Polycentric Finance

Portfolio Choice, Authority, and Investment Design

Research draft · 4 October 2026

Abstract. The investment opportunity set depends partly on how investors organize. Shared implementation can preserve differences in portfolios while imposing common decisions on a particular investment. This perspective connects portfolio choice with formation costs, contractual authority, participation incentives and verification. Its organizing question is which decisions must be common, which may remain separate, and how their costs should be allocated. A companion model supplies a counterexample to monotone decentralized welfare from cheaper individual implementation. The broader research agenda concerns endogenous investment opportunities, conditional capital supply and the institutions that convert information into credible participation. Neither computational decentralization nor a new asset class is required.

1 The object that investors share

An investment is often described as a point in a space of returns. To own it, an investor also needs a route: an account, contract, counterparty, funding arrangement and legally recognized claim. Some routes are cheap at any scale. Others exist only when enough capital can be assembled on compatible terms. The available portfolio is therefore partly a product of financial organization.

This observation does not displace portfolio theory. It specifies another choice that belongs beside portfolio choice. Markowitz's allocation problem remains necessary once the relevant payoffs and constraints have been identified. Security design, intermediation and market formation have long asked how those payoffs become available. The research program considered here studies their intersection with investors who retain distinct authority over their capital. Its question is how much agreement an investment actually requires.

The economic unit of coordination need not be the portfolio. Two principals can share a loan and disagree about every other asset they hold. They can use the same fund as one component of different portfolios. A common underlying investment can support different participation amounts, tax positions or completion holdings. Conversely, apparently similar exposures can require different contracts because one investor cannot bear a particular capital call, transfer

restriction or legal obligation. Personalization should be measured by preserved economic differences, not by the number of bespoke products.

Polycentricity describes independently governed authority. A pension fund, a family office and an individual can each be a principal. Several principals may use one centralized system. An institution with a single investment committee may operate through many computers without becoming polycentric. The relevant distinction is who can authorize an action, on whose behalf, and subject to which restrictions.

This definition permits common governance where it is useful. Investors in a vehicle can delegate decisions concerning that vehicle while retaining authority over future commitments and other holdings. A fund remains a coherent institution rather than an obsolete compromise. The choice is which decisions to place inside a common mandate, and which decisions to coordinate across mandates. Both boundaries can move over time.

The proposed field consequently concerns the organization of feasible portfolios. Its most provocative claim should be stated as a research question: can declining costs of interpreting and administering investor differences make the formation of investments more contingent on the design of participation institutions? The accompanying example answers yes in one transparent economy. It does not show that decentralized investment is generally inefficient, or that a new intermediary necessarily improves on existing ones.

2 From a return frontier to an implementable frontier

A return frontier describes attainable combinations of risk and payoff under specified opportunities. An implementable frontier also depends on access, minimum size, formation costs, funding dates, authority and the allocation of expenses. The additional constraints are not all technological. Some express investor welfare; others express law, contract or the rights of counterparties. A coordinator must distinguish them before attempting to relax anything.

Consider a portfolio allocation to a company. The company name does not specify the claim. Common shares, preferred shares, a fund interest and a derivative can have different contingent payoffs and obligations. Even equivalent economic exposure may be unavailable through the same route. A right of first refusal, seller eligibility rule, investor classification or consent

condition can prevent execution. An exposure label is useful for retrieval and explanation but cannot define the investment object.

Now suppose an acceptable claim is available only if a purchasing program reaches a minimum size. Each principal's choice affects more than its own weights. By helping meet the threshold, it changes the feasible investments of other principals. This is the point at which an ordinary separable allocation problem becomes an organizational problem. A matrix of mandates and investments is insufficient unless it also represents the conditions under which the columns exist.

The complete object is a set of feasible joint arrangements. A row records what a principal can accept and authorize. A column records a program with terms and cash flows. Shared constraints represent capacity, formation, timing and common costs. Individual constraints preserve liabilities, obligations and remaining authority. An allocation is implementable only when it satisfies both kinds. A positive aggregate objective does not establish that every principal accepts its assigned participation.

$$\textit{Joint feasibility} = \textit{individual feasibility} + \textit{shared implementation conditions}$$

This expression is an organizing statement, not an additive scalar formula. In particular, individual maximum capacities cannot generally be summed. The same account may support several mandates, and several offices may act for the same principal. The joint feasible set records the tradeoffs among these uses. A scalar capacity number is a projection under stated assumptions.

Three margins should be kept separate. Reallocation changes holdings among existing claims. Repackaging changes the route or contractual division of an existing opportunity. Formation brings an additional program into existence at a real cost. A reported improvement should identify which margin moved. If the only change is a more favorable forecast, the experiment has not measured the value of organization.

The implementable frontier can expand without investors receiving the gain. A supplier may charge more once coordinated demand becomes credible. A coordinator may consume the saving through delivery costs. A pivotal participant may require compensation. The welfare account must therefore separate gross opportunity, real resource costs and transfers. Investor surplus, supplier surplus and total surplus answer different questions.

3 What established portfolio theory already supplies

Separation and spanning are the first objections to unnecessary customization. Different preferences need not imply different underlying funds. If mixtures of available claims reproduce an investor's desired contingent payoff at no greater cost, inventing another claim has no demonstrated value. Ross's separation analysis and the broader spanning literature are therefore tests that a proposed organizational innovation should survive, rather than intellectual obstacles to dismiss.

Investor differences nevertheless enter familiar portfolio problems. Merton's intertemporal framework makes future consumption and opportunities relevant. Goals-based approaches connect holdings with heterogeneous objectives and liabilities. Black–Litterman incorporates views into portfolio construction; it does not make an expressed belief true. Robustness and estimation error matter because elaborate customization can amplify uncertain inputs. DeMiguel, Garlappi and Uppal's comparison with simple diversification is a useful warning against confusing optimization complexity with economic performance.

Multiple mandates need not be disjoint subportfolios. A holding can support retirement income and reduce an employer-stock concentration at the same time. Attribution to both purposes is explanatory. It does not create two assets or double the budget available for new commitments. A total-portfolio analysis should reconcile the underlying claims before evaluating the marginal contribution of another investment.

Multiportfolio optimization already addresses shared resources, transaction costs and fairness among separately managed accounts. Iancu and Trichakis explicitly study efficiency and fairness in that setting. Distributed coordination can also reduce joint trading costs without requiring identical portfolios. These are close precedents for shared implementation. The departure pursued here is a formation and participation problem: an investor's outside option can change whether a new program exists, not only the cost of executing an existing order.

Illiquid allocation research supplies another essential discipline. A commitment is a schedule of contingent obligations, not an immediately funded asset. Commitment pacing, uncertain distributions and liquidity reserves must enter the decision before the apparent allocation is evaluated. Ang, Papanikolaou and Westerfield, Luxenberg and coauthors, and the capital-

commitment literature provide relevant foundations. Neither a mandate label nor an attractive terminal payoff excuses omission of intermediate funding needs.

An academically defensible contribution must hold these capabilities fixed across counterfactuals. A capable conventional vehicle can allow different participation amounts, shared execution and separate completion portfolios. If it can implement the same arrangement with the same costs and information, the proposed coordinator has no remaining model advantage. This null result is a condition for credible research, not a defect to engineer away.

4 Clubs, varieties and the investment formation margin

A shared investment with a fixed implementation cost resembles a club in one respect: participation changes the cost per member and may determine whether the service is supplied. Buchanan's theory makes membership part of the economic problem. Tiebout's sorting and the Ostrom tradition help distinguish overlapping centers of authority from a single hierarchy. These connections suggest questions; they do not identify investment mandates with local governments.

The analogy has limits that matter. A portfolio claim is generally divisible in payoff even when its implementation is not. Investors may join several programs. An investor can be financially able to participate but legally ineligible. Capital can be committed today and called later. A fund may bind members to common decisions concerning a particular asset while leaving broader investment choices independent. A useful model must specify which of these features generates its result.

Entry and variety provide two opposing forces. An individual route may attract capital away from a program whose remaining members can no longer cover its fixed cost. A differentiated program may instead attract participants who would otherwise remain outside, creating a valuable additional opportunity. Mankiw–Whinston and Spence supply related lessons about the social consequences of entry and product variety. They do not imply that every financial distinction should carry a surcharge.

The four-principal economy in the focused paper isolates the first force. Four investors can place up to two units each in a common program or an individual route, or remain liquid. Formation requires seven units. With an expensive individual route, all-common participation yields $19,000 per principal. With a cheaper route, one principal can obtain $22,000 alone and leaves.

The remaining six units cannot form the program. Aggregate benefit in the unique low-cost equilibrium is $22,000 rather than $76,000 in the high-cost strong equilibrium.

Mixing is allowed throughout. The example is therefore not an assertion that individuality requires another vehicle. It is a case in which an attractive outside use of scarce participation capacity makes an investor pivotal. The comparative static survives continuous quantities over a strict parameter region. It also carries an equilibrium qualification: at the high cost there is another Nash equilibrium with no formation. The best Nash outcome falls, the worst rises, and the unique strong outcome falls. Reporting only the favorable high-cost Nash equilibrium would conceal an essential part of the result.

A second numerical exercise illustrates the opposing force. Allowing a second common program raises total synthetic benefit from $32,000 to $50,000 because different investors value the two programs differently. All investors can mix programs. The increase is aggregate, not automatically Pareto improving: the first two principals receive lower unadjusted benefits and need compensation if their previous outcomes are participation floors. More variety can therefore be valuable and still present an allocation problem.

These examples support a narrower conclusion than over-differentiation as a general law. The relevant distinction is whether a new option removes support from a valuable arrangement, creates another arrangement, or does both. The empirical task is to identify these margins and the investors whose participation is pivotal. Counting bespoke portfolios does not answer that question.

5 The economics of remaining available

An investment commitment can benefit other participants before capital is deployed. A member who remains available may preserve a threshold, reduce failed-closing risk or allow a supplier to investigate a transaction. This service consumes flexibility. The option to use capital elsewhere has value, and the member may give it up even when the ultimate investment never occurs.

The appropriate object is a contingent participation contract. It specifies how much capital may be called, under which terms, until which date, and what happens if conditions fail. Its value depends on alternative opportunities and funding states. An expression of interest has different

economics because it need not survive the arrival of a better use of capital. A recorded approval that remains freely revocable should not be counted as a binding subscription.

This distinction provides a possible extension of the formation model. Let a program require contributions at a future date. Principals can either retain discretion or sell an option on part of their admissible capacity. The program purchases reliability; the principal sells a restriction on future choices. A complete analysis would price that restriction jointly with liquidity and competing opportunities, rather than treating standing mandates as costless capital.

The focused paper does not solve that dynamic contract. It computes funding gaps under specified withdrawal and failure states. Revocable support has a maximum shortfall of $60,000 in its illustrative schedule; binding support still has a $20,000 shortfall. Expected payouts and state-price values differ from those capacities. An external guarantor's price requires a valuation of contingent losses, capital requirements and counterparty risk. Prefunding consumes liquid resources rather than eliminating cost.

The often proposed autonomy–reliability–self-sufficiency trilemma needs a funding premise. Correlated shocks alone do not prove impossibility. Sufficient internal resources can cover them. Conversely, a shortfall can arise with independent shocks. The valid statement is conditional: if some supported state and permitted withdrawal pattern leave admissible resources below the required funding, reliable formation needs an additional resource or a change in the promised arrangement. Self-sufficiency means formation without an outside guarantor; it does not mean that internal reserves are free.

Funding reliability is also different from investment acceptance. An accepted contractual obligation can survive the later revocation of advisory authority. A capital call satisfying that obligation may remain due even when the investor no longer prefers the investment. Payment authority is another distinct permission. A system that treats every change of preference as cancellation of an obligation manufactures flexibility the contract did not provide.

The empirical implication is to measure the states in which support disappears. A failed close caused by investment rejection differs from one caused by documentation, bank processing or a capital call elsewhere. Only some failures can be reduced by interpretation technology. A chronology of opportunities, permissions, reserves, acceptances and payments is more informative than a single conversion rate.

6 Participation payments and the limits of a difference price

The value of a preserved distinction is not its implementation cost. A tax restriction may be expensive and indispensable. A minor preference may be cheap but irrelevant. A principle that charges by the number of constraints rewards wording changes and penalizes careful disclosure. A defensible allocation rule must refer to economic consequences under a specified counterfactual.

In the main example, the pivotal investor's departure costs other investors their program. Under a separately specified unanimous subscription contract, a payment of $4,000 makes remaining individually acceptable. The other three each contribute one third of that amount, leaving all four strictly better off than their jointly feasible low-cost outside outcome. This is compensation for participation. It can also be described as a reassignment of the formation surplus. Nothing in the arithmetic establishes truthful reporting or a unique bargaining division.

The original completion-payment implementation was unstable with continuously divisible choices: a donor could reduce its quantity slightly and avoid paying while the common program still formed. The revised all-or-none contract conditions every subscription on acceptance of the specified four allocations. This changes the contract, not the physical seven-unit formation threshold. It supports formation in a separate subscription game, with nonformation Nash equilibria still present.

The payment does not preserve every investor's previous payoff. The old common allocation gave each $19,000. The new outside option entitles the pivotal investor to at least $22,000 for voluntary participation; preserving $19,000 for the others would require $79,000 in total. The common arrangement produces only $76,000. Cheaper individual implementation has improved an outside option without creating enough additional shared surplus to maintain every historical outcome. The coordinated benchmark must say which participation floor it uses.

Marginal-externality charges offer another approach. One can construct a utility adjustment equal to the change in other participants' welfare, causing unilateral incentives to track changes in total welfare. That alignment is useful but does not guarantee a feasible budget. In the saved four-principal calculation, paying every investor its full external contribution requires an outside

subsidy of $240,000 at the common outcome. Counting overlapping contributions as independently payable claims exceeds the actual surplus.

Even an exact potential game can retain inefficient equilibria. A single investor may be unable to cross the formation threshold, so no unilateral improvement exists from the no-formation outcome. A marginal rule then aligns local incentives without solving coordinated entry. This is why the compensation contract is evaluated for both Nash and strong equilibrium rather than declared an all-purpose implementation mechanism.

Restrictions on transfers can also bind. A program may prohibit cross-subsidies, a participant may lack authority to make a side payment, or an expense policy may allow only specified charges. The transfer-feasibility calculation in the focused paper gives interval conditions for individual participation under a fixed allocation. It is not a theorem about coalition stability. If permitted transfers cannot cover the required compensation, aggregate surplus alone is insufficient.

Combinatorial cost-sharing and mechanism-design research provide the right benchmarks for stronger claims. Budget balance, efficiency, voluntary participation, coalition stability and truthful reporting are different properties. A facility-location core theorem concerns a particular characteristic function and relaxation. The integrality gap of an arbitrary portfolio program does not determine whether its cooperation is stable. Our continuous counterexample deliberately makes no such identification.

7 Authority as an economic constraint

An authority grant identifies a set of permitted contingent actions. It can bound price, quantity, instrument, issuer, time, counterparty or funding source. Grants can be partially ordered by set inclusion. With investment opportunities and prices fixed, broader permission can weakly expand the principal's choice set. The principal can ignore new actions. This is a feasibility observation, not a statement that broader delegation is always desirable.

Prices and offers can respond to the grant. A supplier that observes a credible price ceiling may accept a lower price than it would otherwise demand. The focused paper uses a deliberately small bargaining example to exhibit this effect. The investor values an investment at 15, the

supplier's cost is 9, and an immutable, observable cap at 12 divides the surplus. A cap below cost prevents trade. A freely amendable cap provides no assumed commitment benefit.

Schelling and strategic-delegation research supply the intellectual background. Fershtman–Judd study incentive contracts in oligopoly; their published corrigendum must accompany any reliance on the affected result. Katz qualifies simple claims about observability and commitment. The example here is proved directly under explicit assumptions and does not import an unrestricted theorem from either paper.

A registry can make the current text of a restriction verifiable. It cannot by itself prevent amendment, side agreements or renegotiation. Commitment value depends on which of those possibilities the counterparty expects. Nor should the system reveal a principal's entire feasible set merely to demonstrate eligibility for one offer. A supplier that learns the reservation price or the identity of a pivotal investor may capture more surplus.

This creates a disclosure problem. The coordinator wants to establish that an allocation can close while revealing as little as necessary about willingness to pay and alternative opportunities. A certificate might state that a specified price–quantity offer is presently supportable, subject to named conditions, without publishing the underlying demand schedule. Such a statement must be supported by current evidence and must disclose its expiry. Its market value and potential adverse selection require study.

The appropriate next model has private information and endogenous verification. It asks which predicates a principal can certify, how certification changes supplier behavior, and whether investors strategically withhold information. Ben-Porath, Dekel and Lipman's costly-verification setting excludes monetary transfers; that restriction matters when drawing parallels to investment allocation. The present complete-information model cannot establish incentive compatibility for reported mandates.

8 Four applications and their distinct tests

Private funds combine uncertain calls, distributions and long residual obligations. A participation program might reduce formation cost or expand access, but those savings can be outweighed by reserves and fee layering. Evidence on aggregate private-fund performance does not identify the joint cash-flow process of a particular investor's commitments. Harris, Jenkinson and Kaplan

inform the performance context; research on cash-flow timing addresses another part of the problem. The synthetic schedule here makes no inference that venture or buyout investments carry a free return premium.

The private-fund exercise retains a call after month 120 and reports distributions delayed beyond that date under stress. This prevents a finite horizon from making commitments disappear. Its questions are modest: which cash flows move outside the reported horizon, and how much funding remains due? It is not a calibrated optimal pacing policy. A serious empirical application would require dated calls, distributions, fees and surviving obligations, including funds that have not liquidated.

Secondaries separate a claim's economic value from its reported NAV. A discount can compensate for stale marks, risk, expenses, unfunded commitments or urgency. It is not itself evidence of mispricing. Jefferies's review of 2024 provides a dated market-scale and pricing reference, not a forecast for an individual sale. The sensitivity exercise varies purchase price and unfunded obligations while holding its synthetic payoff process explicit.

The bilateral risk-transfer example makes a different point. Buyer and seller use identical state probabilities. The seller already holds four units of a risky claim; the buyer holds none. A one-unit transfer at $90,000 raises the seller's certainty equivalent by about $23,613 and the buyer's by about $7,501. No better information or expected excess return is introduced. Their different starting books and risk tolerances create a mutually beneficial trade. At $100,000 the buyer loses in this example. Fees, taxes and consent costs would narrow the bargaining interval.

Public securities offer a demanding null benchmark because ordinary funds, separate accounts and shared trading already span many desired outcomes at low cost. Direct indexing can preserve tax-lot distinctions while sharing portfolio exposures. Its benefits depend on realization, available losses, future basis and fees. The illustrative tax calculation records later recapture of a basis reduction, so the benefit is deferral under the assumed horizon. It does not insert a published tax-alpha estimate as an additional expected return.

Credit makes contractual conditions unusually visible. A facility may specify minimum draws, increments, exceptions, permitted uses and conditions precedent. The public-document benchmark includes a minimum amount with an exception for the remaining commitment, and a distinct reimbursement provision with its own treatment. Treating the headline minimum as

universal would reject some permitted actions; treating the exception as general authority would permit too much. Neither the filing nor its extraction establishes that a particular borrower currently satisfies every condition.

Real assets fit the same framework when a property, infrastructure asset or project requires a common implementation arrangement. They are not numerically calibrated here. Construction draws, operational governance, refinancing and sale rights would need explicit treatment. Mentioning another asset class does not establish that the same numerical mechanism dominates there. Across applications, the research question is whether shared formation remains valuable after the actual claim, cash flows and ordinary alternatives are specified.

9 Why the technology margin matters now

Persistent machine interpretation can lower the cost of maintaining a representation of an investor's circumstances across changing documents. A system can propose how a new notice changes obligations, connect an expressed objective with candidate claims, and direct attention to incompatible terms. These tasks have historically required repeated interpretation across institutions, formats and dates. The present study tests a small portion of that process. It does not quantify a historical decline in total intermediation cost.

The important potential change is frequency. An annual policy review can tolerate a manually assembled snapshot. A service that continually reevaluates programs needs to keep holdings, commitments, permissions and available terms current. As interpretation becomes cheaper, the feasible frequency of review may rise. This can make previously ignored differences worth preserving. Whether those differences justify a different investment remains a financial question.

Two technology costs must remain distinct. Interpretation and verification concern the evidence needed to know an admissible action. Implementation concerns the accounts, contracts, vehicles, tax treatment and transactions needed to perform it. Better language models may reduce the first without changing the second. They can also uncover obligations that reduce apparent available capacity. More accurate records need not produce more investment; avoiding an infeasible action is a different benefit.

The architecture should preserve this division. Language interpretation proposes records and unresolved questions. Deterministic checks apply validated constraints to proposed actions. A

portfolio engine evaluates complete bundles rather than ranking documents independently. Transaction processing checks current versions and consumes shared capacity atomically. Subsequent evidence updates the economic book without rewriting the historical basis for earlier decisions.

An attractive user interface is not evidence that these steps work. The system may generate a fluent explanation from an incomplete book. It may use a corrected record before its historical knowledge date. It may confuse permission to research with permission to commit. It may reserve the same principal's capital through two offices. These failures are economic and institutional, even if they arise from software.

The strong why-now proposition is consequently a testable complementarity: lower interpretation costs may make finer investment coordination economical when reliable records and permitted implementation are already available. It is not that AI invents diversification or abolishes financial intermediaries. The organizations best placed to use the technology may be existing institutions with credible execution, supplied with a better way to preserve and reconcile individual constraints.

Cheaper technology can also weaken a valuable coalition by improving outside options. The main comparative static makes that possibility explicit. Coordination can preserve a shared program through compensation, but the outcome depends on transfer authority and participation incentives. AI therefore changes the organizational problem; the current evidence does not establish that a particular organizational answer is universally necessary.

10 What the bounded AI study actually identifies

The evaluation compares an evidence index and a deterministic action guard in a two-by-two design. Sixty synthetic dossiers include twelve development cases and forty-eight held-out cases. Twelve public-document cases add scoped provisions from eleven filings. The study performs 304 evaluation calls, including sixty-four predetermined repeats. Two development calls are separate. Every model invocation uses an isolated prompt; no live financial action is possible.

The evidence index is deliberately modest. It reorganizes the same source paragraphs by their labels and supplies them alongside the original text. It does not infer a complete semantic book or maintain an evolving account over months. The guard checks selected financial and

permission conditions directly from available evidence. It does not read hidden labels. A post-review probe nevertheless defeated its regular-expression source parser with an unauthorized attachment; this is a fixture checker, not a provenance-secure authority gate. This design measures representation and action checking on bounded inputs, not the full value of a production registry.

Results are mixed. Of 304 evaluation calls, 299 completed; five timed out and remain in the denominator. No completed call proposed an unsafe execution under the fixture labels. The guard therefore intercepted nothing. Validation was postprocessing: its model prompts were identical to the no-validation prompts, so differences between their raw answers reflect repeated-call variation. Seven primary synthetic calls chose review by citing the instruction that no live action was authorized; their apparent error is confounded by the question's simulation-scope ambiguity. Indexed calls more often added a USD suffix that was economically correct but missed the frozen exact-string scorer. Strict and unit-normalized scores are both retained so that these format deviations are not mistaken for wrong financial quantities.

Public performance was near the ceiling on the scoped questions. The remaining disagreements concerned whether price quotations alone established equivalence of rights; internal review marked that item disputed, and all other public items were correct in every arm. Labels were audited by agents, not validated by human legal specialists. A public case asking whether a filing alone authorizes an unidentified investor has an easy negative answer. Success on that question provides little evidence about correctly approving a genuinely executable transaction.

The cost account includes reported tokens, latency and construction time for the deterministic index. Human review prices are assumptions. Because the apparent unnecessary reviews are confounded by prompt ambiguity, this study does not assign them lost investment gains. A loss-function specification is retained for a better controlled subsequent study. Model-implied regret requires a financial payoff model beyond a correctness label. The study therefore informs a narrow cost–accuracy comparison rather than estimating the total cost of preserving a distinction.

Errors must also be grouped by their common origins. The synthetic dossiers share one common generation template, divided into twelve administrative groups, and two public cases use the same credit agreement. The sixty-four repeats are repeated measurements, not independent investment opportunities. Their descriptive agreement rate is useful for reproducibility but does

not identify market-wide monoculture. Kleinberg and Raghavan's work motivates a broader concern about common algorithmic judgments; testing that concern would require choices across investors, not just documents.

A larger experiment should compare genuinely maintained semantic records with equally capable retrieval, include the cost of constructing and correcting those records, and test erroneous evidence as well as missing evidence. It should use independent expert labels, different model families, realistic document bundles and recorded intervention time. The current null validation finding is a reason to improve the challenge set, not to manufacture unsafe cases after the fact and report them as held-out evidence.

11 Research designs that could change the argument

The first priority is an observed formation margin. Assemble matched prospective programs, including those that failed, and record the minimum viable size, formation expense, available routes, proposed terms and reasons participants declined. A useful unit of observation is an opportunity–principal decision with a dated feasible outside option. The design must distinguish insufficient demand from unavailable permission, documentation delay and rejection on investment merit.

One hypothesis is that credible conditional participation reduces a supplier's expected failed-closing and servicing costs. It predicts a difference between nominal interest and support that survives verified conditions. The dependent variables should include supplier effort, reserved allocation, quoted price, formation probability and who captures the saving. A supplier can benefit without making a concession. Randomizing information presentation may identify interpretation effects; assigning actual capital obligations requires a separate lawful arrangement and is outside this program.

A second hypothesis concerns the pivotality of participation. The model predicts that an improvement in one investor's outside option can harm other investors when the investor is necessary for formation. Empirical evidence would require identifying the outside-option change without attributing the same change to common investment quality. Thresholds can be negotiated in response to demand, so a discontinuity at a stated minimum is not automatically an exogenous experiment. Documented threshold revisions are outcomes to explain.

A third hypothesis concerns voluntary compensation. Observe whether explicit shared-cost allocations preserve formation relative to ordinary pro-rata charges. Record rejected offers and transfer restrictions. A successful program is not sufficient evidence: compensation might merely move surplus to a pivotal party or subsidize an otherwise unattractive investment. The test should compare net participant outcomes, outside options, real expenses and supplier economics under common information.

A fourth design examines authority disclosure. A supplier could receive a narrowly specified readiness certificate or a broader description of investor constraints, subject to consent. The theoretical question is whether extra disclosure improves execution enough to outweigh bargaining exposure. Before any field experiment, a laboratory mechanism with induced values could test the distinction between verifiable and freely amendable limits. The current cap example is a proof of possibility, not evidence that real suppliers respect such certificates.

A fifth design concerns technology and organizational choice. Measure how interpretation assistance changes review time and which consequential distinctions survive. Then connect those changes to decisions under a separately specified financial model. Record construction errors, correction effort and incidents that force abstention. A lower extraction cost may be offset by more proposals requiring expert attention. Costs should be reported by workflow stage rather than compressed into an unobserved universal parameter.

These studies need a common event vocabulary. At minimum: first observation, mandate version, evidence availability, eligibility decision, provisional allocation, reservation, acceptance, funding and final ownership evidence. Corrections carry both effective and knowledge dates. This chronology makes causal claims more demanding but also more credible. It prevents a later successful closing from retrospectively turning an earlier expression of interest into executable capital.

12 Institutional design and ownership form

A coordinator chooses not only investments but rules governing access and expenses. That creates conflicts when it also supplies products or manages an affiliated fund. A program can be socially valuable while its allocation policy favors the coordinator's other clients. The

counterfactual must include which investor receives scarce capacity, not only the program's total return.

Ownership form can affect those incentives. Hansmann's analysis of enterprise ownership asks when the costs of market contracting and collective governance favor ownership by a particular patron group. Mutual or customer-owned structures may internalize some distributional concerns but incur governance costs. Investor ownership is neither necessary nor sufficient for efficient coordination. The relevant comparison includes decision rights, financing, accountability and the cost of resolving disagreement.

A commercial platform faces a related measurement problem. Model-implied certainty-equivalent gains are not willingness to pay. A participant may value convenience, distrust delegation or already have a capable investment office. Charging a recurring fee can change the participation constraint that made a program viable. The economic service must be evaluated after platform delivery costs and any supplier revenue share, not before them.

An institutional partner can contribute distribution, implementation or investment capacity, but those roles should not be conflated. A party that validates records does not necessarily advise the investor. A custodian does not necessarily control an investment decision. An administrator's register may be authoritative for one legal interest and only informative for another underlying exposure. The system must identify where legal authority resides rather than presenting a consolidated display as consolidated power.

The implementation note applies these distinctions to Argosy and Megalith. The research argument does not depend on those entities, their particular sourcing capabilities or any proposed fees. Its contribution would remain relevant if a conventional institution implemented every useful arrangement. That neutrality is necessary to identify a genuine organizational mechanism.

13 A research program with rejection conditions

The formation thesis should be rejected for an application if ordinary mixtures and existing routes reproduce the outcome at equal or lower total cost. The sourcing thesis should be rejected if suppliers do not respond to credible demand, or if their response captures all investor benefit after execution costs. The interpretation thesis should be narrowed if maintaining structured

records costs more than the review effort they save. These are distinct failures; one need not invalidate the others.

The program should also abandon results driven by incomparable forecasts, favorable private marks or omitted residual obligations. A discount to preferred stock does not establish attractive common-share economics. A positive aggregate certainty equivalent does not justify imposing a loss on an unwilling member. A high completion rate does not establish that the completed transactions were suitable. The underlying financial questions remain demanding even when the interface becomes simpler.

The current package establishes a continuous-quantity formation counterexample, a fixed-allocation transfer test, a funding-gap identity and a limited bargaining example. It supplies exhaustive finite calculations and a bounded interpretation evaluation. The novelty gate did not establish an original general theorem. The appropriate manuscript form is therefore a model-and-counterexamples paper and this perspective, rather than a claim to replace modern portfolio theory.

The next contribution worth pursuing is the joint design of conditional participation and disclosure under private information. It would price the flexibility sold by a principal, the reliability bought by a program, and the information revealed to suppliers. A useful result would identify when verified partial commitments improve formation without requiring full disclosure or unrestricted transfers. A negative result would be equally consequential if it showed that the proposed coordination cannot overcome incentive or commitment limits.

The lasting idea is precise. Investors may need to agree on an implementation without agreeing on a portfolio. Their participation can alter what is available to one another. Technology can lower the cost of preserving those differences, but the distribution of gains depends on contracts, funding and incentives. The investment frontier is partly designed; the task of financial research is to identify the designs that are feasible and worth their cost.

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.