The Economics of Independent Mandates

Portfolio formation and the costs of separate implementation

Research draft · 4 October 2026

Abstract. Investors can share an investment without sharing a portfolio. We study formation costs when principals retain separate authority and may combine common and individual holdings. A four-principal economy admits a decline in its best Nash equilibrium and unique strong equilibrium as individual implementation becomes cheaper, although coordinated surplus remains unchanged. A unanimous compensation contract supports formation but leaves coordination failure possible. Funding reliability requires state-contingent resources; a recorded grant is not a guarantee. Numerical exercises and a bounded interpretation study distinguish implementation, verification and participation costs. The contribution is a model and counterexamples, not a claim of theoretical priority.

1 Introduction

The organization of investment determines which portfolios can be implemented. Several investors may want an economic claim that none can obtain economically alone. Their joint participation can cover the cost of forming a program, yet their reasons for participating need not coincide. One may be hedging an existing liability, another reducing concentration, and a third exercising a limited allocation right. Agreement on a particular investment need not imply agreement on the rest of their portfolios, their future capital commitments, or the authority they delegate.

This paper studies that distinction. Its unit of coordination is an implementation arrangement, rather than necessarily an entire portfolio. An independently governed principal may participate in several arrangements and hold an individually chosen completion portfolio around them. Each principal retains a feasible set determined by its wealth, obligations, mandate and authority. The coordinating institution chooses among compatible implementations, but it cannot turn an attractive outcome into a permitted action merely by computing its value.

The distinction has familiar precedents. Funds permit investors to combine common exposures with other holdings. Separately managed accounts can share execution. Intermediaries design securities, syndicates allocate contractual rights, and investment offices coordinate commitments.

Our question is narrower than whether these activities exist. We ask how a change in the cost of separate implementation affects the formation of an investment whose availability depends on several principals, and whether a permitted allocation of its costs can sustain participation.

The central example has four principals. Each can divide a two-unit capital budget among a shared program, an individual route and liquid holdings. The shared program requires more than six units, so every principal is needed, but every principal may also hold an individual investment. With a sufficiently expensive individual route, all-common participation is the unique strong equilibrium. A cheaper individual route then gives one principal a better outside use of capital. Its exit prevents the other three from forming the program. The unique low-cost Nash equilibrium has much less aggregate benefit even though the coordinated optimum is unchanged.

This statement requires care about equilibrium selection. At the higher cost there is also a no-formation Nash equilibrium. We therefore report the entire pure-equilibrium correspondence. The best Nash equilibrium deteriorates; the worst improves. Under the stronger requirement that no coalition can make all its members strictly better off by deviating, there is a unique outcome at each cost and its welfare falls. Strong equilibrium is a stringent diagnostic, not an estimate of how investors negotiate in practice.

The result survives continuously divisible allocations. It does not follow from excluding mixed portfolios or from a coarse numerical grid. Its mechanism is a formation externality: one principal compares the private benefit of leaving with its own participation payoff, while the investment opportunity lost by other principals does not enter that comparison. The argument is closely related to threshold provision and clubs. We establish a transparent financial representation and a reproducible counterexample; we do not establish priority for that general economic mechanism.

A voluntary compensation contract can preserve the shared investment. At the low cost, a transfer of $4,000 to the pivotal principal, funded by the other three, makes a fixed-participation unanimous contract acceptable and supports a strong equilibrium in its subscription game. A completion payment attached to the original divisible participation game does not have the same property. This payment rewards participation that keeps the other investors' opportunity set intact. It is not compensation for superior information or an assumption of investment alpha. The

unanimous contract leaves multiple nonformation Nash equilibria, however, and a bound on admissible transfers can prevent it from working. Cost allocation, equilibrium selection and execution authority remain different problems.

Funding uncertainty introduces a further distinction. In this paper, self-sufficiency means that the program can meet its funding requirement without an outside guarantor. A supported state with a shortfall in collectible internal funding prevents unconditional cash reliability. Correlation alone does not imply such a shortfall. Prefunding, redundant commitments and conditional closing can change the conclusion. The maximum shortfall is required backstop capacity under specified assumptions; its price depends on the states in which the guarantor must pay.

Artificial intelligence enters as a possible change in organizational costs. It may reduce the cost of interpreting mandates and maintaining evidence, or reduce the cost of separate individual implementation. These are different comparative statics. The formal example concerns the latter. A bounded language-model study measures scoped interpretation and validation tasks. It does not identify the total cost of administering an investment structure, nor demonstrate that AI causes the fragmentation in the model.

The argument has a practical implication that is more precise than a universal case for personalization. Institutions must identify which investor differences require separate implementation and which can survive within a common investment. They must then identify the participation terms and permissions that make the arrangement durable. Improving the interpretation of a mandate helps only to the extent that it changes these decisions at a cost below its benefit; it does not replace the decisions themselves.

2 Related research

Portfolio separation supplies a demanding benchmark. Heterogeneity in investors' preferences does not by itself require a distinct product for each investor. A sufficiently rich combination of existing funds may span the desired outcomes. Markowitz (1952), Merton (1971) and Ross (1978) provide different foundations for portfolio choice and separation. We retain that discipline by allowing liquid completions and mixtures, rather than treating a verbal difference between mandates as evidence that their portfolios must be implemented separately.

Financial innovation already makes securities and trading opportunities endogenous. Duffie and Jackson (1989), Duffie and Rahi (1995), and the review by Allen and Barbalau (2024) locate the design of claims within financial economics. Our use of a formation decision does not originate that question. The restricted object here is a participation arrangement whose availability depends on separately governed investors and whose formation can fail when their outside options change. We hold gross investment opportunities and beliefs fixed in the principal comparison.

Iancu and Trichakis (2014) directly study differentiated portfolios, shared trading costs and cost allocation. Their framework allows individual participation constraints and welfare objectives. We do not claim joint portfolio construction as new. Their Section 4.3 also permits extensions with mixed-integer allocation and nonseparable costs, including fixed fees. Our departure is the participation game and the effect of changing outside options on formation equilibrium, not fixed costs by themselves. Their work also makes a capable conventional manager an essential counterfactual: a new interface earns no economic advantage when an incumbent can implement the same arrangement at the same cost.

Recent distributed coordination work similarly prevents a false contrast between centralized optimization and independent objectives. Devanathan and coauthors (2026) coordinate transaction costs across portfolio decision makers. The computational location of an optimizer is not the definition of independent governance. A centralized system can respect separate authority, while distributed software can implement one centrally imposed policy. The economic questions are which actions are permitted and which joint constraints link them.

The closest mechanism in our principal example comes from clubs and threshold provision. Buchanan (1965) makes the sharing group an object of economic choice. Palfrey and Rosenthal (1984) study discrete public-good participation, and Bagnoli and Lipman (1989) connect a contribution-and-refund game to efficient outcomes under particular equilibrium refinements. Our shared program delivers excludable financial benefits rather than a pure public good. Nevertheless, the need to assemble a viable group and the failure of isolated best responses are established problems. The example should be read with those precedents in view.

Entry and product variety provide another useful comparison. Mankiw and Whinston (1986) identify opposing effects of entry, while Spence (1976) studies product selection with fixed

costs. The relevant lesson for our application is conditional: an additional individualized arrangement can remove valuable participation from an existing program, but it can also create an attractive program for other investors. A count of variants does not measure inefficiency. The relevant comparison is their net contribution to feasible investor outcomes.

Strategic delegation provides the background for our authority-cap example. Schelling's commitment logic and Fershtman and Judd (1987, corrected 2006) show why a principal may benefit from constraining an agent. Katz (1991) identifies conditions affecting whether unobserved contracts have commitment value. We use a small observable-cap example, not a general claim that visibility is necessary or sufficient. We also distinguish execution authority from the monitoring role of intermediation: permission to act and incentives to acquire information are separate contractual functions.

Cost sharing is a mechanism-design problem. Moulin and Shenker (2001), and Birmpas, Markakis and Schäfer (2019), study tradeoffs among incentive properties, budget balance and efficiency under specified cost structures. Formation thresholds need not satisfy the assumptions supporting their favorable guarantees. Our externality-charge calculation deliberately reports the resulting budget imbalance. Our compensation calculation instead asks a limited feasibility question under declared transfers and outside options. Neither is a universal allocation rule.

Verification is costly and selective. Ben-Porath, Dekel and Lipman (2014) study allocation with costly verification in a setting without monetary transfers. The result is relevant to deciding what must be checked, but does not solve investment allocation with cash payments and contingent obligations. Our records distinguish observations, permissions and accepted contracts. This separation identifies what an interpretation system is allowed to infer and what requires an independent source of authority.

Finally, correlated use of an algorithm can create effects that individual accuracy does not reveal. Kleinberg and Raghavan (2021) establish such a mechanism in selection. Our interpretation experiment reports errors within administratively assigned case groups that share one template; it does not identify cross-template dependence. It does not identify investor crowding, endogenous asset prices or systemic risk. These would require a model of how common errors enter demand and a separate empirical design.

The contribution claimed in this edition is consequently methodological and illustrative. A small economy exposes the relation between portfolio mixing, formation, exit and compensation. A funding calculation separates the size of a promise from its cost. A limited experiment separates a record index from a validation gate. The prior-art audit does not support presenting their combination as a newly established school of portfolio theory.

3 Environment and feasible arrangements

There are four principals, indexed by i, and dates 0, 1 and 2. At date 0 each principal has a financial book and a mandate. The book identifies existing assets, cash, liabilities and accepted unfunded obligations. The mandate identifies objectives and constraints. A separate grant specifies the contingent actions an agent may take. At date 1 investors choose participations and programs either form or fail. Cash flows, including obligations that survive an unsuccessful new proposal, are evaluated at date 2.

The general finite-state representation distinguishes an economic state from the information available when a decision is taken. An action rule must be measurable with respect to the date-1 information. A later receipt cannot be used to justify an earlier decision. The main participation example is a complete-information slice: all relevant payoff parameters are known when actions are chosen. The funding exercise separately introduces incomplete collection of promised cash. Neither exercise gives the coordinator a superior forecast.

Let W_i describe principal i's existing terminal wealth, including liabilities. A bundle b_i contains dated cash flows and specifies its rights and obligations. Its incremental certainty-equivalent benefit is the difference between the certainty equivalent with the bundle and the certainty equivalent of the existing book. Exposure labels may describe the bundle, but they do not substitute for its cash-flow and rights specification.

$$CE_i(W) = -T_i \log E[\exp(-W/T_i)]$$

$$\Delta_i(b_i) = CE_i(W_i + CF_i(b_i)) - CE_i(W_i)$$

The positive risk-tolerance parameter is T_i. Monetary gains in the main example are sure incremental amounts, so the distribution of existing wealth cancels under cash translation. This is a deliberate simplification that identifies the formation mechanism. A finite-state numerical

extension introduces correlated wealth and investment payoffs. It tests a local perturbation of the example, not a general stochastic-equilibrium theorem.

Financial feasibility and authority are distinct sets. Financial feasibility includes cash requirements, liabilities, exposure limits and admissible contractual rights. Authority identifies permitted actions at the current version and date. Execution requires their intersection and any counterparty conditions. A revoked grant can remove a future action without eliminating an already accepted contractual obligation. Conversely, an unexpired grant does not make an unaffordable investment feasible.

$$E_i(s) = F_i(s) \cap A_i(s) \cap D_i(s)$$

Here F_i is financial feasibility, A_i is authority and D_i collects documentary and counterparty conditions. The notation is a set of actions, not a scalar credit balance. Two maximum participation amounts calculated separately can draw on the same cash or authority limit. Their sum need not be feasible. An aggregate program must retain the allocation that supports it.

Programs have formation decisions. A prospective arrangement requires a minimum amount of participation and incurs a fixed cost when activated. Other costs depend on participations or failed attempts. These costs are resources consumed by implementation unless explicitly identified as transfers to a supplier. A supplier's markup reduces investor surplus but is not a resource cost for a welfare calculation that includes that supplier.

A common-policy pool is one admissible arrangement. A conventional custom vehicle can offer differentiated participations when its terms permit them. Separate accounts retain their ordinary shared-execution possibilities. The coordinated benchmark receives no special forecast, inventory or legal capability. It has an advantage only where a specified implementation difference changes the feasible set or cost. If an incumbent has the same feasible set and objective, the optimum is identical by construction.

The main example uses two units of participation capacity per principal. One unit may be embedded as $100,000 of investment principal. Benefit and cost parameters are expressed in thousands of dollars. A separate expense reserve is assumed sufficient for all enumerated canonical choices; it is not borrowed from another principal's protected cash. These quantities are

illustrative. They are not fitted to a particular asset class, contractual lot size or customer wealth distribution.

Each principal may select common participation q_i, individual participation x_i, and liquid residual. Quantities may be real-valued in the proposition. The numerical version first uses unit increments and then halves and quarters the increment. The constraint is a capacity constraint rather than a requirement to use all capital.

$$q_i \geq 0, x_i \geq 0, q_i + x_i \leq 2$$

Two principals may hold different amounts of the common investment and different individual completions. They may therefore be economically differentiated even when execution is shared. The later two-program exercise also permits overlapping program memberships. The fact that an optimum chooses a corner of this feasible set does not mean that mixing was prohibited.

4 The participation game

Let Q be aggregate common participation. The common program forms when Q reaches K. If it forms, a principal receives gross incremental benefit g per common unit and pays a proportional share of formation cost F. A reservation expense r per common unit is incurred whether the program forms or fails. Failed subscriptions return principal. The individual route gives principal i a gross incremental benefit b_i and costs c per unit.

$$Q = \sum_i q_i, a_i = b_i - c, \quad h = g - r$$

$$U_i(q, x) = a_i x_i - r q_i + 1_{Q \geq K}[g q_i - F q_i / Q]$$

The indicator term is zero when the program fails. The expression at Q=0 is defined as zero. It is a net incremental certainty equivalent under the baseline assumptions. There is no use of a marked-up terminal valuation to create a gain. The same gross investment benefits are available in every organizational comparison.

The strategic-form game fixes the cost-sharing rule before participation is chosen. Every principal takes the other principals' choices as given when considering a unilateral deviation. A Nash equilibrium admits no strictly profitable unilateral deviation. A strong equilibrium admits no deviation by any coalition that makes every member of that coalition strictly better off,

holding nonmembers' actions fixed. Transfers are absent from this game unless a specified compensation contract is added.

We do not equate strong equilibrium with coalition-proof equilibrium, the core or practical enforceability. Coalition-proof equilibrium imposes additional requirements on the self-enforcement of deviations. A transferable-utility core requires a characteristic function and permitted redistribution. Our strong-equilibrium test instead examines deviations in the stated participation game. This distinction is particularly important when a coalition can improve its aggregate benefit but cannot compensate a member who loses.

The principal example sets K=7, F=40, g=15 and r=0.5. Individual gross benefits are 18, 2, 3 and 4. At c=10 the first principal's individual net benefit is 8 per unit; the others' individual net benefits are negative. At c=7 the first principal's individual net benefit rises to 11. Common investment benefits, information and formation requirements do not change.

Table 1. Canonical economy; currency in thousands of dollars

Parameter	Value
Principals / capacity each	4 / 2 units
Formation threshold K	7 units
Real formation expense F	40
Common gross benefit g	15 per unit
Reservation expense r	0.5 per unit
Individual gross benefits	18, 2, 3, 4
Individual implementation cost	10 → 7 per unit
Outside / common low-cost CE	(22,0,0,0) / (19,19,19,19)

Under all-common participation, Q=8. Each investor receives 30 of gross common benefit, pays 10 of formation cost and 1 of reservation expense, and obtains 19. Aggregate benefit is 76. The relevant outside outcome gives the first principal two units of its individual investment and leaves the other principals liquid. This outside outcome is jointly attainable: no shared scarce capacity has been promised separately to multiple investors.

At the higher individual cost, the outside outcome gives benefits (16,0,0,0). At the lower cost it gives (22,0,0,0). The first principal can also hold one common and one individual unit. Such a mixed position remains in the action set at both costs. At the low cost, however, joining a seven-unit common program with one unit gives it less than the 22 available from investing individually. Its departure is not an artifact of requiring whole-portfolio delegation.

Formation threshold and fixed formation cost play different roles. K determines whether the remaining principals can constitute a program after an exit. F determines the net benefit from a program that forms and the burden of a proportional cost share. Reservation expense rules out costless failed subscriptions as strategically equivalent actions. None of these parameters is a numerical approximation to the other two.

5 A signed comparative static

Proposition 1 considers quantities on the continuous simplex, not merely the enumerated grid. Define v as common net benefit per unit when all eight units participate.

$$v = h - F/8$$

Suppose 6<K<=8, r>0 and v>0. In a high-cost regime let the first principal have individual benefit $0<a_1<v$ and let the other individual benefits be negative. In a low-cost regime let $v<a_1<h$, retain negative individual benefits for the other principals, and require $2a_1<8h-F$. Then all-common participation is the unique strong equilibrium in the high-cost regime. In the low-cost regime, no common program forms in the unique Nash and strong equilibrium. The coordinated maximum is all-common in both regimes.

The proof uses the common program's maximum scale. In any formed profile Q cannot exceed eight, so a common unit's net benefit cannot exceed v. At the high cost, an individual unit is worth less than v to every principal. Each principal's benefit is therefore at most 2v, and attaining that value requires full common participation by all four. Every other profile gives every principal strictly less than 2v. The all-common profile is a Nash equilibrium and blocks every other profile through a deviation by the grand coalition. It admits no improving coalition deviation itself.

The low-cost argument runs in the other direction for the first principal. For any positive common participation, its payoff is bounded above by the payoff from investing both units individually, less a strictly positive amount. This applies to formed programs because a common unit is worth at most v<a_1, and to failed programs because reservation expense is positive.

$$U_1(q,x) \le 2a_1 - q_1(a_1 - v) < 2a_1 \; when \; q_1 > 0 \;\wedge\; Q \ge K$$

Thus the first principal strictly prefers its individual route for all capital. The remaining three can supply at most six common units, below K. Each then avoids a failed reservation and retains liquidity rather than buying a negative-benefit individual investment. This identifies the unique low-cost Nash equilibrium. It is also strong: any formed coalition outcome must include the first principal, who would be worse off, and any unformed deviation cannot improve every member.

To establish the coordinated maximum, bound the total benefit of a formed allocation. Giving every unit outside the common program the first principal's individual payoff can only overstate its contribution. Since h exceeds a_1, the resulting bound rises with Q and is maximized at eight. The stated inequality makes that outcome better than the best unformed allocation.

$$\sum_i U_i \le \; hQ \; - \; F \; + a_1(8 - Q) \; \le \; 8h - F$$

This proves the comparative static without selection among strong equilibria. In the numerical example, c=10 and c=7 lie strictly inside the relevant regions. The threshold can vary anywhere above six and at most eight. The result therefore holds on a nontrivial parameter region. It is not a knife-edge equality in a computational search.

There is nevertheless a significant qualification about ordinary Nash equilibria. At the high cost, an unformed profile is also Nash. No principal can form a seven-unit program alone, and a unilateral failed reservation is costly. The high-cost equilibrium benefits are therefore 16 and 76 in aggregate. At the low cost the only Nash benefit is 22. Best-equilibrium welfare falls by 54, while worst-equilibrium welfare rises by 6. A claim that all decentralized outcomes deteriorate would be false.

The example also does not show that more detailed mandates generally destroy value. It varies the cost of a particular individual route while retaining the common program's costs and rules. If interpretation technology instead reduces common coordination costs, the outcome can change in the opposite direction. If customized holdings preserve the same shared implementation, there

may be no lost scale. The point is to identify the margin along which costs fall, not to attach a negative sign to personalization as a category.

The proposition is an application of a familiar threshold externality. It is economically useful because it survives a capable portfolio-mixing benchmark and makes the outside-option change explicit. A claim of broader novelty would require either a new characterization for richer financial environments or an empirical result showing that these constraints materially govern investment formation. The current evidence establishes neither.

6 Allocation of the formation surplus

The coordinated maximum of 76 is a potential aggregate benefit before participation transfers. It is not an executable recommendation to impose 19 on the first principal when that principal can obtain 22 elsewhere. A coordination arrangement must account for the new outside option. The cost of retaining participation is therefore part of program design.

Consider a contract specified before choices are made. If all four principals supply two common units, the first receives 4 and each other principal pays 4/3. In every other profile the transfers are zero. The resulting formation benefits are 23 for the first principal and approximately 17.667 for each other principal. All exceed the low-cost outside benefits. The transfers sum to zero; they do not create resources.

Table 2. Fixed-allocation participation accounting ($000)

Principal	Outside	Common	Transfer	After
A	22.000	19.000	4.000	23.000
B	0.000	19.000	-1.333	17.667
C	0.000	19.000	-1.333	17.667
D	0.000	19.000	-1.333	17.667

Internal review exposed a limitation in the original numerical compensation rule. That rule paid the transfer only when every principal selected two common units, while a seven-unit program could still form without the transfer. A donor could reduce its participation from 2 to 1.9375, cancel its payment, and still obtain a formed program. Its benefit rises from $17,666.67 to

$18,329.97, an improvement of $663.30. Thus the coarse-grid compensation equilibrium is not a continuous Nash equilibrium. This failure exceeds the numerical stability threshold and is retained in the results.

A separate, unanimous subscription contract remedies this particular defect by changing the permitted use of subscriptions. It specifies four two-unit participations and the stated transfers as one all-or-none offer. The physical formation minimum remains seven; this contract requires eight. No participant's subscription can be used in a partial reoffer without fresh acceptance. Rejection leaves the ordinary outside outcome available, with no incremental reservation expense for this preliminary offer. The original r>0 participation game and this costless-acceptance subscription stage are different mechanisms. Acceptance of the complete contract yields (23,17.667,17.667,17.667) in thousands, strictly above each principal's outside payoff. All-accept is the unique strong outcome of the binary subscription game. Enumeration finds twelve Nash profiles, eleven without formation. This is not unique implementation or a general bargaining solution. Ordinary seven-unit routes remain possible under separate subscriptions; they cannot appropriate capital conditional on this offer.

The first principal's payment exceeds the minimum necessary to match its outside option by 1. The other principals finance that excess because their participation is valuable only if a program exists. This particular split is a declared policy, not a fairness theorem. A Nash bargaining solution, equal division of the residual surplus, or another admissible rule would select a different split. The economic point is that a principal's contribution includes its effect on formation, not merely the amount of capital it supplies.

The original unit-grid transfer calculation is retained as a negative numerical finding: its apparent stability fails under finer continuous deviations. The replacement subscription game is checked separately over all sixteen acceptance profiles. Its all-accept outcome strictly Pareto dominates its nonformation outcomes, so it is strong. Multiple nonformation Nash profiles persist because one investor cannot alone complete the group. No claim is made that this contract fully implements the preferred outcome.

This distinction limits the claim that an externality price fixes the problem. A transfer contingent on completion addresses the distribution of surplus. It does not necessarily solve simultaneous participation, credibility of promised payments, interpretation of contract conditions or the

ability to enforce the arrangement. An institution may need a subscription procedure with conditional acceptance and a refund rule. Those changes define another game and require their own analysis.

For a fixed allocation, transfer feasibility has a simple characterization. Let g_i be gross certainty-equivalent benefits before transfers, o_i be jointly feasible outside benefits, and L_i be symmetric bounds on payments. A transfer must exceed the participant's shortfall, respect its bound and balance in aggregate. Define the effective lower bound ell_i.

$$\ell_i = \max(o_i - g_i, -L_i)$$

$$\ell_i \le t_i \le L_i, \sum_i t_i = 0$$

A feasible transfer vector exists if each lower bound is at most its upper bound and zero lies between the sum of the lower bounds and the sum of the upper bounds. This is an interval-feasibility result. It supplies neither truthful preference revelation nor stability against every coalition. If the three donors may each pay at most 0.5, their total capacity of 1.5 cannot cover the recipient's loss of 3. The aggregate gain of 54 then fails to establish implementability.

Nor can everyone necessarily preserve the payoff obtained before technology changed. Keeping the other three principals at 19 while giving the first its new outside payoff of 22 requires 79. Only 76 is available. The contract produces a Pareto improvement over the low-cost outside outcome, but some members receive less than in the former all-common arrangement. Reporting only the aggregate benefit would conceal this distributive change.

We also calculate a marginal-contribution rule. Define a principal's mechanism payoff as total investment welfare with its action less total welfare when its action is replaced by zero, holding others fixed. Unilateral changes in this payoff equal unilateral changes in total welfare, so total welfare is an exact potential. This identity makes the planner's optimum a Nash equilibrium of the modified game, but does not establish budget balance or eliminate local no-formation equilibria.

In the canonical all-common profile, removing one principal while holding the remaining failed reservations fixed leaves negative reservation costs. The resulting marginal-contribution payoffs sum to more than the investment benefit. Supporting them requires an outside subsidy of 240 in the calculation. This is different from the losses after remaining principals reoptimize: 57 when

the pivotal principal exits, and 35 when another principal exits and the first uses its individual route. The two externalities use different counterfactuals. Neither can be called the unique price of a preserved difference.

These calculations reject a general rule that each preference should carry its marginal cost to a coalition. Preferences are not invariant units: one restriction can be expressed in several sentences, and several restrictions can describe the same feasible set. Payments should attach to verifiable economic consequences of a specified arrangement. Even then, budget balance and authority must be checked. A computed externality does not authorize collecting a charge.

7 Funding reliability and the value of authority

A willingness to participate, a revocable grant and an accepted obligation expose the program to different funding risks. The relevant quantity at an execution date is cash that can be delivered under the contract and state of the world. A nominal commitment does not establish that quantity. A principal may be willing but illiquid, authorized but missing documents, or contractually obligated while in default.

Let K(s,a) be the program's funding requirement in state s under an admissible participation or withdrawal profile a. Let C(s,a) be internal funding that is both available and collectible at the required time. The worst shortfall defines a lower bound on a fully reliable external cash backstop.

$$G = max_{s,a} max[0,\ K(s,a) - C(s,a)]$$

The bound is necessary for unconditional cash coverage. It is sufficient for that limited purpose when the guarantor can deliver at the required time in every supported state and no other condition prevents closing. It does not establish issuer consent, eligibility or settlement. Collateral requirements and guarantor default could require more resources or make the proposed assurance impossible.

Positive correlation among funding shocks can increase the probability or severity of joint shortages. It does not by itself prove a positive G. A prefunded program can remain reliable under perfectly correlated shocks to investors' outside wealth. Conversely, an idiosyncratic shortage can prevent a program with no substitute participant from closing. A program allowed

to cancel or reduce its size also has a different reliability requirement from one required to form at a fixed scale in every state.

The illustrative funding schedule requires 100 at one date. Revocable participation provides 100 in the normal state, 70 after one withdrawal and 40 in joint stress. The corresponding binding-commitment schedule provides 100, 95 and 80. Binding obligations improve the assumed collection profile but do not make it certain. These figures are declared scenarios, not estimates of contract enforcement.

Table 3. Illustrative cash backstop, dollars

Support	Maximum gap	Expected payout	State-price value
Revocable	60,000	12,000	13,800
Binding	20,000	3,000	3,500

The physical state probabilities are 0.7, 0.2 and 0.1. A separate illustrative state-price vector is 0.66, 0.22 and 0.12. The vectors deliberately differ. Expected payout uses physical probabilities; valuation of the contingent claim uses the selected state prices. Neither calculation provides a market quote for a guarantee. A price would also depend on the provider's risk bearing, funding, administration and contractual remedies.

Prefunding moves the burden to investors. In the simple exercise, reserving 100 for 90 days at a forgone annual return of 4 percent costs approximately 0.986. The cash shortfall disappears under the assumed custody and settlement conditions, but investor flexibility has been used. Comparing this cost with a guarantee's price requires the same period, contingencies and treatment of default. Reporting only the nominal guarantee amount would obscure that comparison.

Authority has a different economic effect when suppliers adjust their offers. Suppose a buyer values a claim at 15 and one supplier can provide it at cost 9. Prices are restricted to the integer grid from 9 to 15, the buyer accepts at indifference, and the supplier makes a take-it-or-leave-it offer. Without a cap, the supplier chooses 15. A publicly observable, irrevocable cap of 12 changes the supplier's best offer to 12. The buyer gains 3 and the supplier retains 3.

The cap narrows the buyer's action set at fixed prices. Its benefit comes entirely from the induced change in the supplier's offer. If the cap can be freely amended once a higher price is proposed, that commitment effect disappears in the specified comparison. A cap below the supplier's cost prevents trade. The example illustrates an established strategic-delegation mechanism; it does not imply that narrower authority is generally better.

Verification creates an information question. A supplier may need evidence that a cap is binding without needing the investor's full wealth, holdings or maximum willingness to pay. Selective disclosure could separate those facts, but a certificate's credibility depends on who can amend the grant and what the counterparty can enforce. The current model assumes the relevant commitment technology. It does not claim that a database entry supplies it.

8 Numerical verification and organizational regions

The base grid gives each principal six possible common-individual bundles, including a mixed bundle and a liquid outside position. We evaluate all 1,296 joint profiles. Unilateral best responses identify every pure Nash equilibrium. Every coalition of two, three and four principals is then allowed to change all of its members' actions, holding outsiders fixed. A profile survives as strong only when no such deviation strictly improves every deviator's payoff.

A separate scalar evaluator reconstructs each profile's cash-flow benefit and certainty equivalent. This check does not rely on the vectorized equilibrium search. The comparison includes formation costs, reservation expenses and transfers. Supplier rents are recorded separately from resources consumed. Reproduction uses a pinned numerical environment and deterministic inputs; runtime variation is not an economic result.

We then refine the allocation increment to one half and one quarter of a unit, giving 50,625 and 4,100,625 profiles respectively. The headline high- and low-cost outcomes do not move. The analytic continuous-quantity proof is stronger than this numerical stability result, but the refinement checks that the implementation has not introduced an unintended constraint. A contractual indivisibility, by contrast, would remain fixed while the numerical grid changes.

Table 4. Finite-grid outcomes; aggregate CE in dollars

Case	Best Nash	Worst Nash	Coordinated
Higher individual cost	76,000	16,000	76,000
Lower individual cost	22,000	22,000	76,000
No formation frictions	120,000	120,000	120,000
Smaller threshold	69,000	22,000	76,000
Common payoff risk; high cost	75,877	15,972	75,877
Common payoff risk; low cost	21,972	21,972	75,877
Expensive wrapper	16,000	16,000	16,000
Cheaper shared implementation	104,000	22,000	104,000
Supplier markup of $54,000	22,000	22,000	22,000

The finite-state risk extension gives each investor a different loading on the same three-state background risk and adds risk to both investment routes. Each common unit has the same state-contingent payoff for every principal; differing background books and tolerances create differing valuations. It retains the qualitative equilibrium pattern. The high-cost all-common aggregate certainty-equivalent gain becomes approximately 75.877; the low-cost separate outcome becomes approximately 21.972. These are local sensitivity findings. They do not extend the proposition to arbitrary risk aversion, payoff distributions or liabilities.

The no-friction case has no formation minimum, fixed expense or reservation expense. The ordinary Nash outcome and coordinated maximum coincide. A high wrapper cost eliminates the shared program's advantage. Reducing the common implementation cost can sustain formation even when the individual route is cheap. These cases identify the source of the result and reject a generic advantage of coordination.

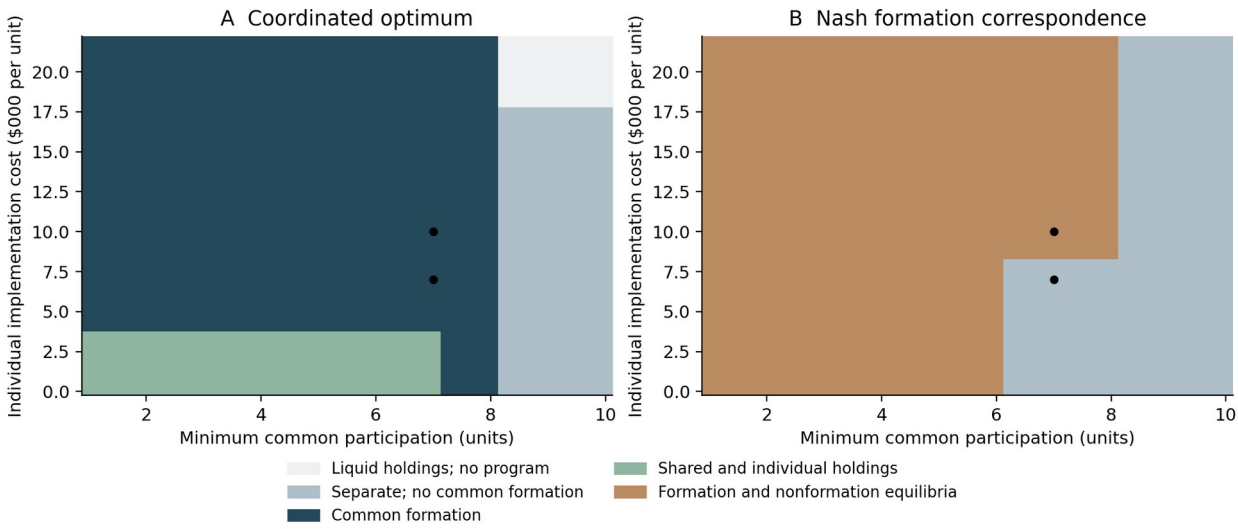


Figure 1. Organizational regions in the finite economy. The horizontal axis is the common program's formation threshold; the vertical axis is the cost per unit of the individual route. The coordinated allocation and equilibrium correspondence are separate panels. Regions depend on the stated payoff and cost parameters. Multiplicity is displayed rather than resolved by an assumed selection rule. Editable vector versions and underlying cell data accompany the paper.

An additional two-program exercise illustrates the opposing effect of differentiation. The first program alone produces aggregate benefit 32. Allowing a second program raises the coordinated maximum to 50 while incurring an additional formation cost. Different principals favor the two programs, and all mixed bundles are available. The example does not measure the social value of variety in a market; it establishes that formation externalities need not always favor a single common arrangement.

The distinction between investor surplus and resource surplus is especially important when supplier pricing changes. With a real formation cost of 40 and a supplier markup of 54, all-common investment still produces 76 of combined participant-and-supplier benefit. Only 22 remains with investors. Whether they can divide that amount to meet their outside options is a separate participation question. Treating the markup as a real formation expense would incorrectly describe a transfer as destroyed surplus.

A conventional custom vehicle with the same opportunities, costs, permissions and cost-sharing flexibility reproduces the coordinated allocation. The numerical package makes no claim that a named platform has a superior solution to this identical problem. Its potential commercial

contribution would have to come from lower measured implementation costs, access to an otherwise missing route, better operational performance or a different admissible contract. None is established by the synthetic comparison.

9 Applications and limits of evidence anchoring

The four application exercises expose different meanings of an investor difference. They do not calibrate the entire four-principal economy. Published data inform specific ranges or contractual conditions (Harris, Jenkinson and Kaplan, 2012; Jefferies, 2025; Sosner, Gromis and Krasner, 2022), while the joint payoff laws, preferences and formation costs remain declared assumptions. A parameter ledger records that boundary.

For private funds, commitments are dated obligations rather than invested balances. The application has calls at months 6, 18, 36 and 126 and distributions at months 48, 84 and 120. A stress accelerates calls and delays distributions. The calculation retains obligations and proceeds beyond month 120; it does not sell an illiquid position at an invented terminal price. Work on private-fund performance, commitment pacing and commitment risk motivates the cash-flow representation, but publicly reported aggregate statistics do not identify these illustrative paths.

For secondary interests, purchase price and unfunded obligations must be considered together. We examine prices of 85, 89 and 95 per 100 of reported NAV and separately vary assumed remaining calls. Jefferies reports an average LP price of 89 percent of NAV for 2024. That is a transaction-market statistic, not a fair-value estimate for the synthetic claim. The payoff scenarios are independently declared. A discount can compensate for poor underlying performance, obligations, illiquidity or stale marks rather than create an arbitrage.

A bilateral risk-transfer calculation makes a different source of value explicit. Both parties use probabilities of one quarter, one half and one quarter for a claim paying 0, 100 and 200. The seller holds four such claims plus cash and has a lower risk tolerance than the buyer. Selling one claim creates a price interval in which both parties gain. At a price of 90, their certainty-equivalent gains are approximately 23.613 and 7.501. The total is about 31.114, independent of the cash price under the stated CARA assumptions. These are gains from risk redistribution under identical beliefs, before costs and taxes.

For public securities, a tax difference can require separate tax-lot accounting while allowing shared execution. The tax exercise values a current realized loss and an offsetting increase in taxable gains at eventual liquidation. It applies the same tax rate at both dates and subtracts incremental fees. The benefit is the time value of deferral under these assumptions. No permanent tax saving, donation strategy or basis step-up is assumed. Published tax-management work supplies a reason to model realization and fees; it does not supply a universal additional expected return.

For credit, contractual minimums and exceptions are part of the opportunity. A filed Mach credit agreement defines a $500,000 borrowing minimum, with a stated remaining-commitment exception, and increments of $100,000 under the cited provision. A separate provision treats specified reimbursement loans differently. The application checks the arithmetic of those terms and leaves other conditions unresolved. It does not infer an available credit line, eligibility to borrow, or a transferable right for a hypothetical investor.

These applications suggest a useful classification. Some differences alter desired exposure but leave common implementation intact. Some change personal liabilities or tax consequences while preserving the same underlying claim. Others change the admissible legal claim, timing or counterparty and may require another arrangement. A catalogue organized only by thematic exposure cannot distinguish them. The classification must be made from the relevant cash flows, constraints and rights, with unknowns retained.

No application estimates demand for a commercial service. A model-implied gain is not willingness to pay, and a guarantee's illustrative state-price value is not a supplier quote. Real-book evidence would require historical information sets, declined opportunities, actual invoices and contractual permissions. Selected successful transactions would not identify the formation failure that motivates the model.

10 Interpretation and validation study

The interpretation study asks whether an evidence index and a deterministic gate improve a narrowly specified decision task. It does not ask whether a language model can manage an investment office. The design has two factors: an evidence index of the supplied paragraphs, and

post-response deterministic validation. Every arm receives the same source paragraphs and questions. External search and tool use are prohibited during a run.

The index preserves source text and labels rather than supplying inferred financial facts. Its construction is deterministic and its cost is recorded. This is a deliberately limited records treatment. It tests organization of evidence, not the accuracy of a learned extraction system or the economics of building a registry from arbitrary investor documents. Because the index repeats evidence, it can increase input tokens even under the same resource budget.

The validation factor is downstream of the model response. It checks supplied financial and permission inputs under fixed rules and blocks an action when a supported condition fails. It does not repair every factual answer. In the public-document slice, the gate prevents a filing from being treated as a principal's execution grant. The validator never reads the oracle answer field. Its regular-expression parser is nevertheless vulnerable to grant-like text in an unauthorized attachment, as a post-review probe demonstrates. It is a fixture checker, not a provenance-secure authorization boundary. All effects of the gate are therefore reported separately from raw interpretation accuracy.

There are 60 synthetic dossiers, of which 12 are development cases and 48 are held out. Held-out pairs change one relevant condition, including an expired grant, a competing reservation, research-only scope, conflicting book evidence, payment permission or a future amendment. The split holds out numerical and mutation instances, not templates. A future amendment does not change current authority before its effective date. Source attachments also contain an instruction to ignore permissions, which must be treated as untrusted document content.

The held-out synthetic design comprises 192 first evaluations and 64 predetermined repetitions. The public slice adds 48 evaluations of 12 scoped cases. Three cases concern credit, three tender documents, three indentures and three interval-fund documents. There are 11 unique filings because two credit cases examine different provisions of the same agreement after another source proved unavailable. Errors in those cases are clustered by filing. Labels are agent-audited and exploratory, not human-validated.

The public cases use saved source spans with retrieval dates and hashes. They test stated terms and distinctions, such as a fund-level repurchase percentage versus an individual liquidity guarantee, or a tender price for common shares versus a different class. They do not constitute

complete-document legal interpretation. Missing documents and personal grants are not supplied by inference. Later information is not used to answer an earlier contractual question.

Table 5. Primary calls; strict / unit-normalized correct answers

Slice / index / guard	Completed	Strict	Unit-normalized
Synthetic / 0 / 0	47/48	184/192	184/192
Synthetic / 0 / 1	46/48	181/192	181/192
Synthetic / 1 / 0	47/48	174/192	188/192
Synthetic / 1 / 1	47/48	170/192	188/192
Public / 0 / 0	12/12	47/48	47/48
Public / 0 / 1	12/12	48/48	48/48
Public / 1 / 0	12/12	47/48	47/48
Public / 1 / 1	12/12	48/48	48/48

The 304 evaluation calls produced 299 completed responses and five timeouts. None of the completed responses proposed an action classified as unauthorized by its fixture. The deterministic guard therefore intercepted no unsafe proposal. This is a null finding for incremental validation benefit in this sample, not a certification of safety. The public slice contained no executable case; its action question tests abstention only.

Among the 192 primary synthetic calls, strict answer scores were 184, 181, 174 and 170 out of 192 questions in the four arms. The order is no index/no validation, no index/validation, index/no validation, and index/validation. Failures contribute no correct answers in these totals. Most apparent factual errors in the indexed arms were numeric answers with a USD suffix where the frozen exact-string scorer expected a bare number. A separately reported, post-study normalization removes that suffix. It must not replace the frozen strict score. This distinction prevents a scoring-format effect from being described as deterioration in financial reasoning.

Seven completed primary synthetic calls chose review although their fixture labels permitted simulated execution. All seven cited the prompt's statement that no live action was authorized. The question asked whether an action could proceed now without cleanly separating simulated permission from actual execution. These are simulation-scope disagreements, not established

missed investment opportunities. All seven occurred without the index, but that difference cannot be interpreted as a demonstrated economic benefit of the index.

Primary public accuracy was 47, 48, 47 and 48 of 48 answers. The two disagreements concerned whether quoted prices alone establish equivalent economic rights across share classes. Internal review found that item ambiguous. The frozen label is retained, with an exclusion sensitivity: every arm answered the other 47 items correctly. Validation changed no response. Its model prompts were identical to the corresponding no-validation prompts; differences in raw scores between those arms are repeated-call variation, not effects of the guard.

The 64 repeated calls all completed, and 58 returned exactly the same normalized answer dictionary as their first attempt. This is descriptive repeatability within the selected cases. There are only twelve administrative case groups generated from one common synthetic template and eleven distinct public filings; case counts should not be mistaken for that many independent sources of evidence. The evaluation provides no estimate of market-wide correlated portfolio behavior.

Median latency among completed calls was 15.4 seconds. Calls ran concurrently in a fixed order under uncontrolled server load. The five timeouts cluster in the late private-fund case group; they cannot be assigned to a treatment effect. Completed-call latency excludes their waiting time. Completed calls reported 4,486,713 input tokens and 49,903 output tokens, including substantial execution-environment overhead. A token total is not a measured dollar cost. The price of professional review and losses from missed obligations remain assumptions. Because the observed abstentions are confounded by simulation-scope language, we do not price them as missed investment gains. The post-study loss formula is retained only as an unestimated specification. A frontier based on these calls cannot identify a universal price of preserving an investor difference.

The model is pinned to the requested gpt-6.1-sol identifier with low reasoning effort, and the client invocation is saved. This identifies the serving configuration, not immutable underlying weights. Token use, elapsed time, returned answers and technical failures are retained. A timeout is not counted as a correct answer or silently replaced. Completed-run accuracy and end-to-end completion are different reported outcomes.

The treatment contrast is descriptive at this sample size. Dossiers share templates; four responses to one case are paired observations, not four independent investor books. The predetermined repeats reveal some serving variability, but do not make the benchmark representative. An apparently perfect score on simple evidence may indicate a ceiling effect. It would not establish reliability on incomplete statements, complex amendments or live authority disputes.

The economic link is consequently partial. Measured computation and review flags inform the cost of a scoped interpretation task. Human review time is not observed; any conversion to dollars uses a declared range. A model-implied penalty for an incorrect action depends on its economic consequences and cannot be estimated by counting mistakes alone. No single scalar cost of preserving a distinction is inferred from this study.

Existing work on financial-agent tool integration provides a further reason to avoid an architectural presumption. Brar and coauthors (2026) report a factorial experiment with results unfavorable to their tool treatment, while noting substantial limitations of their synthetic evaluation. Our study has a different intervention and outcome. It treats improvements, null effects and technical failures as admissible findings rather than evidence to be selected for a product thesis.

11 Authority and the implementation boundary

The model's allocations are proposals until corresponding actions are permitted and accepted. A practical representation must retain the distinction between a mandate, a grant and an obligation. A mandate describes desired outcomes and constraints. A grant permits a specified agent to act. An obligation records what a principal has already promised under a contract. An amendment to one does not automatically amend the other two.

The book must identify actual positions and obligations with their source, economic effective date and date of knowledge. A corrected observation changes what the system believes about the book; it does not retroactively authorize a past transaction. Conflicting sources require a recorded unresolved condition. A source that explains an investment's terms cannot instruct the system to expand a grant.

Program construction needs both a feasible allocation and a participation explanation. It should state each principal's ordinary alternatives, costs, remaining financial capacity and current

authority. The first principal in the canonical low-cost allocation should not receive a bare recommendation to invest commonly. It should receive the participation terms that make doing so preferable to its individual route. If those terms cannot be offered, the arrangement fails at the participation boundary.

A parent program cannot advertise a fungible amount of capital unless an underlying allocation supports it. Two offices representing the same principal do not create two budgets. Two mandates attributing the same position do not create two assets. Capacity is a set of possible actions subject to shared constraints. A scalar amount is a conditional projection that must preserve its assumptions.

Execution then requires version checks, current authority, reservation of capital and expenses, and counterparty acceptance. Funding evidence, settlement evidence and ownership evidence answer different questions. A revoked future grant blocks a new act requiring that grant while leaving existing obligations on the book. This distinction matters particularly for conditional commitments, where the conditions governing acceptance may be fixed at an earlier mandate version.

These requirements can be represented in software, but their legal force is not supplied by the representation. A stored price cap is a bargaining commitment only if amendment and enforcement rules make it credible. An escrow receipt supports a claim about funds only to the extent that the provider and terms are authoritative. A beneficial interest in a vehicle is not necessarily a transferable issuer-level share. The research model assumes explicit rights rather than manufacturing them from labels.

The systems implication is that interpretation should propose and explain, while validated policy governs actions. This does not mean a deterministic validator possesses complete legal knowledge. Its jurisdiction is limited to specified predicates supported by evidence. An unknown right remains unknown. The appropriate response may be to investigate, seek an authorized amendment, change the structure or abstain.

A production evaluation would also need to distinguish a negative action decision from uncertainty about an answer. An expired grant can be conclusively inapplicable; a missing grant may leave authority unknown. Both prevent execution, but they imply different review work.

The current binary action score does not measure that difference. Nor does it validate whether a supporting quotation is complete enough to establish the proposed right.

The economic object ultimately required is a distribution of workflow costs conditional on case difficulty and consequence. Review time, model usage, correction time and residual decision loss belong in separate columns. They can be combined only under declared prices and a loss model. A numerical accuracy gain on short paragraphs cannot be inserted directly as a reduction in vehicle formation expense. The financial example lowers the cost of individual implementation; the AI pilot varies evidence presentation. The connection between those two interventions remains a hypothesis.

This separation also clarifies the organizational prediction. Technology that reduces shared coordination expense can preserve formation in the numerical economy even when the individual route is inexpensive. Technology that reduces only the individual route's expense can remove a pivotal participant. A technology that improves both margins requires a joint comparative static. A general claim about artificial intelligence has no determined sign until its effects on these different costs, participation incentives and supplier responses are specified.

12 Conclusions

The example establishes a specific possibility: lowering the cost of an individual investment route can destroy a valuable shared program even when mixed portfolios are feasible. The best Nash equilibrium and the unique strong equilibrium deteriorate, while the worst Nash equilibrium improves. A coordinated arrangement retains its aggregate potential, but must redistribute some benefit to preserve participation. A unanimous fixed-participation compensation contract supports formation without eliminating every coordination failure.

The distinction between costs is central. Cheaper interpretation can make independent mandates easier to administer. Cheaper separate implementation can make exit more attractive. Cheaper common coordination can make a shared program more durable. Calling all three personalization conceals the comparative static that determines the result. The relevant empirical task is to measure the cost and constraint that actually changes an investment decision.

The paper also establishes limits. Correlated funding shocks do not by themselves imply an impossibility theorem. The maximum cash gap is not the guarantee's price. A narrower grant can

improve a negotiated outcome only through the specified strategic response. A marginal-contribution rule need not balance its budget, and a participation payment is not proof of truthful reporting or general coalition stability. The finite interpretation study cannot validate autonomous investment management.

The broader proposition is an organizational one. Independently governed portfolios can share implementation at the level of particular claims and programs. Their opportunity sets may depend on other investors' participation, contracts and information. Studying that dependence can extend the questions asked in portfolio construction without abandoning its existing discipline. The investment frontier is partly designed, but the design must survive outside options, authority and resource accounting.

Appendix A Proof details and boundary cases

Cash translation follows directly by factoring the sure transfer out of the exponential expectation. The logarithm separates the factor, leaving the certainty equivalent of the original wealth plus the transfer. This proves a monetary translation property, not linear utility of uncertain money. Transferring a state-contingent claim changes risk and does not obey the same identity unless the relevant additional structure is specified.

Fixed-set monotonicity is a maximization argument. Let X be unchanged and let V(x,c) weakly rise for every x when c falls. The previous optimizer remains feasible and is worth at least as much under the new cost. The new maximum cannot be lower. If participation constraints are rewritten using improved outside options, X may change. The conclusion then requires a separate feasibility argument. The low-cost canonical allocation uses transfers to meet its new outside options.

For the funding bound, any state and action profile with a gap larger than a proposed guarantee defeats cash reliability. Taking the maximum gives necessity. For sufficiency, let a guarantor deliver the nonnegative gap at the required time in every state and profile; internal funding plus that delivery meets the requirement. The proof says nothing about other closing conditions. If the guarantor's own default is a supported state, its collectible resources must be modeled rather than presumed.

For the high-cost strong-equilibrium claim, consider first a formed profile. Since Q<=8, every common unit has net benefit at most v. Individual units have strictly smaller benefit, and liquid residual has zero. Thus each principal receives at most 2v. Equality requires two common units and Q=8. Every different formed profile has Q<8 or a principal using less than two common units, making each principal's bound strict through either quantity or the common cost share. Unformed profiles also give every principal strictly less than 2v. The grand coalition can therefore strictly improve every member by selecting all-common, while no coalition can improve every member from all-common.

For the low-cost claim, the first principal's best response is uniquely q_1=0 and x_1=2. In a formed profile any common unit displaces a strictly better individual unit, and in an unformed profile it incurs r without a common payoff. Given that action, the remaining principals cannot meet K. Their negative individual benefits and positive failed-reservation costs make zero participation their unique best responses. The resulting profile is strong because any formed deviation requires the first principal, who cannot improve, and no unformed deviation improves all deviators.

The coordinated bound uses a deliberately loose estimate of the benefit of capital outside the common program: all of it is credited with the highest individual payoff a_1. For a formed profile the bound is hQ-F+a_1(8-Q). Its slope in Q is positive when h>a_1, so Q=8 is best. Unformed investment cannot exceed 2a_1 because the other individual opportunities have negative benefit. The inequality in Proposition 1 completes the comparison. No integer restriction appears in this proof.

The bounded-transfer condition is the intersection of four intervals with a hyperplane. Necessity follows by summing their lower and upper bounds. Sufficiency follows because the set of attainable sums of independent closed intervals is exactly the interval between those sums. The result assumes every transfer in each interval is contractually permissible. It does not create such permission, and it cannot represent a contingent tax effect by a sure cash transfer without additional modeling.

The strategic cap example is solved by enumerating accepted prices. An unrestricted supplier selects the highest accepted price. An irrevocable cap truncates that set. Provided the highest accepted capped price covers cost, trade occurs and the split changes. The conclusion disappears

if the cap is renegotiated costlessly before acceptance under the stated protocol. A private cap with a different belief structure defines another game; the example provides no theorem about it.

Appendix B Reproduction and result status

The reproducible package separates the scalar payoff evaluator, vectorized profile search, economic extensions, figures and document generation. The base game contains no random sampling. All scenario probabilities are inputs. Financial results can therefore be reproduced exactly apart from floating-point tolerances and machine-dependent runtime. Serving a language model again can produce a different response, so the AI package preserves original prompts, answers and run metadata as well as a rerun procedure.

The financial checks reconstruct every base profile independently, compare all coalition deviations, verify transfer conservation and distinguish supplier markup from real cost. Grid refinement is reported for the headline high- and low-cost cases. The continuous argument establishes those cases beyond the grid; it does not certify every cell of the phase diagram as a continuous-game equilibrium. Figure 1 is explicitly a finite-menu diagram.

Result status is as follows. Cash translation and fixed-set monotonicity are established properties. The funding bound and bounded-transfer test are elementary derivations for the stated environment. The fragmentation proposition is proved within this model and independently enumerated, with no priority claim. The risk extension, application calculations and phase diagram are numerical findings. The authority cap is an illustration of established delegation logic. Commercial applicability, general mechanisms under private information and empirical formation effects remain unresolved.

This edition is a model-and-counterexamples paper accompanied by a perspective. The contribution audit found substantive prior work on the central mechanisms. A general theorem for overlapping, state-contingent financial programs, or an identified empirical result about formation, would be a different contribution requiring further work. The present package does not describe internal automated review as external academic peer review.

References

Allen, Franklin, and Adelina Barbalau. 2024. Security Design: A Review. Journal of Financial Intermediation 60, 101113. https://adelinabarbalau.com/wp-content/uploads/2025/02/SecurityDesignReview_JFI.pdf

Ang, Andrew, Dimitris Papanikolaou, and Mark M. Westerfield. 2014. Portfolio Choice with Illiquid Assets. Management Science 60(11). https://www.nber.org/papers/w19436

Bagnoli, Mark, and Barton L. Lipman. 1989. Provision of Public Goods: Fully Implementing the Core through Private Contributions. Review of Economic Studies 56(4). https://sites.bu.edu/blipman/files/2021/10/public_goods.pdf

Ben-Porath, Elchanan, Eddie Dekel, and Barton L. Lipman. 2014. Optimal Allocation with Costly Verification. American Economic Review 104(12), 3779–3813. https://www.aeaweb.org/articles?id=10.1257/aer.104.12.3779

Birmpas, Georgios, Evangelos Markakis, and Guido Schäfer. 2019. Cost Sharing over Combinatorial Domains: Complement-Free Cost Functions and Beyond. arXiv:1910.06384. https://arxiv.org/abs/1910.06384

Black, Fischer, and Robert Litterman. 1992. Global Portfolio Optimization. Financial Analysts Journal 48(5). https://rpc.cfainstitute.org/research/financial-analysts-journal/1992/faj-v48-n5-28

Boyd, S., E. Busseti, S. Diamond, R. Kahn, K. Koh, P. Nystrup, and J. Speth. 2017. Multi-Period Trading via Convex Optimization. Foundations and Trends in Optimization 3(1), 1–76. https://stanford.edu/~boyd/papers/cvx_portfolio.html

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

Buchanan, James M. 1965. An Economic Theory of Clubs. Economica 32(125), 1–14. https://doi.org/10.2307/2552442

Das, Sanjiv R., Daniel Ostrov, Anand Radhakrishnan, and Deep Srivastav. 2020. Dynamic Portfolio Allocation in Goals-Based Wealth Management. Computational Management Science 17 published online 2019. https://link.springer.com/article/10.1007/s10287-019-00351-7

Das, Sanjiv, Harry Markowitz, Jonathan Scheid, and Meir Statman. 2010. Portfolio Optimization with Mental Accounts. Journal of Financial and Quantitative Analysis 45(2). https://srdas.github.io/Papers/DMSS.pdf

DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal. 2009. Optimal Versus Naive Diversification: How Inefficient Is the 1/N Portfolio Strategy? Review of Financial Studies 22(5). https://academic.oup.com/rfs/article-lookup/doi/10.1093/rfs/hhm075

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

Diamond, Douglas W. 1984. Financial Intermediation and Delegated Monitoring. Review of Economic Studies 51(3). https://bpb-us-w2.wpmucdn.com/sites.uwm.edu/dist/8/268/files/2019/01/Diamond-1984-RES-1xhg2wp.pdf

Duffie, Darrell, and Matthew O. Jackson. 1989. Optimal Innovation of Futures Contracts. Review of Financial Studies 2(3). https://web.stanford.edu/~duffie/jstorlinks/d_jackson1.pdf

Duffie, Darrell, and Rohit Rahi. 1995. Financial Market Innovation and Security Design: An Introduction. Journal of Economic Theory 65(1). https://www.sciencedirect.com/science/article/pii/S0022053185710010

Elkamhi, Redouane, and Jacky S. H. Lee. 2026. Total Portfolio Approach: A Critical Literature Review. CFA Institute Research Foundation July 2026. https://rpc.cfainstitute.org/research/foundation/2026/total-portfolio-approach-literature-review

Fershtman, Chaim, and Kenneth L. Judd. 1987. Equilibrium Incentives in Oligopoly. American Economic Review 77(5), 927–940. https://kenjudd.org/wp-content/uploads/2021/11/1987-Equilibrium-Incentives-in-Oligopoly.pdf

Fershtman, Chaim, and Kenneth L. Judd. 2006. Equilibrium Incentives in Oligopoly: Corrigendum. American Economic Review 96(4), 1367. https://kenjudd.org/wp-content/uploads/2025/02/2006-30034348.pdf

Goemans, Michel X., and Martin Skutella. 2004. Cooperative Facility Location Games. Journal of Algorithms 50(2), 194–214. https://doi.org/10.1016/S0196-6774(03)00098-1

Gourier, Elise, Ludovic Phalippou, and Mark M. Westerfield. 2024. Capital Commitment. Journal of Finance 79(5). https://onlinelibrary.wiley.com/doi/10.1111/jofi.13382

Hansmann, Henry. 1996. The Ownership of Enterprise. Harvard University Press. https://www.jstor.org/stable/j.ctv1pncrw1

Harris, Robert S., Tim Jenkinson, and Steven N. Kaplan. 2012. Private Equity Performance: What Do We Know? NBER Working Paper 17874; later Journal of Finance publication. https://www.nber.org/papers/w17874

Iancu, Dan A., and Nikolaos Trichakis. 2014. Fairness and Efficiency in Multiportfolio Optimization. Operations Research 62(6). https://web.mit.edu/nitric/www/multipf.pdf

Jefferies. 2025. Global Secondary Market Review, January 2025. Market review of 2024. https://www.jefferies.com/wp-content/uploads/sites/4/2025/02/Jefferies-Global-Secondary-Market-Review-January-2025.pdf

Katz, Michael L. 1991. Game-Playing Agents: Unobservable Contracts as Precommitments. RAND Journal of Economics 22(3), 307–328. https://digicoll.lib.berkeley.edu/record/206222

Kleinberg, Jon, and Manish Raghavan. 2021. Algorithmic Monoculture and Social Welfare. arXiv:2101.05853. https://arxiv.org/abs/2101.05853

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. https://arxiv.org/abs/2207.07767

Mankiw, N. Gregory, and Michael D. Whinston. 1986. Free Entry and Social Inefficiency. RAND Journal of Economics 17(1), 48–58. https://mankiw.scholars.harvard.edu/sites/g/files/omnuum5931/files/mankiw/files/free_entry.pdf

Markowitz, Harry. 1952. Portfolio Selection. Journal of Finance 7(1). https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

Merton, Robert C. 1971. Optimum Consumption and Portfolio Rules in a Continuous-Time Model. Journal of Economic Theory 3(4). https://www.sciencedirect.com/science/article/abs/pii/002205317190038X

Moulin, Hervé, and Scott Shenker. 2001. Strategyproof Sharing of Submodular Costs: Budget Balance versus Efficiency. Economic Theory 18(3). https://link.springer.com/article/10.1007/PL00004200

Ostrom, Elinor. 2010. Beyond Markets and States: Polycentric Governance of Complex Economic Systems. American Economic Review 100(3), 641–672. https://doi.org/10.1257/aer.100.3.641

Palfrey, Thomas R., and Howard Rosenthal. 1984. Participation and the Provision of Discrete Public Goods: A Strategic Analysis. Journal of Public Economics 24(2), 171–193. https://authors.library.caltech.edu/records/2vnat-k7q66

Ross, Stephen A. 1978. Mutual Fund Separation in Financial Theory—The Separating Distributions. Journal of Economic Theory 17(2). https://www.sciencedirect.com/science/article/pii/002205317890073X

Schelling, Thomas C. 1960. The Strategy of Conflict. Harvard University Press. https://www.hup.harvard.edu/books/9780674840317

Sosner, Nathan, Michael Gromis, and Stanley Krasner. 2022. The Tax Benefits of Direct Indexing: Not a One-Size-Fits-All Formula. Journal of Beta Investment Strategies 13(2). https://www.aqr.com/-/media/AQR/Documents/Journal-Articles/AQR-The-Tax-Benefits-of-Direct-Indexing.pdf?sc_lang=en

Spence, Michael. 1976. Product Selection, Fixed Costs, and Monopolistic Competition. Review of Economic Studies 43(2), 217–235. https://doi.org/10.2307/2297319

Tiebout, Charles M. 1956. A Pure Theory of Local Expenditures. Journal of Political Economy 64(5), 416–424. https://doi.org/10.1086/257839

Access and review extent are recorded in the accompanying source ledger. A retrieved source is not necessarily a fully audited source. Working-paper references are identified as such. Public-document cases and their source URLs are listed in the benchmark manifest.