The Formation Game

A briefing on *Commitment and the Investment Frontier*

Argosy Research

October 2026

Abstract

Portfolio theory takes the available investments as given. For most private investments the assumption fails: they exist only once several investors commit, and each investor's willingness depends on what the others will do. Komkov [2026] treats the frontier as the outcome of that game and asks which of a fund's functions must be shared for investments to form. Commitment support, compatible capital and a bounded amount of priced discretion must be held in common; proportional exposure need not be; and the moment at which participation binds sets the price of the capital a closing requires, bounded when investors bind before the common imbalance is known and growing with the square root of their number when they keep a veto. Cheaper separate accounts and cheaper collective search, both of which software is now delivering, act on these quantities in opposite directions, which is why the question has become practical.

1. The frontier is not a given

Portfolio theory begins after the hard part is over. Markowitz [1952] hands each investor a set of assets with known means and covariances and asks which combination he should hold, and the separation theorems explain when a handful of common funds can serve investors whose circumstances differ, so that pooling is the cheapest way to hold what everyone would have chosen anyway [Tobin, 1958]. Seventy years of practice rest on that picture, and most of it is right for public markets, where the assets exist whether or not any particular investor shows up. The quiet assumption is that the opportunity set is a datum.

For most private investments the assumption fails in a specific way. A growth round, a buyout, a credit facility, an infrastructure project, the first close of a fund: none of these exists until several investors have committed to it, and each investor's willingness depends on what he expects the others to do. The set from which an investor chooses is therefore the outcome of a game, and the frontier he faces is decided by that game before any optimization begins. The company that could absorb a Series B exists whether or not any particular lead exists; the financing that brings it onto anyone's frontier is manufactured by a coalition, and coalitions fail to form even when every member would gain.

The failure has a familiar shape. A commitment that others do not match is wasted, because the investor who made it bears the cost of standing ready, negotiating, reserving capital and forgoing alternatives while the deal fails. So the outcome in which nobody commits is an equilibrium however attractive the investment: each investor refuses because he expects the others to refuse, and the expectation confirms itself. Economists have studied the structure under the names of threshold public goods, crowdfunding and contracting with externalities. The paper moves it to the center of portfolio theory and asks what institution it requires.

The reason to ask now is that the two technologies governing the answer are changing at once. The cost of running an account for a single investor, tailored to his taxes, liabilities, exclusions and existing holdings, falls as software reads mandates, reconciles records and keeps an attributable history of holdings, and the separation-theorem reason for pooling, which was always a matter of cost, weakens with it. The cost of collective search, of finding the investors whose mandates admit an investment and assembling their commitments on compatible terms, is falling for the same reason, and that cost has always set which investors a formation reaches. A theory that treats the frontier as given cannot say what either change does. A theory that treats it as the outcome of a game can, and its answer stops well short of the fund's obsolescence: the fund bundles five things, each can be priced, and the bundle will come apart along a line the theory locates.

2. Five things a fund holds in common

A fund is usually described by what it does: it raises committed capital, gives a manager discretion over where the capital goes, and gives every investor the same proportional slice of the result. The paper takes those three functions apart into five things that a fund holds in common whenever it forms an investment, because each can be priced separately and the prices differ by orders of magnitude.

The question is which of these must be shared for investments to form, and what it costs investors to keep the others separate. The answer has two parts. Commitment itself, under simultaneous commitments, requires support of a size the paper computes, whose cheapest robust form is a promise to pay in the state in which the investment does not exist. A set of investments financed together requires

Held in common	What it does for formation	What the paper finds
Commitment support	Makes the investment form although each investor would refuse if he expected the others to refuse	Must be shared. Its least robust size is computed exactly, it is small, and it is never drawn when willingness is known
Compatible principal	Delivers capital whose permissions allow it to finance what the investment requires	Must be shared. Aggregate capital is the wrong measure; compatibility is the right one
Discretion	Lets a manager place that principal where it is needed	Must be shared in a bounded amount. Its required size diversifies like returns, and its price is bounded when many grant a little
Early binding	Binds each investor before the investments and their terms are known	Sets the price of the capacity a closing needs: bounded if investors bind before the common imbalance is known, growing with the square root of their number if they keep a veto
Common exposure	Gives every investor the same proportional slice of the result	Not required for formation. It economizes on records and incentive pay, and that is all it does

principal whose permissions cover it, discretion whose amount has the structure of diversification, and a decision about when each investor's participation binds, and on that decision the central result turns. Investors who keep the right to refuse until the terms are known impose on the closing a capital requirement that grows with the square root of their number under every honest mechanism for pricing the places; investors who bind before the common imbalance is known, with the premium that the pricing collects returned to them, leave a shortfall that stays bounded however large the book. Common exposure is not required at all. It affects whether investments form only through which investors are willing to commit, and what it saves is the cost of records and of paying a manager.

The arrangement in which investors keep their own exposures while committing together the paper calls a *book:* a set of investors who commit to finance a set of projects as a single event, each holding what he wants, with places in a project that turns out to be crowded carrying a price, and the premium so collected either kept by the arranger or returned to the investors. The book keeps everything a fund needs for formation and drops the one thing it does not.

3. The price of committing together

Start with the simplest case: willingness is known and commitments are made at once. Each investor who commits bears an expense of standing ready, the negotiation, approvals and reserved capital that are lost if the investment fails, and the investment forms only if the investors who commit constitute a coalition able to form it, whether that means a count of seats, a threshold of capital or a required set of roles. Refusal by everyone is then an equilibrium. The question is what removes it.

Two kinds of payment can be written. A payment to an investor for committing, whatever the others do, makes commitment a dominant choice, but it is paid on the path on which the investment forms, so someone finances it every time. A payment conditioned on failure, an escrowed payment that an investor receives only if he committed and the investment did not form, removes the bad equilibrium

with nothing paid in equilibrium. Once enough investors are assured that committing cannot hurt them, their commitment is dominant, and the rest commit because the coalition is now certain to complete. The paper computes the least promise that does this robustly, meaning that formation is the only outcome rational investors can reach by any chain of reasoning about one another. In words: over the coalitions able to form the investment, find the one whose members' expenses of standing ready, less the largest single expense among them, sum to the least; that sum is the least liability any scheme can carry, and an escrow comes as close to it as one likes. The most expensive member of the coalition need not be assured, because once the others are, his commitment completes it and he commits for the gain alone.

The striking part is the optimality. One might expect cleverer contracts to do better, terms that depend on who else commits, that vary with the identities of the others, that pay investors who decline, since the literature on contracting with externalities shows that such terms can make participation the unique outcome at no cost on the equilibrium path. None of them does. Among all payment schemes that depend on the profile of commitments, however elaborate, none makes formation the only rationalizable outcome with a liability below the escrow's bound. What the clever contracts save on the equilibrium path they carry off it, as a promise, and the escrow carries the least promise possible. Schemes that rely on investors coordinating on the good equilibrium can be cheaper, but they leave failure rationalizable, and their predictions rest on beliefs no finite evidence can confirm.

Because the promise is never drawn when willingness is known, it can be reused. A single promisor who stands behind many investments with disjoint investors, honoring claims in an order of seniority, needs capacity equal to the largest single requirement rather than their sum: each investment forms in turn and makes no claim on the capacity that then assures the next. This is an economy of scope that requires a common promisor but not common ownership. Separately owned accounts backed by one promisor enjoy it; separately backed accounts do not; and it is the part of a fund's work that only shared terms can do. Time substitutes for capacity in a second way. If investors commit one at a time in a public order, each commitment binding if the investment forms, the last investor whose commitment completes a coalition commits, backward induction carries the argument to the first, and no promise is needed among investors who believe in one another's future rationality. When dates are scarce, an escrow released as each public round closes needs only the largest round's requirement, so that one hundred indispensable investors committing in ten rounds need a promise of nine expenses rather than ninety-nine. A calendar divides the promise as well as shortening it.

Private willingness changes the problem. A payment upon failure attracts investors who would rather the investment failed, and when such investors can be arbitrarily close to indifference, no scheme that pays only committing investors and charges nothing screens them out. Two instruments do. A charge, whether a premium paid upon formation or a fee paid upon committing, falls in the state an unwilling investor hopes for. An order lets an early investor make do with a payment equal to his expense, which no unwilling investor gains by collecting. An order also wastes less, because a failure is discovered before others have spent much waiting for it, and the investors likely to be unwilling, or whose expense of standing ready is small, should commit first.

The calculation yields a puzzle. For a $300 million fund raised from thirty investors whose expense of standing ready is one percent of each commitment, the promise that eliminates strategic nonformation when every investor is required is about one percent of the fund. Carrying it for a year costs half a basis point if a bank supplies it as an undrawn line and about ten basis points if a sponsor holds the reserve itself, against organizational expenses capped at twenty to twenty-five basis points and management fees whose present value runs to ten or fifteen dollars per hundred committed. If contingent capacity

is this cheap, escrows that pay investors more than their expenses upon a failed closing should be common. They are not. Practice relies on leads whose commitments are informal, on orders where earlier commitments bind, on refunds, and, in public offerings, on anchors. The theory supplies two of the three explanations. Among sponsors and investors who meet repeatedly and can rely on one another's future rationality, an order costs nothing, so the escrow is needed only where commitments must be simultaneous, which is why offerings that cannot be sequenced use firm-commitment underwriting, an anchor. Under private willingness, moreover, an escrow above expenses is drawn whenever some investor proves unwilling; in the calibrated fund its expected payments are about $730,000 per attempt, against $360,000 of expense wasted in an order. Capacity that is never drawn when willingness is known is drawn on the path of play when it is private, and its cost is then the waste it insures rather than its carry. The third explanation lies outside the model. A lead whose term sheet does not bind is believed because reneging would cost it future business, so that reputation does the work of capacity within the limit set by the future surplus of the relationship. Whether explicit, recorded commitments can replace the rents that make informal ones credible is one of the questions on which the institutional implications turn.

4. Compatible capital and the diversification of discretion

The second part of the paper separates commitment from exposure. Suppose a book of investors commits to finance a set of projects as a single event, each investor holding only what his mandate permits. The first result is a feasibility condition. The projects can be financed exactly when every subset of them is covered by the capital permitted to finance that subset. Whenever permissions are incompatible, aggregate capital overstates the population's capacity to invest, and the shortfall is the largest gap, over all sets of projects, between what they require and what the investors permitted to fund them can supply. Two projects needing one unit each cannot both be financed by two investors of one unit each if both permit only the first, though there is capital enough in aggregate; let one of them also permit the second and the shortfall vanishes, with nothing about the projects, the capital or the forecasts changed except the compatibility of the commitments. This is a theorem on flows, Hall's condition in its divisible form [Hall, 1935, Gale, 1957], and its content for finance is that the value of an additional investor or an additional permission is what it makes feasible for everyone else. An investor whose mandate duplicates abundant capital adds nothing; one whose permissions cover the binding set of projects is worth the whole gap. Missing principal and missing assurance are different quantities with different remedies.

The second result concerns the discretion needed to absorb the mismatch between what investors permit and what projects require, and it has the structure of diversification. If each investor's restrictions differ from the others' idiosyncratically, the imbalances across projects partly cancel, and the share of capital that must be flexible, free to be moved from the project its owner favors to the one that is short, shrinks like the square root of the number of investors. With ten thousand investors and two projects, about two and a half percent of the capital suffices for ninety-nine percent reliability. If instead a common component runs through the mandates, because a share of investors follow the same view, the same adviser or the same recommendation, the required flexible share converges to a fixed fraction however many investors there are; one mandate in five following a single recommendation between two equal projects requires a sixth of all capital to be flexible. Permissions are therefore a dimension of diversification distinct from returns, and their systematic part determines how much discretion must be

held in common. A manager whose judgment moves every investor's mandate together is the limiting case.

Discretion has a price, and the mechanism that charges it matters. An efficient book moves only the investors needed to balance its projects, and among those the ones who least mind being moved. With two projects, the mechanism that implements this in dominant strategies is a uniform price on the places in the crowded project, paid by each investor who keeps such a place and spared by those moved to the other. Under idiosyncratic mandates the expected total cost of the discretion the book uses stays bounded as the book grows, because the number of investors who must be moved is of the order of the square root of the population while the cost of each move, taken from the cheapest, is of the order of one over the square root. A fund that imposes pro rata exposure instead imposes on the average investor a cost that does not shrink. The flexible capital a book needs is thus of order $\sqrt{n}$, while the flexible population needed to finance it cheaply is of order n, so that a pool of discretion should be broad and shallow, granted by many investors and used on few. Such a book finances every project that a blind pool finances, admits the narrow investors whom the pool cannot, and does so at a bounded cost. Separate consent, under which each investor finances only what he favors, finances the whole set with a probability that tends to zero; thirty investors each willing with probability ninety-nine percent form a deal-by-deal syndicate only seventy-four percent of the time.

The price must be charged on the right side. Paying investors for flexibility, rather than charging those who keep a crowded place, invites every investor to declare the crowded project in order to be paid for leaving it, and a declared preference is a claim that a register can record but cannot verify. Pricing the crowded project leaves no such incentive. One caution closes the section. Discretion should absorb imbalances that reflect differences in fit, not imbalances that reflect knowledge; when investors avoid a project because they know something about it, flexible capital inherits the investments that informed investors decline. The paper gives the band of imbalances within which absorbing the mismatch still has positive expected value, and in a large population almost every imbalance should be respected. The scale of an imbalance, of order $\sqrt{n}$ for independent differences in fit and of order n for shared information or shared taste, is evidence about its source.

5. The right to decide later

Everything so far takes the investors as given. The central result asks which investors remain willing once a book's terms are known, and the answer is that the right to decide later has a financing cost whose size depends on one choice: whether participation becomes binding before or after the information common to all investors arrives. The benchmark is deliberately simple: two equal projects, every investor flexible, each with a private margin of value from a place in the book, a private cost of being moved, and a favored project that is private too.

Consider first a book in which investors keep a veto. They report their favored projects and costs, the assignment and its price are announced, and each then accepts his place or withdraws; replacement capital, committed in advance to take any vacated place on that place's terms, fills the vacancies, and the book closes if the vacancies do not exceed it. The investors who refuse are those who hold a place in the crowded project, pay the premium for it, and have a margin smaller than the price, and their number is proportional to how crowded the project turned out to be. That is the whole mechanism. Preferences were drawn separately for each investor, yet one clearing price, applied to all of them, links decisions that differences in fit alone would have left unrelated. Independent differences become correlated refusals,

and the shortfall, of the order of the square root of the number of investors, fluctuates on the scale of its own mean rather than concentrating around it. At critical load the probability that the book closes tends to zero as it grows. Replacement capital proportional to $\sqrt{n}$ buys a confidence level and nothing more, and the capital a prudent arranger must reserve is set by a high quantile of the common shock, not by its average.

Consider instead a book in which investors bind before the favored projects and costs are known, each knowing only his own margin, and in which more investors are invited than the book needs. Each investor's decision to enroll now depends only on his own margin and on an expectation that is the same for everyone, so that refusals are independent, their number concentrates, and the replacement capital, here simply the extra invitations, need cover the mean. To first order the expected shortfall is the same as under the veto when the premium is kept by the arranger, but the capital needed for ninety-nine percent reliability falls toward thirty-one percent of what the veto requires. If the premium is also returned to the investors within the book, something different happens. The premium was a transfer among the investors, of order $\sqrt{n}$ in total; the real cost of the book's discretion, the cost of the moves actually made, is bounded. Returning the premium leaves each investor an expected net burden equal to his share of that bounded cost, of order $1/n$, so that the expected number who decline is bounded however large the book, and a book in which every investor is indispensable closes with a probability bounded away from zero, about seventy-eight percent in the benchmark, which two or three extra invitations raise above ninety-five or ninety-nine percent. The conclusion survives when investors know their own favored projects and costs at enrollment, because the share of the premium an investor receives when the other project is crowded nearly offsets the burden he bears when his own is. A rebate that leaves out each investor's own report, the mechanism-design counterpart of asking what the price would have been without him, makes truthful reporting exactly dominant at the cost of a bounded premium that cannot be returned [Bailey, 1997]; it lowers the benchmark closing probability from seventy-eight to forty-seven percent, and because no investor in the book receives it, it can pay for standby capacity supplied from outside without disturbing the incentives inside.

The third result closes the door on the hope that some cleverer pricing of the places could escape the veto. Under every mechanism that makes the efficient assignment with dominant-strategy incentives, the capacity a closing draws on, whether replacement capital for the places refused or cash behind the mechanism's own payments, is of order $\sqrt{n}$ in expectation, and it is of that order for any reliability above one half; below one half a bounded reserve can suffice, and the threshold is exact. The reason is a payment identity that every such mechanism obeys [Green and Laffont, 1977]. An investor's transfer may depend on everyone's report but his own only through which side he is placed on, and the two transfers must differ by the price he faces. The mechanism must therefore guess, from the others' reports, the side on which each investor's own report will fall, and two states that differ only in one investor's favored project cannot both be guessed right, so the states in which a mechanism is right about nearly everyone cannot carry more than half the probability. In every other state it is wrong about a positive share of its investors, each of whom either bears a burden and refuses with a probability proportional to it, or is paid in cash, and a burden of order $1/\sqrt{n}$ on a positive share of n investors is a burden of order $\sqrt{n}$ in total. The threshold rests on the privacy of favored projects. Where mandates are fixed by something outside the book and recorded in a register, a regulatory or tax constraint rather than a stated preference, the book becomes a procurement auction that pays the investors it moves: nobody who keeps his place pays anything, no investor withdraws, and the cash the book needs is bounded whatever veto its investors retain. The capital requirement of a retained veto arises from the interaction

of private preferences, truthful allocation and a retained right to refuse, and an institution that removes any one of the three escapes it.

One more mechanism deserves mention because it is the one an economist reaches for first. A mechanism that pays each investor the expected externality of his report [d'Aspremont and Gérard-Varet, 1979] balances its budget and keeps the shortfall bounded under truthful reports, but a veto undoes it. The payment rewards a claimed willingness to be moved, and an investor who claims zero cost can collect it when his favored project is not crowded and refuse the move when it is. The veto turns a payment for flexibility into an option, and the investors with a reason to exercise it, having reported the lowest costs, would be the first moved. The principle the three results share is that at the moment an investor can still refuse, his net burden must not move with the common imbalance and must not rest on a claim that his refusal lets him withdraw.

Four rules for anyone who arranges such a book follow. Bind participation before the common imbalance is known. Price the crowded position rather than paying for flexibility. Return the premium within the book, leaving out each investor's own report where exact incentives are required. Reserve replacement capital to a quantile of the shortfall, which is of order $\sqrt{n}$ under a retained veto and a few invitations under binding enrollment. Read together with the diversification result, the theory gives an investor two forms of control, each with a financing cost: a narrow mandate is accommodated by flexible capital, of order $\sqrt{n}$ when mandates are idiosyncratic, and a late decision by replacement capital, of order $\sqrt{n}$ times a quantile of the common imbalance under a veto and of order one when the burden is accepted in advance and the transfers are returned.

6. What common exposure buys

If proportional exposure is not required for formation, the fund's defenders are entitled to ask what it is for, and the answer is specific. It economizes on two things, the records that individual exposure requires and the incentive pay that a manager serving many separate accounts requires. Both are real, both can be large, and neither is a condition of committing together.

Take incentive pay first. A manager who investigates before investing earns his fee only if his pay depends on whether he investigated, and the records available to write that pay on determine how much it costs. When records are complete, so that each account's own history reveals what the manager did to it, pay written on each account's own record can make investigation the manager's best choice whenever any pay can; joint control over his actions is not needed, although it can make the same incentive cheaper. The cheapening is where pooling earns its keep. Suppose a manager's research improves twenty independent bets, each held in a different account, from a coin flip to fifty-five percent. Pay on each account's own record must reward the manager for each bet separately, and the least expected pay that induces investigation comes to eleven times the research cost however many accounts there are. Pay written on the pooled record, on the number of bets that came right across all twenty, can be concentrated on outcomes that only an investigating manager produces, and the least expected pay falls to about one point seventeen times the research cost. The difference is a factor of nine, and it decides whether investors who together value the research at twice its cost can afford it at all. The advantage belongs to pooling across bets rather than across investors, since an account holding an equal slice of every bet sees the count of right bets in its own profit and loss; it is Diamond's (1984) argument for delegated monitoring in a new place. Paying a manager together is, moreover, itself a formation problem,

in which refusal by everyone is an equilibrium exactly when no single account values the research above the least pay that makes the manager do it.

Records are the second economy. A book that holds many assets for many investors and offers each a combination of a few exposure profiles keeps a record for each interest and each position; a pooled fund is the case of a single profile. For investors who differ in risk tolerance and in the exposures their outside wealth already carries, the loss from restricting everyone to a few common profiles is the tail of a spectrum, the part of the dispersion of investors' optimal holdings that the profiles cannot span, and the mutual fund theorem is the case in which that tail is zero. The number of profiles a book should offer rises as the cost of a record falls and falls as the number of classes of investors who cannot share a vehicle rises. The losses from a common profile fall on the investors whose circumstances differ most from the population's: the founder whose wealth sits in one sector, the institution whose liabilities carry a particular exposure, the family whose holdings already supply what the fund would add. Technology that lowers the cost of forming vehicles, documenting interests and reporting positions moves the book toward an exposure of each investor's own.

A third result is less obvious and in practice often decisive. Investors in a shared vehicle do not pay one another; their economics are set by the program's terms, which most-favored-nation clauses extend from one investor to every investor entitled to his terms. Such a clause is a commitment in the economist's sense, it makes a selective concession expensive, and it bites at exactly the moment the formation problem operates, when a concession is needed to secure a late, pivotal commitment and must then be extended to everyone who signed earlier. Securing a pivotal investor whose shortfall is some amount costs, under rate matching, that amount multiplied by the ratio of all the commitments entitled to his terms to his own. A pivotal investor committing $20 million who requires a quarter-point fee reduction costs $250,000 a year when an $80 million investor is entitled to the same reduction, and $50,000 when the larger investor's existing fee is already below the pivotal investor's new one. The entitled investors could waive their claims, but renegotiating terms held in common is itself a formation problem, each waiver a contribution refunded on failure and universal refusal an equilibrium whenever no single waiver covers the deficit. Individual terms are constrained in the same way as individual commitments, by the need for a collective agreement; individual exposure is constrained only by the cost of records and incentive pay. The arrangement that individualizes exposure while holding terms common has chosen the cheaper of the two things to individualize, and equal-treatment clauses, not records, are the expensive residue of the fund.

7. Better separate accounts, and the investments that disappear

Commitment support and discretion make an investment form when its investors want it. They cannot make an investor commit to an investment he would rather not join, and the technology that cheapens separate accounts raises the alternative against which each investor decides. The paper's third set of results asks what happens to shared investments as that alternative improves, and the answer is a theorem about the range of equilibria rather than about any one of them.

The worst equilibrium of a formation game is the one in which the investment fails and every investor keeps his outside value; the best is the formed equilibrium that leaves the most surplus on the table. Between them lies what the paper calls the selection gap, surplus that better coordination could recover without changing anyone's terms. Under a property the paper calls completion, which holds whenever any allocation every investor would accept can be improved into an equilibrium, raising one investor's

outside value by a given amount raises the worst equilibrium by exactly that amount, since the floor is the sum of outside values, and changes the best equilibrium by an amount between that improvement less the selection gap and the improvement itself. The range above the floor can never widen; it can only compress. Completion holds for threshold programs with scale economies, where an additional member weakly improves every existing member's position, and for divisible contributions from indispensable investors; it fails under congestion, where an entrant can harm incumbents.

The consequence is that an investment which still creates surplus can disappear. In the paper's five-investor economy, outside values rise with a common technological improvement while the program's terms stay fixed; the best equilibrium holds until the first member leaves, falls discontinuously when he does, and the selection gap closes well before the program ceases to create surplus, after which the entire value of coordination is distribution gap, surplus recoverable only by new terms. In a four-investor economy with divisible capital, a fall in the cost of investing separately from ten to seven lowers best-equilibrium welfare from $76,000 to $22,000 although the shared program, which still needs the capital of the one investor whose separate investment has become attractive, would continue to produce $76,000. He declines because proportional cost sharing would pay him less than he now earns alone. The investment did not become worse. The investor who was pivotal acquired a better alternative, the fixed terms could not accommodate him, and nobody could pay him to stay because the equal-treatment clauses of the previous section made the concession cost more than the rent available to fund it.

Two corollaries follow for the instruments of the first part. An escrow is worth least where better separate accounts have already compressed the range above the floor, because failure payments change no incentive at a formed profile and can recover at most the selection gap; where compression has removed every formed equilibrium an escrow recovers nothing, since what remains is distribution gap that only new terms can close. And search is the mirror image of a better outside option. A larger pool of investors that a program can reach leaves the floor where it is and weakly raises the best equilibrium, by exactly the gain in the best acceptable surplus; a better outside option shrinks the family of willing memberships, and a wider search enlarges the set from which they are drawn. In the five-investor economy, reaching one more investor whose alternative is no better than the cheapest member's keeps the program formed over a range of outside improvements in which the original pool had already failed. This is why search, rather than computation, is the margin on which coordination at fixed terms can raise what is attainable: a coordinator who computes the best allocation for a given pool can do no better than its best equilibrium, while one who finds the investors who belong in the pool can.

8. Why now: agents and the cost of collective search

The paper distinguishes two kinds of term that an investment can carry, and the distinction makes the technological backdrop legible. An own term can be written and enforced account by account: a management fee on an account's assets, a performance fee on its returns, a list of permitted positions, a record of its holdings. A joint term depends on what other investors do: a commitment that binds only if enough others commit, a payment conditioned on the failure of the whole, a price on a crowded place, a manager's pay written on a pooled record. Technology lowers the cost of both, and the two reductions act on the formation game at opposite ends.

Lowering the cost of own terms, which is what software that reads a mandate, reconciles an amendment and keeps an attributable record of holdings does, raises every investor's outside value, because a separate account tailored to his taxes, liabilities and existing holdings becomes cheap. That is the

force of the previous section. It raises the floor of the equilibrium range, compresses what lies above it, shrinks what assurance can recover, and can remove from equilibrium an investment that remains jointly valuable; it also completes records, which lets delegation be arranged account by account and leaves to joint terms the roles of lowering its cost and of committing accounts together. Lowering the cost of joint terms, which is what agents that can find eligible and authorized investors, establish that each may commit, hold commitments conditionally, compute what an escrow would owe in every profile and pay a manager on a joint record do, acts on the ceiling. It enlarges what joint terms can achieve, widens the pool a program can reach, finds complementary permissions and compatible substitutes, and lowers the expense of standing ready.

Collective search is the joint-term task that agents change most, and the one on which the two parts of the theory agree least. On the side of commitment, search adds willing investors and widens permissions, which lowers both the shortfall of principal and the commitment support a book needs, and it lowers the expense of standing ready, which lowers the support again; search also supplies replacement capital, since capital identified before the imbalance is known as willing to take a vacated place on stated terms is exactly what a book whose investors keep a veto must hold. On the side of discretion, search has an edge the enthusiast will not have noticed. A technology that leads many investors to the same recommendation, because they consult the same agent, the same adviser or the same ranking, creates the common component of mandates that makes discretion expensive: the required flexible share stops shrinking with the population and converges to a fixed fraction, and the cost of discretion per investor stops vanishing. Whether a search technology diversifies a population's mandates or concentrates them is therefore the question that decides its effect, and the same agents can do either.

Incidence completes the picture. A technology that lowers both costs at once leaves an investor more or less willing according to where its savings land relative to the slack he already has, the saving that reaches him outside the program net of the saving that reaches him inside it. A common percentage improvement does not produce common dollar savings, so a uniform improvement can have a concentrated effect on participation, removing one pivotal investor while leaving the rest untouched. The paper's companion measurement, an experiment on AI-assisted investment administration, measures one component of that comparison and reports it as a datum rather than a verdict.

The secular case does not rest on agents making funds obsolete. It rests on the two costs on which the fund's bundle stood falling together. The fall in the first erodes the separation-theorem reason for pooling, the fall in the second lowers the price of every joint term the theory has priced, and the arrangement that captures both, individual exposure with shared commitment support, compatible principal, priced discretion and early binding, can now be specified rule by rule. What such an arrangement needs in order to be more than a design proposal is the thing the theory keeps returning to: a record of who is authorized, eligible, available on time and unencumbered for a given closing, kept with the same discipline before commitment, at closing, during ownership and at any transfer. Agents lower the cost of keeping that record. They do not remove the need for the capacity, the binding moment and the priced discretion that the record makes it possible to organize.

9. Fund or book: what the theory predicts

A theory that decomposes the fund ought to say when the fund is the right answer, and this one does, along a line it locates in the composition of mandates. Where investors' mandates differ idiosyncratically, the book's expected cost of discretion stays bounded while a pooled fund's grows in proportion to the

population, so that the book's advantage per investor does not vanish as the book grows, and the theory predicts the book. Where a common component moves a share of mandates together, the flexible share the book needs converges to a positive limit and its cost per investor to a positive limit, so that both arrangements' costs grow in proportion to the population, the book's smaller by a factor that depends on how dispersed the costs of being moved are. The cost of the records and incentive pay that common exposure economizes is proportional to the population as well, and whether the book's smaller cost of discretion pays for the records it keeps separate becomes a comparison of two quantities of the same order, which the paper does not make. A manager whose judgment moves all of his investors' mandates together is the limiting case of a common component, and the fund is the arrangement the theory expects around him.

An investor who chooses a manager for his judgment is choosing to let one view move his mandate, and a population of such investors is a population with a large common component; for them the book's economies are small and the fund's are real, and nothing in the theory says they should want anything else. An investor whose constraints are his own, the exclusion his charter imposes, the sector his existing wealth already covers, the tax class that requires a vehicle of its own, is an investor whose mandate differs idiosyncratically, and the fund serves him at a cost the theory can measure and the book avoids. The two populations coexist within every large investor base. The institution the theory suggests is not a replacement for the fund but a way of letting the second population take part in the relationships the first population sustains, with their differences accommodated by priced discretion rather than by excuse rights and side letters negotiated one at a time. Excuse rights, sector exclusions in side letters and co-investment vehicles are partial books already; the comparison that matters in practice is between those arrangements and a book with priced places, at equal costs of administration and records.

The model does not derive secondary-market prices, so nothing in it says what a transferable interest in a book is worth or that transferability lowers the cost of formation. It does not treat a fund's other functions, liquidity insurance, tax and legal structure, limited liability, and it does not say the book is cheaper once those are counted. Its benchmark for the central result has two equal projects, a single closing, equal commitments and every investor flexible; the constants depend on that symmetry, the orders of magnitude do not, and successive closings that draw on the same investors are left open. And it says that the capacity a closing needs is of a certain order, not who should hold it, at what price, or how a market in contingent capacity would be organized.

10. What would test it

The contracts of special-purpose acquisition companies contain the model's objects in a standardized form, which makes them the natural first test. A SPAC is an anchor with a refund attached: its offering is underwritten on a firm-commitment basis, so that the underwriter forms it alone, and its trust account returns the proceeds with interest if no combination is completed, which makes the subscription credible without paying anyone for the failure of others. Its minimum cash condition is a formation threshold, sponsor forfeitures and non-redemption incentives are transfers, and amendments and waivers change terms. The redemption right is the retained veto of the central result. Every public holder decides after the same target and the same terms are announced, so that redemptions should fluctuate across deals on the scale of their mean rather than concentrate, and the capital a sponsor must arrange to close should

scale with a high quantile of that distribution rather than with its average, which is what forward-purchase agreements and the PIPEs signed with the merger agreement supply, while non-redemption agreements bind the participation of the holders who sign them. The data behave as the theory says they should: redemptions averaged more than eighty percent of trust capital for deals closed in 2022 and ninety-five percent in March 2023, the share of merger agreements amended rose from thirty-seven to sixty percent, and sponsors forfeited founder shares in eighty-five percent of 2022 deals against fifty-six percent the year before [Barlow et al., 2023]; the costs of formation fall on the shareholders who do not redeem [Klausner, Ohlrogge and Ruan, 2022]. The theory makes four predictions for the cross-section: that amendments, waivers and transfers concentrate in deals whose minimum cash condition binds; that transfers from sponsors are largest where the marginal holder's slack is smallest; that success-contingent incentives change who remains without removing a deal's dependence on coordinated redemption decisions; and that the replacement capital a sponsor arranges, and what it costs, scales with a high quantile of the cohort's distribution of redemptions rather than with its mean.

Private markets lack standardized contracts, so a test there must be prospective. It should classify each opportunity's constraints before outcomes are known, retain every considered opportunity including refusals and failed closings, and record for each decision date the outside route then available, the requested allocation, net participation terms, authority, existing obligations and funding state, distinguishing refusal from diligence, documentation, consent, funding and settlement as reasons for failure. The central empirical problem is to separate two reasons for declining that look alike in transaction data: an investor may decline because his outside route improved, which is the compression mechanism, or because the shared investment became less attractive, which is not. A persuasive design identifies changes that cheapen separate implementation, measures their effect on the shared route as well, and asks whether formation failures concentrate in cohorts with little participation slack, little headroom between the capital investors would accept and the threshold, and a high cost of retaining the pivotal investor relative to the rent available for concessions, which requires measuring the commitments entitled to his terms and not merely his size.

The costs the paper prices are borne by someone, and in an equilibrium in which investors supply contingent capacity at a price they should appear in required returns at formation, as the cost of trading appears in a security's expected return [Amihud and Mendelson, 1986]. The paper does not derive that premium, which would require a market in contingent capacity that it does not model, and private-market returns are measured too coarsely to identify it across investments. The prediction it does support concerns selection rather than price: at a given expected return, the investments that form are those whose formation costs are smallest relative to the surplus they create, so that compression, cheaper standby and earlier binding should change the composition of what forms before they change what it pays, and a test on returns should look for the change in composition first.

11. Objections, and what the paper concedes

Funds already do this. Excuse rights, side letters that exclude sectors and co-investment vehicles alongside the fund already let some investors decline particular investments and others take more of them. The paper agrees, and calls them partial books. Its claim is that they individualize by negotiation, one investor at a time, with the equal-treatment clauses of Section 6 taxing every concession, whereas a book with priced places individualizes by a mechanism whose cost is bounded and whose incentives

are dominant. The comparison that matters is between the two at equal costs of administration, and the paper proposes the measurement rather than asserting the result.

The relationship is the asset. An issuer who wants a particular lead investor is buying judgment, reputation and continuing involvement, and an allocator may deliberately want one manager's discretion across a portfolio. The prediction of Section 9 accommodates this directly: a population whose mandates one judgment moves together is the case in which the theory expects the fund, and nothing in it suggests making those relationships harder to maintain. What it adds is that the investors who do not share that judgment are served by the fund at a measurable cost, and that an arrangement which preserves the manager's judgment while accommodating their differences by priced discretion is cheaper for them and no worse for the manager, provided his capital is counted only when it is authorized, eligible, available and unencumbered.

The escrow is not observed. The escrow is a design proposal. Practice uses orders, leads, refunds and anchors instead. Two of its three explanations come from the theory itself, the third from reputation, which the model does not contain. The paper does not claim that credible contingent capacity is as cheap as its carry suggests; it identifies that as the friction a measurement would have to resolve.

The benchmark is stylized. The central result assumes two equal projects, every investor flexible, independent private margins and costs, and a single closing; the constants depend on the symmetry and the orders of magnitude do not. Unequal commitments change the scale of the result by an effective population size, so that concentrated commitments or accounts that share a decision maker make replacement capital a larger share of the book; a share of narrow investors changes the constant; successive closings that draw on the same investors, and mechanisms that cancel when a finite reserve is exhausted so that the cancellation rule itself enters the incentive problem, are open. The lower bound on capacity is proved for mechanisms with dominant-strategy incentives; whether some Bayesian mechanism keeps the shortfall bounded under a retained veto is not known.

Transfers and secondary prices. The model includes neither transfers nor secondary prices. It does not establish that transferable interests lower the cost of formation or that secondary discounts measure the price of a retained veto. A register that records the transfer rights actually obtained is useful for formation whether or not any transfer occurs; what transferability does to willingness, and to the expense of standing ready, needs a model the paper does not supply.

Welfare and rents. The paper's welfare is a sum of certainty-equivalent dollars in which transfers net out and real expenses do not; supplier rents, taxes and external effects need their own accounting. A promisor with capacity just above the required bound could, under the model's contracting assumptions, collect nearly the whole surplus of formation; whether any observed charge is such a rent is an empirical question about competition among promisors that the model does not answer.

12. The frontier, decided

Investments that need several investors to commit are formed by holding some things in common, and they need not be formed by holding everything in common. Under simultaneous commitments they need support, whose cheapest robust form is a promise to pay in a state in which the investment does not exist, and an order substitutes for it among investors whose earlier commitments bind and who trust one another's future rationality. A set of projects financed together needs principal whose permissions cover it, measured by compatibility rather than by aggregate

capital, and discretion to absorb the mismatch between what investors permit and what projects require, whose amount has the structure of diversification and whose expected cost stays bounded when many investors grant priced permission and their mandates differ. It needs decisions surrendered at defined times: investors who keep the right to decide after the terms are known must be backed by replacement capital sized to the tail of a shock common to all of them, while investors who accept the burden before that shock is known, with the transfers it generates returned within the book, leave a shortfall that does not grow with their number. Exposure enters only through who is willing to commit. Proportional exposure is not required, and what it buys is an economy in records and incentive pay whose value rises with the dispersion of investors' optimal holdings and falls with the cost of a record.

The forces that technology sets in motion act on these quantities in different directions, which is why the picture is a frontier and not a trend. Cheaper separate accounts raise the outside values against which investors decide, compress the range above the floor, shrink what assurance can recover and can remove from equilibrium an investment that remains jointly valuable, while equal-treatment clauses make the distributive repair expensive and cheaper records make individual exposure affordable. Cheaper collective search widens the pool from which shared investments draw their members, finds complementary permissions and compatible substitutes and lowers the expense of standing ready, but it raises the cost of discretion when it leads many investors to the same recommendation. Whether cheaper search and organization expand the set of investments that can be executed, or mainly intensify competition for investments that would have formed anyway, is the empirical question on which everything institutional turns.

The surprise in looking at portfolio theory through this lens is not that the fund is wrong but that it was never one thing. It was five, priced together because the cost of pricing them apart was prohibitive, and the technology now lowering that cost is the same technology lowering the cost of the search on which every formation depends. The frontier that investors attain is decided where these forces meet: at the investor whose commitment is pivotal, in the composition of the population's mandates, and in the moment at which their decisions become binding.

References

Amihud, Yakov, and Haim Mendelson. 1986. Asset Pricing and the Bid-Ask Spread. *Journal of Financial Economics* 17(2), 223–249.

Bailey, Martin J. 1997. The Demand Revealing Process: To Distribute the Surplus. *Public Choice* 91(2), 107–126.

Barlow, Christopher, Michael Chitwood, Howard Ellin, Raquel Fox, Michelle Gasaway, Edward Micheletti, and Gregg Noel. 2023. De-SPAC Transaction Trends in 2023. Skadden, Arps, Slate, Meagher & Flom LLP.

d'Aspremont, Claude, and Louis-André Gérard-Varet. 1979. Incentives and Incomplete Information. *Journal of Public Economics* 11(1), 25–45.

Diamond, Douglas W. 1984. Financial Intermediation and Delegated Monitoring. *Review of Economic Studies* 51(3), 393–414.

Gale, David. 1957. A Theorem on Flows in Networks. *Pacific Journal of Mathematics* 7(2), 1073–1082.

Green, Jerry, and Jean-Jacques Laffont. 1977. Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods. *Econometrica* 45(2), 427–438.

Hall, Philip. 1935. On Representatives of Subsets. *Journal of the London Mathematical Society* 10(1), 26–30.

Klausner, Michael, Michael Ohlrogge, and Emily Ruan. 2022. A Sober Look at SPACs. *Yale Journal on Regulation* 39, 228.

Komkov, John. 2026. The Formation Game: Commitment and the Investment Frontier. Working paper, October 2026.

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

Tobin, James. 1958. Liquidity Preference as Behavior Towards Risk. *Review of Economic Studies* 25(2), 65–86.