq K$, the vector d is an equilibrium. If $\sum_i d_i < K \leq \sum_i R_i$, continuity permits choosing quantities from those intervals summing exactly to K , again an equilibrium. Necessity follows because no individually acceptable quantity can exceed R_i .

Take an arbitrary acceptable formed allocation. If $\sum_i d_i \geq K$, choosing d weakly improves every investor. Otherwise raise all below-peak quantities to their peaks, then reduce above-peak quantities continuously until their sum is K . Raising toward a peak and reducing toward a peak both weakly improve the owner's payoff. The final profile satisfies the threshold equilibrium conditions. The acceptable domain is compact and convex because each f_i is concave. Its total objective is strictly concave, so its maximum, when the domain is nonempty, is unique. It is a Nash equilibrium by the preceding improvement argument. This proves (23).

Under pointwise deterioration, $f'_i(q_i) \geq 0$ implies $f_i(q_i) \geq 0$, so the new acceptable domain is a subset of the old. Its objective is weakly smaller at every common point. The equilibrium gap therefore contracts. All-zero gives $L = \sum_i O_i$, and individual acceptability bounds all other equilibrium welfare below by that sum. Adding its change proves continuous compression. The

proof neither assumes a lattice of contribution vectors nor extends to arbitrary quantity-dependent cost shares.

A.5 Reproduction and numerical boundaries

The threshold verification enumerates 78,732 exact rational comparisons, preserving all pure equilibria, entry deviations and ties. The new increasing-differences test checks 864 complete games, including failed subscriptions, and exactly matches the transition condition. The Bayesian checks enumerate 972 four-investor, two-type games under correlated full-support priors. One has no pure equilibrium, which is retained rather than excluded. Pure enumeration tests sufficiency and explicit failure cases; the mixed proof supplies the all-equilibrium result. An additional exact mixed witness has a unique all-sign pure equilibrium but only .75 subscription probability in a mixed equilibrium.

The continuous solutions use analytic concave optimization and direct deviation checks. A numerical grid is used only as a separate verification device. Saved allocations preserve individual payoffs, expenses, transfers, liquidity caps, information sets and proof status. Tests distinguish private reports, which are absent from the declared mechanism, from observable signatures. Contract-expiry and authority cases retain accepted obligations while blocking unauthorized new actions. Full reproduction also rescored the unchanged AI archive and independently recalculated its downstream allocations. No new inference response enters this edition.

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.
- Buchanan, James M. 1965. An Economic Theory of Clubs. *Economica* 32(125), 1–14.
- 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 Rohit Rahi. 1995. Financial Market Innovation and Security Design: An Introduction. *Journal of Economic Theory* 65(1), 1–42.
- Fang, Lily, Victoria Ivashina, and Josh Lerner. 2015. The Disintermediation of Financial Markets: Direct Investing in Private Equity. *Journal of Financial Economics* 116(1), 160–178.
- 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*.
- Halac, Marina, Ilan Kremer, and Eyal Winter. 2020. Raising Capital from Heterogeneous Investors. *American Economic Review* 110(3), 889–921.
- 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.
- Lerner, Josh, Jason Mao, Antoinette Schoar, and Nan R. Zhang. 2022. Investing Outside the Box: Evidence from Alternative Vehicles in Private Equity. *Journal of Financial Economics* 143(1), 359–380.
- 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.
- Milgrom, Paul, and Chris Shannon. 1994. Monotone Comparative Statics. *Econometrica* 62(1), 157–180.
- Milgrom, Paul, and John Roberts. 1990. Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities. *Econometrica* 58(6), 1255–1277.
- Morgan Stanley Wealth Management. 2023. Exchange Funds: An Important Alternative for Your Asset Allocation. August 2023, *CRC 5903090, especially pages 3–6 and 8*.
- Morris, Stephen, and Hyun Song Shin. 2003. Global Games: Theory and Applications. In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress, Volume I*, 56–114. Cambridge University Press.
- 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.
- Segal, Ilya. 2003. Coordination and Discrimination in Contracting with Externalities: Divide and Conquer? *Journal of Economic Theory* 113(2), 147–181.
- 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.

- Sákovics, József, and Jakub Steiner. 2012. Who Matters in Coordination Problems? *American Economic Review* 102(7), 3439–3461.
- Tabarrok, Alexander. 1998. The Private Provision of Public Goods via Dominant Assurance Contracts. *Public Choice* 96, 345–362.
- Topkis, Donald M. 1979. Equilibrium Points in Nonzero-Sum n -Person Submodular Games. *SIAM Journal on Control and Optimization* 17(6), 773–787.
- Vives, Xavier. 1990. Nash Equilibrium with Strategic Complementarities. *Journal of Mathematical Economics* 19(3), 305–321.

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 characterizations and the preserved benchmark. 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

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

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 $K_{w,c}^m$ 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

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

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.

VI. Preserved continuous benchmark and overlapping programs

The preserved benchmark has four investors with two capital units each. Investor i chooses continuous quantities $q_i, x_i \geq 0$ with $q_i + x_i \leq 2$, retaining the remainder as cash. The common program needs $Q = \sum_i q_i \geq 7$, gives gross value 15 per unit and costs 40, allocated in proportion to common quantities. Reservation costs $.5q_i$ whether or not it forms. Individual investments have gross per-unit values (18, 2, 3, 4) and common per-unit implementation cost c , compared at $c = 10$ and $c = 7$. Values are thousands of dollars; each capital unit is \$100,000. A separate \$20,000 expense account per investor is available under every route. Failed common capital is returned; reservation expense remains. This linear-value example is a benchmark, distinct from the concave interior-choice model in Section 6.

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

$$11(2 - q_1) + vq_1 = 22 - q_1(11 - v) < 22. \quad (27)$$

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

$$14.5Q - 40 + 11(8 - Q) \leq 14.5 \cdot 8 - 40 = 76. \quad (28)$$

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.

VII. Model access and analysis provenance

The recorded experiment used the product identifier `gpt-6.1-sol`, low reasoning effort, and Codex CLI 0.159.2 in read-only ephemeral sessions, accessed October 4–5, 2026. Exact UTC timestamps, prompts, rendered evidence, schema hashes, retrieved excerpts, returned usage and failures are preserved in the evaluation archive. The identifier is not an immutable model-weight snapshot. This edition replays archived records and performs deterministic financial extensions; it makes no new model inference calls for the evaluation. Frozen primary scores are preserved, and abstention-equivalence scoring is labeled as a post-hoc diagnostic. The synthetic renderer and scoring ground truth are executable, independent of the evaluated model’s outputs.