A Unified Theory of Delegated Capital Management*

Jonathan B. Berk
Peter M. DeMarzo

Stanford University and NBER

August 26, 2026

Abstract

Alternative funds differ sharply from mutual funds: institutional investors dominate, fees are performance-based and investor-specific, access is rationed, and some investors earn positive alpha. We show that these differences arise endogenously in competitive capital markets when managerial skill is uncertain. With liquid, transparent assets, investors optimally learn from performance, capital adjusts through flows, investors earn zero alpha, and the mutual fund contract is optimal. With illiquid, opaque assets, costly ex ante due diligence becomes efficient, with institutions as information producers. Delegation creates a free-rider problem, requiring preferential terms for informed investors. The standard alternative-asset contract, featuring management fees, carried interest, preferential terms, and capped fundraising, implements the first best. Finally, delegation is necessary. Bilateral contracting avoids free riding but is inefficient because informed capital must both reveal information and determine fund size. Delegation separates these roles, allowing informed capital to reveal information while uninformed capital achieves efficient scale.

*We thank Pete Briger, Willie Fuchs, Barney Hartman-Glaser, Vincent Glode, Ben Hebert, Ana Marshall, Marcus Opp, Adriano Rampini, Myron Scholes, Andy Skryzpacz, Paul Oyer, Vish Viswanathan, Michael Weisbach and seminar participants at the Federal Reserve Bank of New York, FT Webinar, the Jackson Hole Finance Group and the Utah Winter Finance Conference for comments and insights.

Capital markets have seen a seismic shift in the last 70 years. Rather than investing directly, most market participants invest indirectly through intermediaries known as money managers. In broad terms, money managers are divided into two sectors, publicly available investments, mainly in the form of mutual funds, and privately available investments known as "alternatives." Although, superficially, managers in both sectors appear to do the same thing — manage money on behalf of investors — important differences between the two sectors exist, which is puzzling given the apparent similarities between them.

In the mutual fund sector, there is little evidence that investors can systematically earn positive risk-adjusted returns by selecting managers. In alternatives, by contrast, some investors appear to earn superior returns, particularly where access and due diligence are important. Why do managers appear to leave rents with investors rather than extracting them through higher fees or larger fund size? A successful manager should be able to capture the value of their skill by raising fees or expanding the fund until diminishing returns eliminate any investor alpha.

Other puzzling differences exist. Mutual funds draw heavily on retail capital, offer largely standardized terms, and compensate managers through fees based on assets under management. Alternative funds are dominated by institutional investors, rely heavily on performance-based compensation, offer different terms to different investors, and ration access even to existing investors. Taken together, these facts suggest that the information and contracting frictions in alternatives differ fundamentally from those in the mutual fund sector.

We show that these facts have a common explanation within the same competitive-market paradigm that explains the mutual fund sector (Berk and Green (2004)). As in mutual funds, the informational problem is that the quality and scalability of a manager's investment strategy are initially uncertain to everyone, including the manager. A sophisticated investor can resolve this uncertainty upfront through costly due diligence. Whether doing so is valuable depends on the nature of the underlying assets. When assets are liquid and performance is informative, investors can learn over time and capital can respond through flows. When assets are illiquid and opaque, waiting is costly and ex ante due diligence becomes valuable. We derive the optimal contract in this environment and show that it reproduces the central features of private-market funds, including upfront due diligence, performance fees, preferential investor terms, positive ex post alpha for informed investors, and binding limits on fundraising. Surprisingly, observed contracts are optimal and implement the first-best outcome. Moreover, delegation itself is essential. Directly hiring a manager, or outsourcing information production, would distort the level of invested capital and thus would be inefficient.

To understand what drives these results, begin with the logic of Berk and Green (2004). Neither the manager nor investors initially know the quality and scalability of the manager's

investment strategy (hereafter, managerial "skill"). When the underlying assets are liquid and realized performance is informative they need not resolve this uncertainty before investing. Investors can learn from performance over time, while capital flows repeatedly resize the fund as beliefs change. Any initial misallocation is therefore corrected relatively quickly, making costly upfront due diligence unnecessary. Illiquidity and opacity weaken different parts of this mechanism. Illiquidity prevents investors from readily adjusting capital, so the consequences of an initial mistake persist. Opacity makes realized performance a noisier signal of managerial skill, so that mistakes are discovered more slowly. Both frictions therefore increase the value of learning before capital is committed. When this value exceeds the cost of acquiring information, ex ante due diligence is efficient.

Due diligence is costly. Because this cost is incurred upfront, it is economical only when the potential capital commitment is sufficiently large. This creates a natural role for institutional investors, rather than dispersed retail investors. It also changes the interpretation of investment returns. An investor who bears the cost of due diligence must expect to earn positive alpha conditional on investing; otherwise, no investor would acquire the information in the first place. In a competitive equilibrium, this alpha compensates the investor for the cost of becoming informed. In our homogeneous investor benchmark, investors earn no rents net of information costs; more generally, investors with a persistent advantage in acquiring information will earn persistent rents.

However, when the manager raises capital from multiple investors, but only one investor incurs the cost of due diligence, a free-rider problem emerges. Uninformed investors can wait for another investor to acquire information and then invest without bearing the cost themselves. If they can invest on the same terms, no investor will be willing to conduct due diligence in the first place. Thus, although information acquisition is socially valuable, it cannot be sustained in a delegated fund in which all investors are treated identically. The informed investor must receive preferential terms.

The most direct way to eliminate this free-rider problem is to dispense with a multi-investor fund. In our setting, the informed investor is not capital constrained, so she could contract bilaterally with the manager and supply all of the capital herself. At first glance, bilateral contracting therefore appears to sustain due diligence while avoiding the problems created by delegation.

We show, in fact, that bilateral contracting is *not* an efficient solution. The informed investor's capital commitment must both reveal what she has learned (which determines the fund's efficient size) and supply the required capital to achieve efficiency. If the manager is paid more when the informed investor reports a more able manager, she has an incentive to understate his ability. Hence, any bilateral contract that gives the manager positive rents

will distort investment: some managers who should be funded are not, and every fund that is launched is too small.

Delegation solves this problem by introducing uninformed capital. To achieve efficiency, the informed investor's capital commitment must reveal what she has learned. Using this information, the manager can raise additional capital from uninformed investors. Thus, delegation separates the capital used to convey information from the total capital invested. Although uninformed capital is not needed to finance the fund, it is essential for efficiency: this separation allows the manager to capture the surplus without distorting investment, thereby implementing the first-best outcome.

The optimal contract has three components. First, the informed investor receives preferential terms. The management fee is set so that the fund will launch if and only if funding the manager is efficient (i.e. it has positive NPV after operating costs), while uninformed investors pay a higher fee. This fee differential rewards the informed investor for costly diligence and prevents investors from instead waiting to invest after someone else has produced the information. Second, the carry-based performance fee allows the manager to capture the remaining rents without distorting the launch decision. Finally, the contract imposes a binding cap on fundraising, set to the efficient level of capital. Without the cap, the manager would have an incentive to raise excessive additional capital and dilute the informed investor's return. Together, the fee differential prevents investors from free riding on information production, while the cap commits the manager not to undermine the informed investor's incentive by expanding the fund after her information has been revealed.

The optimal mechanism therefore resembles an anchor-investor arrangement: the anchor is rewarded for costly diligence through preferential terms, her investment conveys the resulting information to other investors, and the fund-size cap preserves her incentive to become informed. This contract differs from the mutual fund contract because a fee based only on assets under management cannot simultaneously induce efficient information production, implement efficient investment, and extract the rents.

Although some investors acquire information, managers ultimately bear the cost by giving them preferential terms. Informed investors earn a positive ex post alpha that compensates them for the cost of investigating and validating managerial skill. Other investors pay higher fees and earn zero net alpha in equilibrium. Hence, consistent with perfectly competitive markets, managers extract all rents once information costs are taken into account. That said, full extraction of investor rents is not essential for implementation; what is essential is a sufficient fee differential to preserve the incentive to become informed.

It is important to emphasize that, in our framework, the excess returns earned by informed investors do not arise from a premium for holding illiquid assets. Uninformed investors also

hold illiquid assets but earn zero alpha. Moreover, if managerial skill were directly observable, all investors would earn zero net alpha even though capital would still have to be committed for an extended period. Instead, the source of positive alpha in our model is the limited ability of capital markets to endogenously communicate information when underlying assets are illiquid and opaque. That is not to say that a liquidity premium does not exist in practice, only that it is not needed to explain the contractual and return patterns we study in this article.

Delegation succeeds in this setting because it separates information provision from the total capital allocation. The informed investor must put her own capital at risk, so that she benefits from discovering a poor manager and bears the cost of failing to investigate. An independent certifier cannot replicate this arrangement. A certifier paid only when the fund launches cannot be rewarded for discovering that it should not launch, and paying her upfront would remove the incentive to collect the information in the first place.

The assumption that the manager does not already know the fund's productivity is essential to delivering the result that delegation is necessary. We interpret managerial ability broadly as the profitability and scalable capacity of a particular investment strategy, rather than simply the manager's knowledge of his own intelligence. A manager may understand his capabilities yet remain uncertain about how they will translate into abnormal returns in a new fund, with a new team, or at a different scaleInstitutional investors that evaluate many managers may be better positioned to make this assessment. Moreover, if the manager already knew the fund's productivity, the distinctive delegated structure we study would no longer be necessary. If the manager could credibly communicate this information, outside due diligence would serve no purpose. If he could not, a single investor could verify it and then contract bilaterally with the manager without distorting investment. Thus, the prevalence of delegated funds in which sophisticated investors conduct due diligence provides indirect support for the informational premise of the model.

The paper is organized as follows. We build the theory in four steps. First, in Section 2, we abstract from delegation and analyze efficient information acquisition and capital allocation with a single decision maker. We explain why the liquidity and opacity of the underlying assets determine whether it is optimal to acquire information in advance or instead to learn gradually from performance. We then turn, in Section 3, to the delegated setting and derive the optimal contract using a linear carry schedule that does not impose limited liability. There, we characterize the roles of fees, carry, and uninformed capital, and show how differential investor terms and a binding fundraising cap implement the first-best allocation. In Section 4, we then impose the limited liability constraint by replacing linear carry with the option-based performance contract observed in practice. We show that the central economics

are preserved and the first-best allocation is still attained. In Section 5 we show why bilateral contracting fails to achieve first best and that delegation is necessary. We also explain why external certification fails in this setting. Section 6 discusses the broader implications of the model, including investor heterogeneity, the role of uninformed investors, and the open questions the theory leaves unresolved. We also consider the implications of assuming that the manager has an informational advantage assessing his own skill level. Section 7 concludes.

1 Background

One of the most influential studies of performance in the alternatives industry, Kaplan and Schoar (2005), documents substantial heterogeneity and persistence in private equity performance, together with evidence of outperformance in some specifications. Interpreting these findings is complicated by important measurement issues, because much of the early evidence relies on voluntarily reported data rather than the more comprehensive disclosures available in the mutual fund sector. Even so, Kaplan and Sensoy (2015) survey the subsequent literature using alternative data sources and conclude that the broad finding of persistence is robust, although the strength of that evidence appears to vary across segments of the industry.[1] Kosowski, Naik and Teo (2007) and Fung, Hsieh, Naik and Ramadorai (2008) document evidence of outperformance in hedge funds although the same reporting concerns arise in that sector.[2] Somewhat surprisingly in light of these concerns, Barth, Joenväärä, Kauppila and Wermers (2023) find that hedge funds that do not report to commercial databases outperform those that do, suggesting that the direction of voluntary-reporting bias is itself ambiguous.

Sorensen, Wang and Yang (2014) model key contractual features of private equity, including management fees, carried interest, waterfall provisions, and the illiquidity borne by investors. Calibrating their model to the data, they show that conventional performance measures can make PE look more attractive than it is from the LP's utility perspective, and substantial GP alpha can be required merely for LPs to break even after accounting for risk, illiquidity, management fees, and carried interest. They therefore caution against interpreting conventional measures of outperformance as abnormal risk-adjusted returns.

Several recent papers provide direct evidence for the information-production and skill-allocation mechanisms in our model. Brown, Gredil and Kantak (2023) provide direct evidence that institutional allocators acquire costly private information about manager skill

[1] Korteweg (2023) provides a comprehensive review of both the empirical findings and the conceptual difficulties in measuring performance in private markets.

[2] See Jagannathan, Malakhov and Novikov (2010), Agarwal, Fos and Jiang (2013) and Edelman, Fung, Hsieh and Naik (2012).

and capacity through due diligence. Using detailed records of an allocator's interactions with 860 hedge fund managers, they show that this information helps the allocator identify skilled managers before other investors do. Relatedly, Goyal, Wahal and Yavuz (2026) show that LPs frequently invest with first-time GPs, for whom a fund-level performance history is limited or absent. The tendency to select first-time managers is particularly strong among larger LPs and those with formal investment committees, a pattern consistent with a role for institutional capacity in evaluating managers ex ante. Fu and Taylor (2025) show that VC investors undertake costly due diligence when evaluating prospective portfolio companies. As the cost of diligence rises, diligence declines and outcomes become more uncertain. Finally, Braun, Dorau, Jenkinson and Urban (2026) document that private equity (PE) managers have considerable skill, that skill is persistent, and that higher-skill managers are entrusted with more capital, a central implication of our model.

One of the first articles to systematically analyze compensation contracts in the venture capital industry is Gompers and Lerner (1999), which also notes that similar contractual terms are common across private equity funds more broadly. In these settings, general partners are typically compensated through both a management fee based on assets under management and carried interest. Gompers and Lerner (1999) develops learning- and signaling-based interpretations of these contracts, taking the basic contractual form as given. In their learning model, investors gradually learn about the general partner's ability over time, so the contract reflects the manager's evolving reputation. That is, information is revealed gradually through realized performance. Metrick and Yasuda (2010) provide a detailed empirical analysis of the economics of these contracts, documenting the importance of both management fees and carried interest and quantifying how their value varies with fund characteristics and performance. This approach leaves open the central issue we address here: why this particular contractual form is optimal, and why it differs systematically from the contracts used in liquid public markets. In our model, the illiquidity and opacity of the underlying investments make ex ante due diligence more important and thereby alter the optimal contract relative to the mutual fund sector.

Even setting aside the direct evidence on fund-level performance, there is indirect evidence that some investors systematically earn higher returns in alternatives. Lerner, Schoar and Wongsunwai (2007) show that certain investors, particularly endowments, earned superior returns in private equity, suggesting that they were better able to identify and access high-quality managers. Subsequent work qualifies this pattern, indicating that such outperformance was concentrated in particular periods and segments of the market, rather than being uniformly present across all private-market investments.[3] More directly, Cavagnaro,

[3]See Lerner, Schoar and Wang (2008), Brown, Garlappi and Tiu (2010).

Sensoy, Wang and Weisbach (2019) estimate persistent differences in investors' ability to select private-market funds and find economically meaningful heterogeneity in investor skill, particularly in venture capital and in the earlier part of their sample. One prominent explanation, emphasized by David Swensen, is that some institutions possess a persistent advantage in the human capital devoted to manager selection and due diligence. Recent evidence is consistent with this view. Binfarè, Brown, Harris and Lundblad (2023) find that investor human capital and governance are associated with better manager selection and investment performance. Similarly, Martinez and Qian (2026) find that asset owners advised by more selective investment consultants subsequently earn higher returns, a pattern they interpret as consistent with more selective consultants possessing superior information or skill in manager selection. Mittal (2026) shows that more underfunded public pensions disproportionately allocate capital to lower-quality managers and subsequently earn lower returns. Taken together, this evidence is consistent with the idea that, in alternative assets, some investors, or the specialized intermediaries they employ, have a comparative advantage in acquiring and acting on information about manager quality.

As Kaplan and Sensoy (2015) emphasize, persistence in alternatives is puzzling relative to the mutual fund benchmark, where competition implies that investors should not earn persistent abnormal returns. In a competitive delegated-management setting, successful managers should be able to extract those rents through higher fees or larger fund size, so the continued presence of abnormal returns to investors requires explanation. More recent evidence suggests that persistence remains more pronounced in venture capital than in buyout funds.[4] Hochberg, Ljungqvist and Vissing-Jørgensen (2014) propose that persistence may arise because incumbent investors acquire an informational advantage, giving rise to a hold-up problem. That mechanism may help explain persistence within existing relationships, but it does not by itself explain how a new fund can initially raise capital on terms that leave rents to some investors in a perfectly competitive capital market.[5]

Begenau and Siriwardane (2024) show that investors in alternatives do not all earn the same returns: fees differ substantially across investors, and these differences imply persistent heterogeneity in net performance. This contrasts sharply with the mutual fund sector, where

[4]Subsequent work continues to find evidence of persistence in venture capital funds, but suggests that persistence in buyout funds has weakened; see Harris, Jenkinson, Kaplan and Stucke (2023).

[5]Relatedly, Howell, Parker and Xu (2024) show that general solicitation allows venture capital managers with weaker traditional investor networks to raise capital from a more retail investor base, providing a route through which new managers can establish a track record before obtaining institutional capital. This possibility is particularly relevant in our setting. When a manager's ex ante prospects are sufficiently weak, the expected value of learning his quality may not justify the cost of due diligence, so that he cannot initially attract an informed investor. Retail or crowdfunding capital can then provide an alternative route to establishing a public track record, after which institutional investors may find it worthwhile to evaluate and fund the manager.

fee schedules are largely standardized, publicly disclosed, and primarily a function of the size of the investment. In alternatives, by contrast, Begenau and Siriwardane (2024) show that some investors consistently obtain better terms, and that this persistence is not explained simply by investment size.[6] Moreover, funds managed by GPs without an established track record are substantially more likely to use multiple fee schedules. This pattern is particularly consistent with our mechanism, since uncertainty about manager quality—and hence the potential value of investor information—is likely to be greatest for less established managers. More broadly, it suggests that any positive and persistent net alpha in alternatives may accrue disproportionately to a subset of informed investors rather than uniformly to all capital in the sector, and that differential terms across investors are a central feature to be explained rather than a peripheral detail.

Lerner and Schoar (2004) study the role of liquidity in private equity, but their focus is on the liquidity of LPs' claims and their exposure to liquidity shocks, rather than on the liquidity and opacity of the underlying assets. More recently, Maurin, Robinson and Strömberg (2023) develop a theory in which private equity contracts and returns reflect investors' exposure to illiquidity risk. Although liquidity plays a central role in our model, our mechanism is different. We assume, for simplicity, that investors do not require any premium for holding illiquid investments, or equivalently, that there is sufficient capital available to fund any investment opportunity with positive net present value regardless of how long capital must be committed. In our model, the key role of illiquidity is not to generate a premium directly, but to limit the ability of capital markets to reveal information and reallocate capital, thereby making ex ante due diligence valuable. But our setting does not rule out the possibility that investors also demand compensation for illiquidity. In that case the benchmark asset would also contain this illiquidity premium and our results can be interpreted as the expected return, in excess of this illiquidity premium, associated with the informational advantage of some investors.[7]

We abstract from moral hazard in the agency relationship between investors and managers. Moral hazard is clearly important in practice, and a number of papers study its implications in private equity contracting. For example, Pegoraro (2023) derives an optimal dynamic contract to align the manager's incentives in a delegated asset-management setting. Axelson, Strömberg and Weisbach (2009) show that agency considerations can help explain observed contractual features of the industry, including the use of debt financing in underlying investments, and

[6]Preferential fee terms are commonly offered to large, early-closing, or anchor investors. Practitioner surveys and case studies document management-fee discounts, reduced carried interest, "early bird" discounts, and founder share classes; see MJ Hudson LP Unit (2018), Rainey, MacArthur, Creedon and Nicholson (2026), and Pan and Mason (2019).

[7]Pegoraro, Shive and Zambrana (2025) use Interval Funds to study the liquidity premium.

can also imply that limited partners retain rents in equilibrium. Our mechanism is different. In our model, rents left to informed investors arise because costly information acquisition must be sustained in competitive equilibrium, which in turn requires differential terms across investors and a binding cap on fundraising. More broadly, our point is not that agency frictions are unimportant, but rather that many of the central features of private-market contracts can be explained even in their absence.

The papers closest to ours are those that extend Berk and Green (2004) to frictions that are especially relevant in alternatives. Like us, Marquez, Nanda and Yavuz (2015) extend Berk and Green (2004) to private markets, but they focus on matching with entrepreneurs rather than on investor information acquisition. In their model, entrepreneurs seek to match with high-ability managers, but they cannot perfectly infer ability from observed fund performance because performance also reflects the quality of the firms the manager attracts. Managers may therefore optimally limit fund size in order to be more selective and thereby attract better entrepreneurs.

Brown et al. (2023) also extend Berk and Green (2004) by allowing a small institutional allocator to acquire costly private information about manager skill. Their setting is hedge funds, where the underlying assets are liquid and performance is observed frequently, so the allocator can identify skilled managers before the market and earn temporary excess returns until public learning and capital flows eliminate them. Unlike in our setting, the informed allocator is small and takes both fund size and contractual terms as given, so her investment does not affect scale or create the contracting and dilution problem that is central to our analysis.

Glode and Green (2011) also extend the Berk and Green (2004) benchmark, but in a different direction. In their model, investors and managers learn over time about the profitability of an innovative strategy or sector rather than about manager-specific ability. Because that information can spill over and attract imitation, incumbent managers may allow initial investors to share in the rents so that they are less likely to reallocate capital to competing managers and thereby dissipate those rents.

2 Optimal Information Acquisition and Capital Allocation

We begin by abstracting from the delegation aspect of asset management and consider the optimal information acquisition decision and capital allocation to a potentially skilled manager. In this simple setting, the manager's skill (or the quality of her investment strategy) is uncertain (to everyone, including the manager herself), but can be either learned upfront via costly due diligence, or over time from the strategy's realized performance. In this context

we derive the first-best efficient allocation and the optimal information acquisition decision, which will serve as the benchmark for the delegated setting that is the focus of this article.

The value of conducting initial due diligence depends critically on the ability to reallocate capital ex post, and therefore on the liquidity and the opacity of the investment strategy. When the manager can rebalance frequently, such as in mutual funds holding liquid public securities, the opportunity to update the portfolio using realized returns substitutes for upfront due diligence. By contrast, when rebalancing opportunities are infrequent due to the illiquidity of the underlying assets, or when returns are not easily observable, as in private equity, the benefits of acquiring information in advance are much greater.

To formalize this tradeoff, we model investment skill as in Berk and Green (2004). Managers' investment strategies differ in their ability to generate expected returns in excess of those provided by alternative investment opportunities with similar risk. We also assume there are diminishing returns associated with the scale of investment, as investment opportunities may be limited in scale or subject to price impact.

Specifically, we denote the excess return on the first dollar invested by a given manager/strategy by $\alpha \geq 0$, and assume this return falls by $\beta > 0$ for each additional dollar invested. Given total invested capital q, the realized return in period t is therefore given by

$$R_t(q) = \alpha - \beta q + \epsilon_t. \tag{1}$$

Note that this return is in excess of the expected return on a market investment of equivalent risk (i.e., the appropriate passive benchmark). That is, we assume that the strategy's exposure to any systematic sources of risk is known and hedged via benchmarks. The residual risk, ϵ, is idiosyncratic and mean zero conditional on α and any prior information.

There is also a fixed cost $\delta \geq 0$ of implementing this strategy, which represents the manager's opportunity cost and any necessary resources (data feeds, trading infrastructure, etc.). We assume both β and δ are known and independent of the manager's ability α. Incorporating these costs, the total value added of the manager at time t with committed capital q_t is given by

$$\pi_t = q_t \, R_t(q_t) - \delta. \tag{2}$$

The manager optimally allocates capital in order to maximize the expected value added based on his perception of his ability. Specifically, let $\hat{\alpha}_s$ denote the manager's expectation of his ability given the information he has available at time s:

$$\hat{\alpha}_s \equiv E\left[\alpha | I_s\right].$$

Naturally, the optimal allocation will depend on the perceived profitability, $\hat{\alpha}$, of the strategy given the available information, relative to its expected price impact, β.

Suppose the manager must decide, on date s, the capital allocation for a future date $t \geq s$. Then, maximizing (2) with respect to q_t, we have the following result regarding the optimal capital allocation and the expected value added:

Proposition 1 (Optimal Capital Allocation)**.** *The optimal capital allocation on date t, given the information available on date $s \leq t$, is given by*

$$q^*_{t|s} \equiv \frac{\hat{\alpha}_s}{2\beta}.$$

With that allocation, the fund's expected return is $\hat{\alpha}_s/2$, and the (per period) expected value added of the fund is

$$E\left[\pi_t | I_s\right] = \frac{\hat{\alpha}_s^2}{4\beta} - \delta.$$

Proof of Proposition 1. Maximizing the expectation of (2) with respect to q_t, we have the first-order condition $E[\alpha - 2\beta q_{t|s}|I_s] = 0$, which is solved by $q^*_{t|s} = \hat{\alpha}_s/(2\beta)$, and therefore $E[\pi_t|I_s] = E[\alpha|I_s]q^*_{t|s} - \beta(q^*_{t|s})^2 - \delta = \hat{\alpha}^2/(4\beta) - \delta$. □

The optimal allocation increases with the expected alpha of the strategy. Consequently, to optimally manage money, the manager must adjust the amount of invested capital in response to new information about alpha. An important difference between public- and private-market investments is the frequency with which the manager can respond to new information by adjusting the capital allocation.

Suppose that the capital can only be adjusted every T periods. That is, the capital allocation decision q_t on date t is fixed until period $t + T$. The optimal allocation will then be determined based on the information available on each "reallocation date" according to Proposition 1. To keep the model simple, we will make two further assumptions. First, we assume that the excess return that is earned at any point in this interval only depends on the initial capital investment q_t.[8] Second, we assume that the manager does not have the option to shut down the fund (we relax this assumption shortly). Under these assumptions the per period expected value added is constant between adjustment periods, so the expected value added of the strategy will then be a sequence of T-period annuities, which can be evaluated as follows:

[8]The economic justification for this assumption is that between adjustment periods the scale is fixed, that is, and any intermediate excess returns earn a passive, zero alpha return (or are consumed).

Lemma 1 (Value with Reallocation)**.** *Suppose the amount of capital invested can be adjusted every T periods. Given risk-free rate r, the present value of investing is given by*

$$E\left[\int_0^\infty e^{-rt}\pi_t dt\right] = \left(\frac{1-e^{-rT}}{r}\right)\sum_{n=0}^{\infty} e^{-rnT}\left(\frac{E\left[\hat{\alpha}_{nT}^2\right]}{4\beta} - \delta\right). \tag{3}$$

Proof of Lemma 1. Recall that $q_t = q^*_{t|s}$ for $t \in (s, s+T]$. The expected value added is therefore constant over this period, and equal to $\hat{\alpha}_s^2/(4\beta) - \delta$. Because all risk is idiosyncratic, we compute the present value using the annuity formula starting on date s using risk-free rate r:

$$e^{-rs}\left(\frac{1-e^{-rT}}{r}\right)\left(\frac{E\left[\hat{\alpha}_s^2\right]}{4\beta} - \delta\right).$$

The result then follows by summing this result over each adjustment date nT for all n. □

Now suppose the manager can pay a cost $C \geq 0$ to learn his true α. To compute the gain associated with paying this cost, we can compare equation (3) with and without the ex-ante information. When the manager pays the cost and acquires the information, he will learn his true alpha upfront and allocate capital optimally at all times. In that case, the term $E\left[\hat{\alpha}_{nT}^2\right]$ in (3) becomes $E\left[\alpha^2\right]$ when calculating the expected value added of the fund. Thus, the value of information depends on the difference between $E\left[\alpha^2\right]$ and $E\left[\hat{\alpha}_{nT}^2\right]$.

By the law of total variance, for any s, we have

$$E\left[\alpha^2\right] - E\left[\hat{\alpha}_s^2\right] = Var(\alpha) - Var(\hat{\alpha}_s) = E\left[Var(\alpha|I_s)\right].$$

In words, the gain from knowing the true α at any date is proportional to the remaining uncertainty that would be eliminated. This gain declines with s since the remaining uncertainty falls as additional returns are observed.

So, with full information, capital can be efficiently deployed at the optimal level. Without full information, capital is inefficiently deployed, but this inefficiency will be reduced on every adjustment date by using information that is learned in the interim to refine the estimate of α. The longer the interval of illiquidity, the more costly the inefficiency. The cost also increases when returns are noisier, slowing the rate of learning. The following proposition formalizes these insights:

Proposition 2 (Gain from Information)**.** *Given illiquidity period T, the gain from learning the true alpha of the strategy upfront is given by*

$$G(T) = \left(\frac{1-e^{-rT}}{4\beta r}\right)\left(\sum_{n=0}^{\infty} e^{-rnT} E\left[Var(\alpha|I_{nT})\right]\right). \tag{4}$$

The value of information, G, is higher when the illiquidity period T is longer, the manager's skill is more uncertain, or when returns are noisier. In particular, let $\sigma_\alpha^2 \equiv E\left[Var(\alpha|I_0)\right]$, then note that

$$G(T) \to \frac{\sigma_\alpha^2}{4\beta r} \quad as \quad T \to \infty.$$

Proof of Proposition 2. The expression for G follows immediately from (3) and the law of total variance. Because the residual variance of alpha, $Var(\alpha|I_t)$, declines with t (as more information is learned), we can think of (4) as the left Riemann sum approximation of the integral of a decreasing function, which declines with a shorter step size (T). Thus, G is monotone in T. G also increases if the initial uncertainty σ_α^2 is higher, or returns are noiser (in the usual sense), since that will raise the residual uncertainty that remains at any point. Finally, when T approaches infinity, there is no opportunity to update, so the residual uncertainty is constant in perpetuity. □

Note that the value of information is maximal when it is impossible to adjust capital. This limiting value of G provides a necessary condition to pay for information. So long as

$$C < \frac{\sigma_\alpha^2}{4\beta r}, \tag{5}$$

it will be profitable to learn α when the rebalancing horizon T is sufficiently long.

Thus far, we ignored the option the manager has to shut down the fund. In Appendix A we generalize the above results to include this option and show that again, the value of information increases with the illiquidity period.

From Proposition 1, it is profitable to operate the fund given the information available on date s if

$$\frac{\hat{\alpha}_s^2}{4\beta} - \delta \geq 0. \tag{6}$$

We can use (6) to define the threshold belief, α_M, that determines when the value added from investing exceeds the manager's opportunity cost δ.[9] Solving (6) gives an explicit expression for the threshold α_M:

$$\alpha_M \equiv 2\sqrt{\beta\delta}. \tag{7}$$

Given this threshold, it is convenient to define the *first-best surplus* generated each period by a manager of known skill α under the efficient policy — operate if and only if $\alpha > \alpha_M$, at

[9]Note that this condition is sufficient but not necessary, as it might be optimal to operate the fund when $\hat{\alpha}_s < \alpha_M$ because operating affords the option to learn about ability. Obviously when $T = \infty$ the option does not exist and so (6) is also necessary.

scale $\alpha/2\beta$:

$$S^{FB}(\alpha) \equiv \left(\frac{\alpha^2}{4\beta} - \delta\right) \mathbf{1}_{\{\alpha>\alpha_M\}} = \left(\frac{\alpha^2 - \alpha_M^2}{4\beta}\right) \mathbf{1}_{\{\alpha>\alpha_M\}}. \tag{8}$$

Indeed, consider the case with $T = \infty$, so that there is no opportunity to adjust capital. Then from Proposition 1, without acquiring information, operating the fund generates an expected present value of

$$E\left[\int_0^\infty e^{-rt}\pi_0\, dt\right] = \frac{1}{r}\left(\frac{\hat{\alpha}_0^2}{4\beta} - \delta\right). \tag{9}$$

This expression need not be positive: if $\hat{\alpha}_0 < \alpha_M$, the manager will choose not to launch, and the fund's value is zero. Going forward, we assume this to be the case.

Now consider the decision to acquire information. Once alpha is learned, the manager will only invest if $\alpha > \alpha_M$. Therefore, investigating is optimal only if the expected gain from running the fund when it is worthwhile to do so exceeds the cost of acquiring information:

$$\frac{1}{r}E\left[\left(\frac{\alpha^2}{4\beta} - \delta\right) \mathbf{1}_{\{\alpha>\alpha_\mathbf{M}\}}\right] - C = \frac{1}{r}\, E\left[S^{FB}(\alpha)\right] - C > 0. \tag{10}$$

Again, this expectation captures the fact that the fund will only launch when the manager learns that $\alpha > \alpha_M$.

To summarize, we assume the horizon T is sufficiently long, and prior beliefs are sufficiently poor ($\hat{\alpha}_0 < \alpha_M$), so that the fund is not worth operating without prior investigation. We also assume it is optimal to investigate (that is, (10) holds), so that the fund will operate if it is revealed that $\alpha > \alpha_M$.

3 Delegated Capital Management: A Linear Benchmark

We now turn to the delegation problem between managers and investors. We consider a potentially skilled manager with no capital, who must raise capital from investors. As in the previous section, initially neither the manager nor the investors know the skill of the manager. But any investor has the ability to conduct costly due diligence and become informed.

The case in which the underlying assets are liquid and returns are sufficiently informative so that it is suboptimal to acquire information upfront is the focus of Berk and Green (2004). In that setting, investors rationally wait to learn from performance, capital flows adjust over time, and the standard mutual fund contract implements the first-best allocation. Our focus here is the complementary case in which the underlying assets are sufficiently illiquid and opaque so that it is optimal to acquire information before capital is committed.

The goal of this section is to derive the optimal contract for the manager in this delegated equilibrium. We begin with the simplest setting that makes the economics transparent: linear contracts and unlimited liability on the part of the manager. This benchmark allows us to isolate the free-rider problem created by costly information acquisition, to characterize the roles of differential investor terms and fundraising caps, and to show how delegation can implement the first-best allocation.

The result is surprising: without sacrificing any rents, the manager can dictate an optimal contract that raises the optimal level of capital while still inducing an investor to both become informed and reveal her information. In Section 4, we show that this result is robust to imposing limited liability. In a limited liability setting, the optimal contract is non-linear and matches the standard option-based private-market contract with carried interest and a hurdle rate, while still maintaining efficiency without sacrificing rents.

3.1 An Optimal Fully-Revealing Linear Contract

We begin by deriving an optimal fully revealing contract with linear carry, meaning that the manager receives a constant share of the fund's net return. Importantly, linear carry implies that the manager pays the investor if the net return is negative. This linear contract allows us to derive an equilibrium in which the underlying economics are transparent. It attains the first-best allocation given the information constraints and delivers the same payoff to the manager as would arise if the manager invested his own capital. It also makes clear why the mutual fund contract, which uses only a fee proportional to assets under management, cannot implement the efficient outcome in private markets. In Section 4, we then show how these same economic forces lead, under limited liability, to the standard option-based private-market contract with carried interest and a hurdle rate.

As before, we assume that α is the manager's skill, with decreasing return to scale determined by β, and subject to a fixed cost δ, so that the total value added is given by (2). The manager, who does not know α, proposes the following contract to an investor who is willing to become informed:

- A management fee of f_i per dollar invested by the informed investor;
- A fractional share ("carry") s of the net return that the manager will retain;
- A cap $\bar{q}$ on overall fundraising that depends on the amount of capital q_i that the informed investor commits.

Within this framework, we show that

- There exist terms $(f_i, s, \bar{q})$ that induce an investor to become informed, implement the first-best allocation, and represent an optimal contract for the manager;
- Given the capital commitment q_i of the informed investor, the manager's contract with uninformed investors has the same carry, s, but has a higher management fee f_u, and the capital constraint $\bar{q}$ is binding.

With the linear sharing rule, we will show that first-best allocation can be obtained with a cap on overall fundraising that is linear in the amount of informed capital, q_i:

$$\bar{q}(q_i) = q_0 + k \cdot q_i \tag{11}$$

We will show later that the manager will raise funds up to this cap, so that q_0 represents the minimum fund size, and $k > 0$ limits how much the fund can expand relative to the contribution of the informed investor. Importantly, the choice of k, hereafter called the *participation factor*, limits the amount of uninformed capital the manager can raise along side the informed investor. Specifically, if we subtract the level of informed capital from (11), the cap on the uninformed capital that the manager can raise is:

$$q_u = \bar{q}(q_i) - q_i = q_0 + (k-1)q_i \tag{12}$$

For now, we restrict attention to the case with $k \geq 1$, so that the amount of uninformed capital is increasing in the amount of informed capital. Later we will discuss the equilibrium determinants of k.[10]

3.1.1 Informed Investor's Problem

Consider the informed investor's decision regarding whether to become informed and the amount of capital to commit. In analyzing this decision, we assume for now that the informed investor expects that the manager will always choose to raise the maximum amount of capital $\bar{q}$, and show later that it will indeed be optimal for the manager to do so.

Suppose the investor chooses to investigate and learn alpha. Once informed, the investor will then choose an amount q_i to invest to maximize her net profits from investing:

$$q_i(1-s)E\underbrace{[R_t(\bar{q}(q_i)) - f_i]}_{\text{pre-carry net return}} = q_i\,(1-s)\,(\alpha - \beta(q_0 + kq_i) - f_i)\,. \tag{13}$$

[10]The case $k \geq 1$ is certainly more consistent with practice, and we will show later that there is a degree of indeterminacy in k so that $k > 1$ may be without loss of generality. We will also show that this linear form of the cap is uniquely optimal when carry is linear — that is, a linear cap is the only functional form that supports the efficient outcome.

Solving the first-order condition for q_i,

$$q_i(\alpha) = \frac{\alpha - (\beta q_0 + f_i)}{2\beta k} = \frac{\alpha - \alpha_I}{2\beta k}, \tag{14}$$

where α_I is defined as

$$\alpha_I \equiv \beta q_0 + f_i. \tag{15}$$

Note that the expected pre-carry return is positive only if $\alpha > \alpha_I$. So, the informed investor will only invest if $\alpha > \alpha_I$; that is, α_I is the minimum level of managerial ability required for the informed investor to commit capital, hereafter the *investment threshold*.

Given this investment rule, the net return of the informed investor when she chooses to invest is equal to

$$(1-s)\left(\alpha - \beta(q_0 + kq_i) - f_i\right) = (1-s)\,\frac{\alpha - \alpha_I}{2}. \tag{16}$$

Multiplying (16) by the optimal investment, (14), the informed investor's expected profit is given by:

$$\pi_i(\alpha) = \begin{cases} \left(\frac{1-s}{k}\right)\frac{(\alpha-\alpha_I)^2}{4\beta} & \text{if } \alpha > \alpha_I \\ 0 & \text{if } \alpha \leq \alpha_I \end{cases} \tag{17}$$

Recall that to become informed, the informed investor incurs a cost C upfront. In order to ensure that she can recoup this cost over the T-period life of the investment, the informed investor's per-period expected profit must be at least c_T, where,[11]

$$c_T \equiv \left(\frac{r}{1 - e^{-rT}}\right) C, \tag{18}$$

that is, the equivalent per-period annuitized cost.

Using our earlier notation $\hat{\alpha}_0$ to represent the expected ability of the manager prior to investigating, the informed investor therefore finds it optimal to investigate as long as

$$E\left[\pi_i(\alpha)\right] - c_T \geq E\left[\pi_i(\hat{\alpha}_0)\right], \tag{19}$$

where $E\left[\pi_i(\hat{\alpha}_0)\right]$ is the expected profit from investing without information. Note that if $\hat{\alpha}_0$ is below the investment threshold α_I, then the right-hand side of (19) is zero, and the optimality

[11] Because the informed investor's information will be fully revealed in equilibrium, we assume he must recoup his investigation cost before the subsequent funding round in period T. In principle, the manager could write a longer term contract and the informed investor could recoup the investment over a longer horizon. In this case the net alpha will remain positive and will persist over funding rounds.

condition for investigation becomes,

$$E\left[\pi_i(\alpha)\right] \geq c_T. \tag{20}$$

3.1.2 Manager's Problem

Having characterized how the informed investor responds to the contract, we now turn to the manager's problem of choosing contract terms so that the induced investment rule implements the first-best allocation. We will show that by choosing the contract parameters (s, f_i, q_0, k) appropriately, the manager can achieve the first-best allocation and extract all surplus. The reason the first-best allocation is attainable is that once the manager and uninformed investors observe the informed investor's capital commitment q_i, they can infer α from (14) and information becomes symmetric. Importantly, because the first best is attained, the resulting contract must be optimal for the manager.

To set the terms of the contract to achieve the first best, there are two conditions that must be satisfied. First, the fund should operate whenever the value added from investing exceeds the manager's opportunity cost, that is, whenever $\alpha > \alpha_M$ as defined by (7). Hence, the manager's and the informed investor's investment thresholds must coincide:

$$\alpha_I = \alpha_M = 2\sqrt{\beta\delta}. \tag{21}$$

Second, from Proposition 1, the total capital $\bar{q}$ that is invested should match the efficient level. Therefore,

$$\bar{q}(q_i) = q_0 + kq_i = \frac{\alpha}{2\beta}. \tag{22}$$

Substituting the optimal investment function $q_i(\alpha)$ of the informed, (14), into the above, we have

$$q_0 + \frac{\alpha - \alpha_I}{2\beta} = \frac{\alpha}{2\beta}. \tag{23}$$

Solving for q_0 and applying (21), we have

$$q_0 = \frac{\alpha_I}{2\beta} = \frac{\alpha_M}{2\beta} = \sqrt{\frac{\delta}{\beta}}. \tag{24}$$

Finally, using $\alpha_I = \beta q_0 + f_i$ from (15) implies that the optimal management fee for the informed investor is

$$f_i = \frac{\alpha_M}{2} = \sqrt{\beta\delta}. \tag{25}$$

Together, (24) and (25) determine the (f_i, q_0) that ensure the capital allocation is optimal.

The remaining contract terms (s, k) and the management fee f_u paid by the uninformed are determined so that all investors break even, thereby ensuring that all rents accrue to the manager. Consider first the uninformed management fee. It must be set to ensure that the uninformed investor's earns zero net return:

$$(1-s)\left(\alpha - \beta\bar{q}(q_i) - f_u\right) = 0, \tag{26}$$

which, using (22), requires the fee to be

$$f_u = \alpha - \beta\left(\frac{\alpha}{2\beta}\right) = \frac{\alpha}{2}. \tag{27}$$

That is, the uninformed management fee is set equal to the fund's (gross) expected excess return so that the uninformed earn a zero expected pre-carry return. Because carry is symmetric, it has zero expected cost, and so the uninformed investor's expected post-carry return is also zero. While α is not known initially, it is revealed by the size of the informed investor's commitment q_i (or equivalently by the size of the cap, $\bar{q}$, on the total capital raised). Thus, the fee f_u can be implemented conditional upon the informed investor's committing to invest.[12] Note that if the fund is launched, because $\alpha > \alpha_M$, the uninformed management fee exceeds that of the informed, $f_u > f_i$, implying "most favored nation" treatment of the informed investor.

Lastly, the manager sets the contract so that the informed investor earns zero rents net of her investigation cost. Equivalently, since any investor can choose to become informed, investors must be indifferent between acquiring the information or not. That is, under the assumption that $\hat{\alpha}_0 \leq \alpha_M$ the participation constraint, (20), holds with equality, implying that the final parameters (s, k) satisfy:

$$E\left[\pi_i(\alpha)\right] = \left(\frac{1-s}{k}\right) E\left[\frac{(\alpha-\alpha_M)^2}{4\beta} \cdot \mathbf{1}_{\{\alpha>\alpha_M\}}\right] = c_T. \tag{28}$$

Notice that (s, k) are not uniquely determined. There is a one-dimensional family of solutions that satisfy (28). That is, the manager can either set the ratio of informed to uninformed capital by specifying k and then let the carry be determined by the distribution of skill in the market, or he can set the carry and let the ratio be a function of the distribution of skill in the market.[13] Empirically, the cross-sectional variation of the carry charge appears to be

[12] It is common in practice for funds to use the commitment of a high reputation investor to raise additional capital.

[13] This indeterminacy mirrors a similar result in the mutual fund space. As Berk and Green (2004) show, a mutual fund manager can either choose the management fee and let the market determine the size of the fund

low, likely implying that the participation factor k, which controls the ratio of informed to uninformed capital, adjusts based on variation in market perception of managerial ability. Later, we discuss factors that can impact the tradeoff between s and k.

3.1.3 Why the Capital Constraint Binds

Thus far, we have assumed that the manager raises capital up to the cap and does not exceed it. We now show that this assumption is satisfied in equilibrium. Moreover, we show that the cap is indeed binding: without it, the manager would continue raising capital by lowering uninformed fees.

Note first that because the manager sets q_u and f_u after α is revealed, f_u will be set by the manager so that the uninformed earn zero expected net return. As a result, no matter what q_u is, the manager captures all rents attributed to uninformed capital. Therefore, the manager's payoff for a given choice of q_u is given by

$$\begin{aligned}
\pi_m(q_u) &= E[\text{ informed fees} + \text{uninformed rents} - \delta\ ] \\
&= q_i(f_i + s\left[\alpha - \beta(q_i+q_u) - f_i\right]) + \underbrace{q_u\left[\alpha - \beta(q_i+q_u)\right]}_{\text{uninformed rents}} - \delta \\
&= (sq_i + q_u)\left[\alpha - \beta(q_i+q_u)\right] + (1-s)q_i f_i - \delta. \qquad (29)
\end{aligned}$$

The first-order condition for the manager's choice of q_u is then

$$\begin{aligned}
\pi_m'(q_u) &= \left[\alpha - \beta(q_i+q_u)\right] - \beta\left(sq_i + q_u\right) \\
&= \alpha - 2\beta(q_i+q_u) + \beta(1-s)q_i \qquad (30) \\
&\geq \beta(1-s)q_i \\
&> 0.
\end{aligned}$$

The first inequality follows when total capital is below the efficient level — that is, $q_i + q_u \leq \frac{\alpha}{2\beta}$. The second inequality holds because $\beta(1-s)q_i$ is always strictly positive given $s < 1$. Therefore, the manager has an incentive to increase capital even beyond the efficient level of capital, and so the cap strictly binds.

The reason the cap binds even at the efficient level is as follows. The manager could lower the uninformed fee to raise additional capital. Doing so would reduce total surplus because additional capital is inefficient. But this loss is second order. At the same time, increasing

he manages or he can choose the size of the fund and let the market decide the management fee. Because the cross-sectional variation in fund size is much larger than management fees, in that space it appears that managers opt to set the fee. It appears that a similar result holds in the alternative space.

capital imposes a first-order transfer from the informed investor — who suffers a reduction in returns of $\beta(1-s)$ without any increase in quantity — to the manager. The cap prevents the manager from exploiting this externality.

This result, that the capital constraint binds, explains one of the most perplexing features of private market investing. Unlike open-ended mutual funds, private investment funds often have binding limits on the capital raised. It might therefore appear that the fund manager is giving up rents by restraining his own ability to raise additional funds. Our model shows that instead such constraint is necessary the manager to maximize his payoff — without such an ex ante commitment, the manager would be unable to raise capital from investors.

Now that we have established that the capital constraint binds, we have can state our main result:

Proposition 3 (First-Best Linear Contract)**.** *A contract (s, f_i, q_0, k) satisfying (24), (25), and (28) is optimal for the manager and achieves the first-best capital allocation. Under this contract, the manager starts the fund and raises capital $q_u = \bar{q}(q_i) - q_i$ from uninformed investors, charging $f_u > f_i$ given by (27), if and only if the informed investor commits $q_i > 0$. All investors break even, and the manager earns an expected per period profit of* [14]

$$E\left[\left(\frac{\alpha^2}{4\beta} - \delta\right)\ \mathbf{1}_{\{\alpha > \alpha_{\mathbf{M}}\}}\right] - c_T, \tag{31}$$

consistent with (10).

Proof of Proposition 3. From the analysis in the text, the proposed contract achieves the first-best payoff given by (10). Therefore, it is an optimal contract. □

3.2 The Economic Role of Fees, Carry, and the Capital Constraint

Having established the linear benchmark and shown why the fundraising cap binds, we can now interpret the economic role of each component of the contract. To achieve the first best, two conditions need to be satisfied. First, the set of states in which the fund invests must be efficient. Second, all extracted rents must accrue to the manager. Because the informed investor determines the set of states when the fund invests, his contract must satisfy this dual purpose. As shown in (25), the management fee is set to ensure that the set of states in which investment occurs is optimal by determining the critical alpha threshold. The role of carry and the cap on total capital, ensures that the manager extracts all the rents, that is, that the informed investor breaks even.

[14] As a reminder, the fact that the fund is only launched when $\alpha > \alpha_M$ is equivalent to $\alpha^2/(4\beta) > \delta$.

As we have already noted, there is an indeterminacy between k and s, that is the participation factor and the carry fee. The manager can directly reduce the rents to informed capital by raising the carry s. But he can also reduce rents by increasing the ratio of uninformed capital to informed capital. Intuitively, because free riding by uninformed investors comes at the expense of the informed investor, the manager can extract rents from informed investors either by increasing the amount of free riding (by increasing k), or limiting the free riding (by keeping k low) and instead increasing the carry, s.

Notice that in this equilibrium, the management fee paid by the informed is not set to extract rents. Instead, it is set to ensure that the manager makes at least his opportunity cost even in the worst-case scenario when the manager turns out to have ability near α_M. The form of the optimal contract also makes it clear why the manager has to take uninformed capital at all. Because the informed investor pays fees, his optimal investment is always less than the efficient level. It is zero when $\alpha \leq \alpha_M$, and rises continuously when $\alpha > \alpha_M$. On the other hand, the efficient level of capital jumps discontinuously at α_M, from 0 to $q_0 = \frac{\alpha_M}{2\beta}$. This gap must be filled by uninformed capital.

This role of the management fee accords well with the common justification given in the industry for charging a management fee — to ensure that the fixed costs of managing money are covered. If the fee was not set in this way, there would be states in which either the manager would choose to quit when his ability is revealed, or the informed investor would choose not to invest even though starting the fund is optimal.

In contrast, in a public-market mutual fund the assets are liquid and investors have insufficient incentive to investigate ex ante. Because managers and investors are therefore symmetrically informed, so they always choose to invest in the same states. As a result, the manager can use the management fee to extract the rents. In private markets, when the manager and investor are asymmetrically informed, an additional instrument is required for rent extraction.

3.3 High Investigation Costs and the Limits of Uninformed Capital

A technical issue we have not yet addressed is when (28) admits a solution with both $s \geq 0$ and $k \geq 1$. As long as the investigation cost satisfies

$$c_T \leq c^* \equiv E\left[\frac{(\alpha - \alpha_M)^2}{4\beta} \cdot \mathbf{1}_{\{\alpha > \alpha_M\}}\right], \tag{32}$$

then (28) implies $(1-s)/k \leq 1$ and such a solution exists. But investigation remains efficient over the wider range

$$c^* < c_T < E[S^{FB}(\alpha)] \equiv E\left[\frac{(\alpha^2 - \alpha_M^2)}{4\beta} \cdot \mathbf{1}_{\{\alpha > \alpha_M\}}\right], \tag{33}$$

where the upper limit is the total surplus generated by first-best investment using the definition of $\delta = \alpha_M^2/(4\beta)$.[15]

Reducing k below one changes the role of the cap. From (12), the uninformed position $q_u = q_0 + (k-1)q_i$ is now decreasing in the informed stake, since each informed dollar displaces $1-k$ dollars of uninformed capital. Indeed, uninformed capital is extinguished once the informed investor commits

$$q^* \equiv \frac{q_0}{1-k}, \tag{34}$$

at which point $\bar{q} = q_i = q^*$. We call q^* the *solo-investor threshold*: the fund size beyond which the informed investor is the fund's only investor. Since the efficient fund size is $\alpha/2\beta$, it corresponds to a skill level

$$\alpha^* \equiv \frac{\alpha_M}{1-k} = \alpha_M \frac{q^*}{q_0}. \tag{35}$$

If skill never exceeds α^* the contract is unchanged. Otherwise it requires one modification, because the cap (11) can no longer describe fund size: the uninformed position cannot be negative. The solution is to revise the cap to allow the informed investor to continue to invest above q^* while keeping uninformed capital at zero.

$$\bar{q}(q_i) = \max\{q_i,\, q_0 + kq_i\}. \tag{36}$$

In addition, once the informed investor is the fund's only investor, her incentives are fully aligned with maximizing the fund's expected surplus. Hence, any marginal fee makes her internalize a cost that is not social. We therefore replace the constant fee with the schedule[16]

$$f^*(q_i) = f_i \cdot \min\{q_i,\, q^*\}, \tag{37}$$

so that the informed investor pays the management fee only up to the solo-investor threshold and invests without a fee beyond it. Finally, for $k \geq 1$ we set $q^* = \infty$ and nothing changes.

[15] To gauge magnitudes, suppose α/α_M is normal with mean one and volatility v. Then $c^* = \delta v^2/2$ and $E[S^{FB}(\alpha)] = \delta\left(v^2/2 + \sqrt{2/\pi}\, v\right)$. At $v = 1.5$, entering this range requires the annualized cost of investigating a manager to exceed the operating cost of the fund by a factor of 1.125, with an upper limit of 2.32. In this range, the manager must set $k < 1$ to adequately compensate the informed investor.

[16] While a linear cap is optimal with a linear fee schedule, a kinked cap is consistent with a kinked fee schedule. We explore the relation between the cap and the fee schedule further in Section 4.

Proposition 4 (Efficiency with a Solo Investor)**.** *Fix any $s \in [0,1)$ and $k > 0$, and let the manager offer the contract (f_i, s, q_0, k) of Proposition 3 with cap* (36) *and fee schedule (37). Then the informed investor's commitment is strictly increasing and fully revealing, the fund raises the first-best capital $\bar{q} = \alpha/2\beta$ at every skill level, and the contract remains optimal for the manager.*

Proof. The investor maximizes $(1-s)\left[q_i(\alpha - \beta\bar{q}(q_i)) - f^*(q_i)\right]$. For $q_i \leq q^*$ marginal profit is $\alpha - \beta q_0 - 2\beta k q_i - f_i = \alpha - \alpha_M - 2\beta k q_i$, using $\alpha_I = \beta q_0 + f_i = \alpha_M$; for $q_i \geq q^*$ it is $\alpha - 2\beta q_i$. At $q_i = q^*$ both equal $\alpha - \alpha^*$, because the jump in the marginal cost of crowding, $\beta q^*(1-k) = \beta q_0$, is exactly the fee waived. Marginal profit is therefore continuous and strictly decreasing, the objective is strictly concave, and the first-order conditions give $q_i = (\alpha - \alpha_M)/2\beta k$ for $\alpha \leq \alpha^*$ and $q_i = \alpha/2\beta$ above, which is continuous and strictly increasing. Skill is thus revealed, uninformed investors break even when they participate, and substituting into (36) gives $\bar{q} = \alpha/(2\beta)$ on both branches — the first-best level. Below q^* the manager's incentive to raise the maximum uninformed capital is unchanged, above q^* none is available. Because the allocation is first best and all investors break even, the manager earns $E[S^{FB}(\alpha)] - c_T$, the most attainable. □

When $s = 0$ the manager's fee income above the threshold is $f_i q^* = \delta/(1-k)$, independent of α: beyond the solo-investor threshold he sells access at a fixed price and no longer participates at the margin, which is precisely why the allocation remains efficient and delegation does not fail. Indeed, under the revised contract, $E[\pi_i]$ rises continuously from $(1-s)c^*$ to $(1-s)E[S^{FB}(\alpha)]$ as k falls from one to zero, so every investigation cost in (33) is implementable with a non-negative carry $s \leq 1 - c_T/E[S^{FB}(\alpha)]$. Only in the limit $c_T \to E[S^{FB}(\alpha)]$ does it degenerate: $k \to 0$, $q^* \to q_0$, and the informed investor funds the entire fund at every skill level while the manager's fees just cover his fixed cost. Section 5 shows that direct contracting between an informed investor and the manager cannot be efficient except in exactly this limit; as long as the manager earns positive rents, delegated management is strictly Pareto superior.

3.4 Implementing the Linear Contract

The contract we have outlined can be implemented in a simple sequence. First the manager calculates the minimum skill level at which he would still choose to manage money, α_M, from his opportunity cost, δ, and the technology parameter β using (7). Next, he uses this parameter to calculate f_i using (25) and the minimum fund size q_0 using (24), and picks any reasonable s. He then calculates k using his choice of s based on the investigation cost and (28). Using these parameters, he presents potentially informed investors with the contract

(f_i, s, q_0, k). On seeing this contract, an investor chooses to incur the cost C to acquire information, and if the resulting α exceeds α_M she chooses to invest q_i as given by (14). Seeing this investment level, uninformed investors (and the manager) infer α by inverting (14) as follows:

$$\alpha = \alpha_M + 2\beta k q_i.$$

Thus, implementation does not require the informed investor to communicate her information directly: her investment choice reveals it. The manager then offers the remaining (uninformed) investors the contract $\{f_u, s\}$, where f_u is given by (27). Note that since $\bar{q}(q_i) = \alpha/2\beta$, we can equivalently set the uninformed management fee proportional to the maximum fund size:

$$f_u = \beta \bar{q}. \tag{38}$$

The manager then solicits the maximum allowable capital from uninformed investors, which will equal $\bar{q}(q_i) - q_i$ as specified in (12). In practice, the manager could set a maximum fee ex ante and offer discounts as needed to fill the remaining capacity.[17]

This implementation has a natural counterpart in private-market fundraising. New managers almost always seek a cornerstone investor with a successful reputation. This investor is offered a larger capital allocation and a break in fees in exchange for committing capital early. The preferential terms compensate the anchor investor for undertaking costly diligence, while the size of her investment communicates the resulting information to other investors. The manager can then use this information to raise capital from investors who did not themselves incur the cost of becoming informed.

The fund-size limit is crucial to this implementation rather than an incidental feature. By committing ex ante to limit subsequent fundraising, the manager preserves the informed investor's incentive to perform diligence. Once her investment has revealed the fund's quality, the manager would otherwise have an incentive to raise additional capital from uninformed investors, undermining the preferential treatment that induces information production in the first place. Thus, the combination of preferential terms for an anchor investor and a commitment to limited fund capacity implements the delegation mechanism at the heart of the model.

The main difference between this benchmark contract and those observed in practice is

[17] If after the manager's skill is revealed, the maximum fee turns out to be below that in (38), uninformed investors retain some rents and the manager earns less, but the informed investor's incentives and the first-best allocation are unchanged. Equation (38) sets the uninformed fee to extract all rents from uninformed investors, but full rent extraction is not essential to the implementation. The informed investor's incentives depend on her own contract and the fund-size limit, not on how the remaining surplus is divided between the manager and uninformed investors. Thus, any fee that induces uninformed investors to supply the remaining capital up to $\bar{q}$ implements the same allocation.

the linear carry charge. Because the manager would have to pay investors when realized returns are sufficiently negative, the contract requires unlimited liability. In practice, of course, managerial liability is limited by both personal wealth and bankruptcy law, so the linear contract is not directly implementable.

The next section shows that this limitation does not overturn the underlying economics. We replace the linear carry with the option-based carry used in practice, under which carry is earned only when returns exceed a hurdle. We show that limited liability then rationalizes the standard private-market contract, while preserving the central economic logic of the model and, most importantly, the first-best allocation.

4 Option-Based Performance Contracts

Although the linear benchmark in Section 3 attains the first-best allocation, it requires unlimited liability for the manager. When the fund's return is sufficiently negative, the manager must reimburse investors for a fraction s of the loss. Such a contract is therefore not directly implementable in practice, both because the manager's ability to make such payments is uncertain and because bankruptcy law limits investors' ability to enforce them.

In practice, private-market contracts restore limited liability by making carry asymmetric: the manager receives a share s only of returns in excess of a target hurdle rate. As long as the hurdle rate is non-negative, the manager's minimum carry payment is also non-negative, and liability is limited. Typical hurdle rates in private equity range from zero to ten percent, depending on the asset class, with zero common in U.S. venture funds and 8% the most common in private equity.[18] We now show that replacing the linear carry with this option-based performance fee continues to implement the first-best allocation.

Let $h \geq 0$ be the hurdle rate. Investors are charged the carry s on the portion of their return net of management fees that exceeds the hurdle h. When the investor's net return is below h, carry is zero (the manager does not rebate the investors). That is, the carry owed to the manager by an investor with management fee f, given (gross) return R, can be calculated as[19]

$$\text{Carry}(R) = s(R - f - h)^+.$$

[18] See Goodwin Insight, November 2023. Note that hurdles may be soft or hard; hard hurdles compensate the manager only on the portion of the return that exceeds the hurdle, as modeled here, whereas soft hurdles compensate the manager based on the entire return once the hurdle is met.

[19] As a simplification, we compute carry each period. In private equity, carry is computed based on the current cumulative return, and may include clawback provisions if the fund subsequently underperforms. Hedge funds pay carry periodically but often contain a high water mark, effectively paying carry on cumulative returns only.

We refer to this carry formula as "option-based" carry, as opposed to the "linear" carry of the prior section.

This modification changes the shape of the informed investor's incentives, but not the basic structure of the problem. As we show below, all of the main insights of the prior section continue to hold. The important contract features, such as the cap on uninformed capital, and a lower management fee for the informed investor, continue to be optimal. The main qualitative change is that the amount of capital invested by the informed, as well the cap on uninformed capital, are no longer linear in the manager's perceived alpha.

4.1 Gross versus Net Returns

The key new object in the option-based setting is the mapping from expected pre-carry returns to expected (post-carry) net returns. Recall that a manager whose trading strategy has quality α, and who invests total capital q, generates a gross return equal to

$$R(q) = \alpha - \beta q + \epsilon. \tag{39}$$

For an investor with management fee f, the expected *pre*-carry return is therefore

$$\bar{R}^{\text{pre}} = E\left[R(q) - f\right] = \alpha - \beta q - f. \tag{40}$$

Given carry s with hurdle h, the investor's expected post-carry net return is given by the function $\mathcal{R}_{s,h}$, defined by

$$\bar{R}^{\text{net}} = \mathcal{R}_{s,h}\big(\bar{R}^{\text{pre}}\big) \equiv \bar{R}^{\text{pre}} - s\,E\Big[\big(\bar{R}^{\text{pre}} + \epsilon - h\big)^{+}\Big]. \tag{41}$$

When no ambiguity arises, we suppress the subscripts and write $\mathcal{R}$. For example, with this definition of $\mathcal{R}$, the expected net return for the informed investor given a manager of quality α managing a fund at the efficient investment level, $\frac{\alpha}{2\beta}$, is equal to $\mathcal{R}(\alpha/2 - f_i)$.

In the linear carry model, $\mathcal{R}(r)$ would equal $(1-s)r$, proportional to the pre-carry return. With option-based carry, the net return function $\mathcal{R}$ is strictly increasing and concave. The concavity arises because better performance raises the probability that carry will be charged.[20]

The concavity of $\mathcal{R}$ has two important implications. First, because expected net returns are no longer linear in alpha, the informed investor's optimal capital commitment will no longer increase linearly with alpha. Second, the concavity of $\mathcal{R}$ raises the break-even level of the pre-carry return. To see why, note that $\mathcal{R}(0) < 0$, since even with a zero expected pre-carry

[20] To see this precisely, $\mathcal{R}'(r) = 1 - s\Pr(r + \epsilon > h)$.

return, expected carry is positive — the manager will earn carry when the idiosyncratic shock is above the hurdle and will not pay carry when it is below (that is, the option component of compensation always has positive expected value).

The second implication is central for the contract design problem. In the linear model, investors invest whenever the pre-carry return is positive. With option-based carry, a positive pre-carry return is no longer sufficient. Instead it must exceed the break-even pre-carry expected return, r_0, which is strictly positive, and is defined as the unique solution to

$$\mathcal{R}(r_0) \equiv 0. \tag{42}$$

Note that, solving this equation using (41) implies that r_0 is equal to the expected carry charge at break even,

$$r_0 = sE\left[(r_0 + \epsilon - h)^+\right] > 0. \tag{43}$$

Intuitively, $r_0 > 0$ is the pre-carry return that just offsets the value of the manager's carry option. The informed investor will invest whenever his expected pre-carry return exceeds r_0 (rather than zero, as in the linear case). Managers set fees so that uninformed investors earn zero expected net return, which requires their pre-carry expected return to equal r_0 (again, rather than zero, as in the linear case). Notice that r_0 is an increasing function of the carry rate, s, and does not depend on the manager's ability, α. Figure 1 plots the expected post-carry net return as a function of the pre-carry expected return and illustrates the difference the option contract makes. As we will show, the contract fees will adjust to ensure that the set of states in which the fund is launched remains efficient.

4.2 Informed Investor's Problem

With linear carry, the informed investor's expected net return, and therefore her optimal capital commitment, varied linearly with α. With option-based carry, expected net returns are concave in α, so the informed investor's capital commitment is no longer linear. As a result, we can no longer solve directly for the manager's inference and the associated fundraising cap as functions of the informed investor's quantity choice.

We therefore proceed differently from the linear case. We begin by assuming the informed investor follows an investment schedule $q_i(\alpha)$ that is strictly increasing in managerial ability, $q_i'(\alpha) > 0$ whenever $q_i(\alpha) > 0$, with $q_i(0) = 0$. Given such a schedule, the manager can infer α by inverting q_i and can then set the fundraising cap $\bar{q}(q_i)$ equal to the efficient fund size. We then characterize the schedule $q_i(\cdot)$ and show that it is incentive compatible for the informed investor to follow this schedule in equilibrium. Once that schedule is determined,

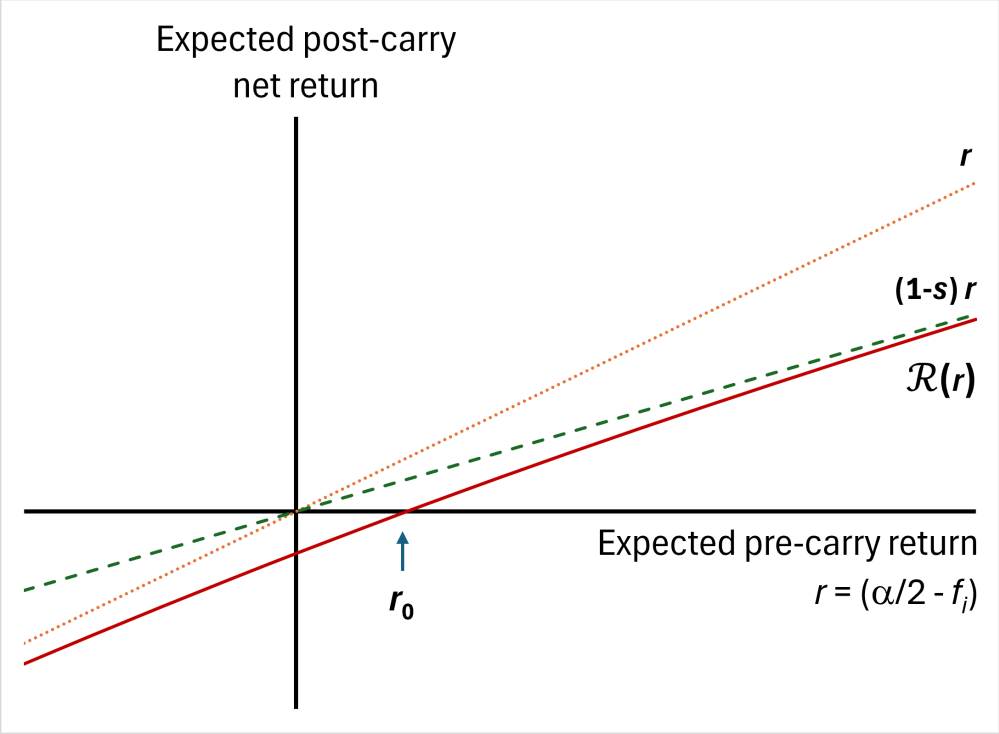


Figure 1: Break-even Expected Pre-Carry Return r_0. The dotted line r is the pre-carry expected return. The dashed line is the expected (post-carry) net return in the linear case. The solid line is the expected net return in the option-based carry case (with hurdle rate normalized to zero). While the optimal investment policy in the linear case is to invest when ever the expected pre-carry return exceeds zero, with option-based carry it is only optimal to invest when the expected pre-carry return exceeds r_0.

the remaining contract terms can be chosen to implement the efficient allocation and extract the rents.[21]

To determine the informed investor's incentives, consider such an equilibrium investment schedule $q_i(\alpha)$. Suppose that after learning the true α, the informed investor deviates and instead communicates $\hat{\alpha}$ by choosing to invest $q_i(\hat{\alpha}) > 0$. In that case, from Proposition 1, the manager will infer that the efficient size of the fund is $\hat{\alpha}/(2\beta)$. Assuming the manager then runs the fund at that inferred efficient size, the informed investor's expected pre-carry return is

$$\alpha - \beta\left(\frac{\hat{\alpha}}{2\beta}\right) - f_i = \alpha - \frac{\hat{\alpha}}{2} - f_i.$$

The informed investor will choose $\hat{\alpha}$ to maximize her net investment payoff:

$$\max_{\hat{\alpha}} \quad q_i(\hat{\alpha})\mathcal{R}(\alpha - \hat{\alpha}/2 - f_i). \tag{44}$$

Incentive compatibility requires that (44) be maximized when $\hat{\alpha} = \alpha$.[22] The first order condition is therefore

$$q_i'(\alpha)\mathcal{R}(\alpha/2 - f_i) - \frac{1}{2}q_i(\alpha)\mathcal{R}'(\alpha/2 - f_i) = 0. \tag{45}$$

When $q_i(\alpha) > 0$, we can equivalently rewrite (45) as follows:

$$\frac{q_i'(\alpha)}{q_i(\alpha)} = \frac{1}{2}\frac{\mathcal{R}'(\alpha/2 - f_i)}{\mathcal{R}(\alpha/2 - f_i)}.$$

Integrating both sides in the range when $q_i > 0$, we get $\ln q_i(\alpha) = \ln \mathcal{R}(\alpha/2 - f_i) + \ln \theta$, where $\ln \theta$ is an arbitrary constant of integration. Exponentiating both sides then gives

$$q_i(\alpha) = \theta\ \mathcal{R}(\alpha/2 - f_i). \tag{46}$$

For future ease of comparison with the linear case, we rewrite the constant of integration in terms of k, defined as,

$$k \equiv \frac{1}{\beta\theta(1-s)}. \tag{47}$$

[21]This part of the model can be viewed either as a signaling problem or as a screening mechanism in which the manager commits ex ante to the mapping from the informed investor's capital commitment to total fundraising. Methodologically, however, it is closer to signaling, as the key object is the incentive-compatible capital-commitment schedule $q_i(\alpha)$ and the relevant deviations are deviations in that signal.

[22]Given $s < 1$, the objective in (44) satisfies the Spence-Mirrlees single-crossing property $\frac{\partial^2}{\partial q_i \partial \alpha} q_i \mathcal{R} = \mathcal{R}' > 0$. Thus, we can use the approach of Mailath (1987) to characterize the separating equilibrium in our setting with continuous types, which we follow here.

Substituting this expression leads to the following general form for any incentive compatible investment schedule:

$$q_i(\alpha) = \begin{cases} \frac{\mathcal{R}(\alpha/2 - f_i)}{\beta k(1-s)} & \text{if } \mathcal{R}(\alpha/2 - f_i) > 0 \\ 0 & \text{if } \mathcal{R}(\alpha/2 - f_i) \leq 0 \end{cases} \tag{48}$$

where the scale factor $k > 0$, the arbitrary constant of integration, is still to be determined.[23]

Equation (48) states that for any contract in which the informed investor reveals α and the manager implements the efficient investment level $\alpha/2\beta$, the informed investor commits capital proportional to her expected net return. Note also that the investment threshold α_I is defined by

$$\mathcal{R}(\alpha_I/2 - f_i) = 0 = \mathcal{R}(r_0),$$

and therefore,

$$\frac{\alpha_I}{2} - f_i = r_0. \tag{49}$$

The left side of (49) is the expected return of the fund, $\alpha_I/2$, less the management fee, that is, the informed investor's pre-carry expected return. So the investment threshold is the level of managerial ability at which the investor's pre-carry expected return is equal to r_0, the break-even pre-carry expected return.[24]

4.3 Manager's Problem

Having characterized the incentive-compatible investment schedules of the informed investor, we now turn to the manager's problem of choosing contract terms that extract the first-best surplus. To do so, the manager again commits to the following upfront:

- A management fee of f_i per dollar invested by the informed investor,
- A fractional share ("carry") $s \in [0, 1)$ of the net return in excess of a hurdle rate $h \geq 0$ that the manager will retain;

[23]Let $\alpha_I = \inf\{\alpha : q_i(\alpha) > 0\}$. Because the informed must earn a non-negative return, continuity and incentive compatibility imply that the marginal participating investor earns zero net return, so $\mathcal{R}(\alpha_I/2 - f_i) = 0$. Otherwise, a slightly lower type could profitably imitate a participating type.

[24]To see that this model collapses to the linear model, note that in the linear model, $\mathcal{R}(\alpha/2 - f_i) = (1-s)(\alpha/2 - f_i)$. Hence in that case, when $\alpha > \alpha_I = 2f_i$, (48) becomes

$$q_i(\alpha) = \frac{(1-s)(\alpha/2 - f_i)}{\beta k(1-s)} = \frac{\alpha - 2f_i}{2\beta k} = \frac{\alpha - \alpha_I}{2\beta k}. \tag{50}$$

matching (14). (49) matches the linear case, because $r_0 = 0$.

- A cap $\bar{q}(\cdot)$ on overall fundraising that depends on the amount of capital q_i that the informed investor commits.

First, we must ensure the informed investor's investment threshold α_I matches the efficient level α_M defined in (7). From (49), this requires

$$f_i = \frac{\alpha_M}{2} - r_0. \tag{51}$$

Because the schedule $q_i(\alpha)$ derived in (48) is incentive compatible, the manager and other investors can invert the schedule to assess the informed investor's determination of α based on his capital commitment. The cap on total capital is then set to the efficient level:

$$\bar{q}(q_i) \equiv \frac{q_i^{-1}(q_i)}{2\beta} = \frac{\alpha}{2\beta}. \tag{52}$$

Note that because q_i increases with alpha, $\bar{q}$ will increase with q_i. To compare with the linear model, we can compute the slope of $\bar{q}$ using the definition of $\mathcal{R}$ in (41) as follows:

$$\begin{aligned} \bar{q}'(q_i) &= \frac{1}{2\beta q_i'(\alpha)} = \frac{k(1-s)}{\mathcal{R}'(\alpha/2 - f_i)} = \frac{k(1-s)}{1 - s\Pr(\alpha/2 - f_i + \epsilon > h)} \\ &= \frac{k(1-s)}{1 - s\Pr(\text{carry}|\alpha)} \end{aligned} \tag{53}$$

Therefore, $\bar{q}$ is convex in q_i, with a slope that increases and approaches k as the likelihood that carry will be charged approaches one (matching the linear model).

Uninformed investors, upon observing the capital cap $\bar{q}$, can also infer $\alpha = 2\beta\bar{q}$ and will therefore correctly anticipate a gross return of the fund equal to $\alpha - \beta(\alpha/(2\beta)) = \alpha/2$. Given that, the manager will set the fee f_u to make the uninformed investors indifferent. That is, the uninformed expected net return will be set to zero, implying that the uninformed investors' pre-carry expected return must be set equal to r_0, the break-even pre-carry expected return:

$$\frac{\alpha}{2} - f_u = r_0,$$

implying

$$f_u = \frac{\alpha}{2} - r_0. \tag{54}$$

Note that, like f_i, the uninformed fee is similarly reduced by r_0, the break-even expected carry. Consequently, the difference in management fees,

$$f_u - f_i = \frac{\alpha - \alpha_M}{2},$$

is the same as in the linear case, and because $\alpha > \alpha_M$ if the fund launches, we have $f_u > f_i$.

Thus, given any parameters (s, h, k), by choosing $(f_i, \bar{q}, f_u)$ according to (51), (52), and (54), the manager will be able to raise funds equal to the efficient level of capital. The remaining condition for efficiency is that the informed investor finds it worthwhile to investigate. The informed investor's profit upon investigating is given by

$$\pi_i(\alpha) = \begin{cases} q_i(\alpha)\mathcal{R}(\alpha/2 - f_i) = \frac{\mathcal{R}(\alpha/2-f_i)^2}{\beta k(1-s)} & \text{if } \alpha > \alpha_M \\ 0 & \text{if } \alpha \leq \alpha_M \end{cases} \tag{55}$$

Therefore, for a given option-based carry (s, h), we can choose k so that the expected profit of the informed investor from investigating equals the investigation cost c_T:[25]

$$E\left[\pi_i(\alpha)\right] = \frac{1}{\beta k(1-s)} E\left[\mathcal{R}(\alpha/2 - f_i)^2 \, \mathbf{1}_{\{\alpha > \alpha_M\}}\right] = c_T \tag{56}$$

The final step is to show that, in this equilibrium, the manager will indeed choose to accept capital until the cap, $\bar{q} = \alpha/2\beta$, is reached. As in the linear model, the manager will find it optimal to do so. Because the manager earns all rents attributed to the uninformed investors, the manager's expected profit is equal to the total value added of the fund less the amount captured by the informed investor. If the total capital is below the efficient level — i.e., the cap — then each dollar raised both improves efficiency, and so increases the total value added of the fund, while it reduces the net return earned by the informed investor, both of which improve the manager's payoff. We prove this formally as part of the following proposition:

Proposition 5 (First-Best Option Contract)**.** *For a given (s, h), if the manager chooses $(k, f_i, \bar{q}, f_u)$ according to (51)-(56), then the informed investor will choose to investigate and invest according to (48), and the manager will raise uninformed capital until the cap $\bar{q}$ binds at the first-best allocation. This contract is optimal for the manager, who will earn the first-best surplus (31).*

Proof of Proposition 5. The essence of the proof is contained within the main text. Here we formalize two details. First, we verify the global optimality of the informed agent's optimal investment choice. Recall that the first-order condition (45) holds by construction of $q_i(\alpha)$. To see that the first-order condition is sufficient, note that $\mathcal{R}$ is increasing and concave, and therefore q_i, which is proportional to $\mathcal{R}$, is also increasing and concave. The second derivative

[25]Note that we continue to assume $\hat{\alpha}_0 < \alpha_M$, so that the investor would never invest without investigating first. (Indeed, with option-based carry, an investor who does not investigate earns an even lower expected payoff than $\pi_i(\hat{\alpha}_0)$, since their higher residual uncertainty raises the expected cost of the carry option.)

with respect to $\hat{\alpha}$ of the objective in (44) is then:

$$q_i''\mathcal{R} - q_i'\mathcal{R}' + \tfrac{1}{4}q_i\mathcal{R}'' < 0.$$

Second, we verify that the manager will choose raise capital until $\bar{q}$ binds. Here we use the fact that, because the uninformed earn zero rents, the manager's expected payoff is equal to the total surplus net of the informed agent's profit:

$$\pi_m = (q_i + q_u)(\alpha - \beta(q_i + q_u)) - \delta - \pi_i.$$

Therefore,

$$\tfrac{\partial}{\partial q_u}\pi_m = (\alpha - 2\beta(q_i + q_u)) - \tfrac{\partial}{\partial q_u}\pi_i.$$

Because higher q_u reduces the expected return for the informed investor, the second term, $-\frac{\partial}{\partial q_u}\pi_i$, is strictly positive. The first term is weakly positive as long as $q_i + q_u \leq \alpha/(2\beta)$. Therefore, the cap $\bar{q} = \alpha/(2\beta)$ strictly binds. □

4.4 Quantitative Implications

Having characterized the option-based contract analytically, we now illustrate its quantitative implications with a simple example. We identify the main parameters as follows. First, we consider a minimum fund size $\bar{q}(0) = \$50$ million, with manager opportunity and operating cost $\delta = \$1$ million per year. Because the minimum efficient alpha is α_M, the minimum fund size must equal $\frac{\alpha_M}{2\beta}$, with total value added $\bar{q}(0)\alpha_M/2 = \delta$. Therefore, we can calibrate

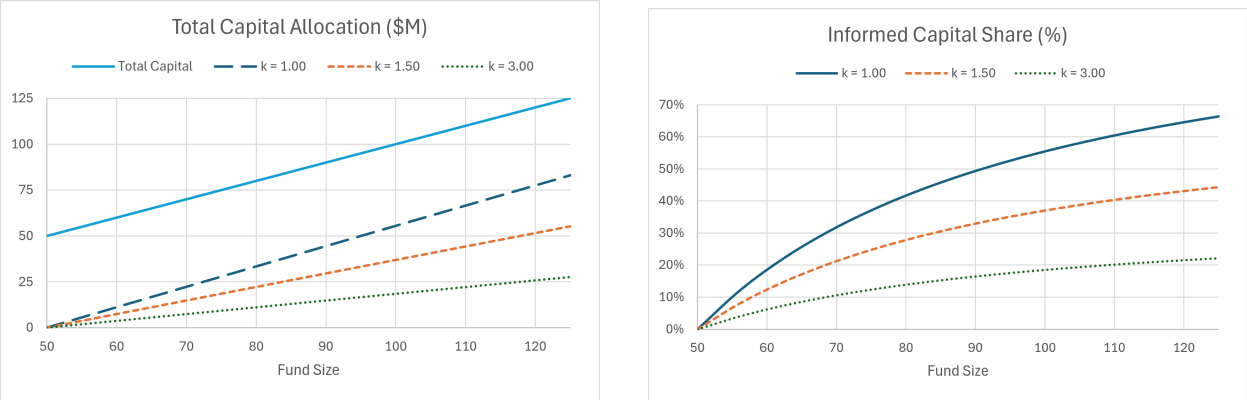
$$\alpha_M = \frac{2\delta}{\bar{q}(0)} = 4\% \quad \text{and} \quad \beta = \frac{\alpha_M}{2\bar{q}(0)} = 0.04\%.$$

Next we assume that the idiosyncratic risk $\epsilon \sim N(0, \sigma_\epsilon^2)$. We begin with the standard carry rate of 20%, idiosyncratic volatility of 15% and a hurdle rate of 0% which implies an r_0 (using (43)) of 1.335%. The optimal informed management fee is $f_i = \alpha_M/2 - r_0 = 0.665\%$. Table 1 computes r_0 for different carry rates and idiosyncratic volatilities.

Next, we consider the realized capital allocation. Figure 2 plots the capital allocation as a function of fund size, which scales with the manager's skill α (recall $\bar{q} = \alpha/(2\beta)$). Given $s = 20\%$, the division of capital between informed and uninformed investors varies based on the co-investment factor $k = 1, 1.5, 3$. When $k = 1$, then the level of uninformed capital approaches a constant for large α (see (53)). (Note that, as in the linear case, when $k < 1$ uninformed capital is decreasing in q_i.) For $k > 1$, the level of uninformed capital increases as q_i rises, limiting the rents of the informed investor.

	σ_ϵ		
s	10%	15%	20%
10%	0.420	0.630	0.841
15%	0.648	0.972	1.297
20%	0.890	1.335	1.780

Table 1: Implied break-even pre-carry expected return r_0 (in percent) given carry s, zero hurdle, and idiosyncratic risk σ_ϵ.



(a) Total Capital Allocation

(b) Informed Share

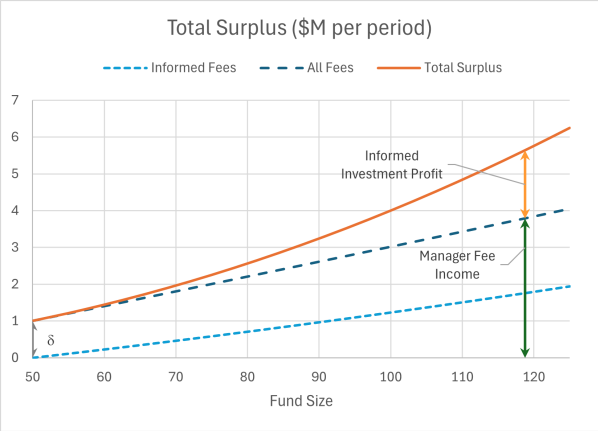
Figure 2: Total Capital Allocation and Informed Capital Share versus Fund Size

Finally, Figure 3 plots the total surplus, fee income, gross alpha, and net alpha as a function of realized fund size for the case $k = 1$ and $s = 20\%$. At the minimum fund size, the total fees cover the fund's operating costs δ. For larger funds, the manager and the informed investor share the net surplus. The expected investment profit of the informed investor covers the upfront investigation cost.

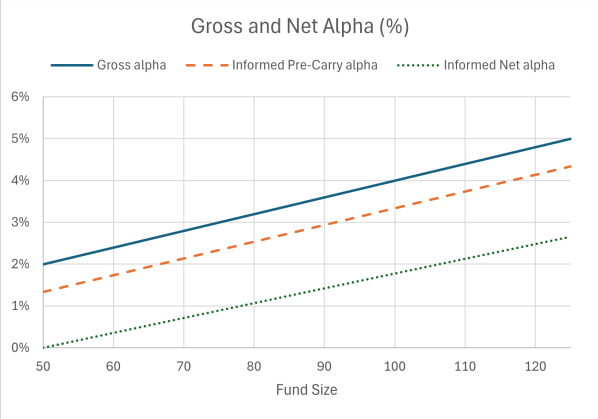
To calculate this expected profit, we need to assume a prior distribution of investment skill α. Table 2 does so for a normal prior with mean $\hat{\alpha}_0 = 2\%$–4% and $\sigma_\alpha = 2\%$–6%.

	$\hat{\alpha}_0$		
σ_α	2.0%	3.0%	4.0%
2%	0.019	0.052	0.123
4%	0.206	0.323	0.489
6%	0.622	0.830	1.090

Table 2: Expected informed-investor profit $E[\pi_i(\alpha)]$ (in \$ millions per year) given prior mean $\hat{\alpha}_0$ and prior standard deviation σ_α.



(a) Total Surplus and Fees



(b) Gross and Net Alpha

Figure 3: Total Surplus, Fee Income, and Gross and Net Alpha versus Fund Size

5 Why Delegation is Necessary

The money management industry is characterized by delegation, that is, investors pool their capital with other investors and delegate the management of that capital to a skilled third party manager. This naturally raises the question of why other forms of market organization are largely absent. In this section we answer this question by demonstrating that the two most obvious alternative ways to organize the industry lead to inefficient outcomes.

5.1 Bilateral Contracting

At first glance, uninformed investors seem unnecessary in our setting. The manager has the skill, and the informed investor has the information needed to determine whether the fund should be launched and, if so, at what scale. Moreover, we do not assume that informed investors are capital constrained — so why are uninformed investors necessary at all? Put differently, why can't the informed investor simply hire the manager and provide all the capital?

The answer is that a purely bilateral relationship asks the informed investor's capital commitment to do two jobs at once. It must reveal what she knows — she has to be willing to commit more when the manager is better — and it must also determine the size of the fund. Under delegation these are separate instruments: her stake reveals α, while uninformed capital fills the fund out to its efficient size. Bilaterally, they are the same number, and the manager cannot use it for both purposes at once.

The conflict is easiest to see from the informed investor's side. Her report determines the manager's compensation as well as the fund's size. If the manager is paid more when she

reports a better manager, she will shade her report downward. Shading costs her almost nothing on her own position — she was at her optimum, so the loss is second order — but it saves her the manager's fee, which is first order. The only way to make shading expensive is to make investment respond to her report, which means launching fewer funds than is efficient and making each one smaller than it should be. Manager rents and efficient investment are therefore incompatible.

Proposition 6 (Bilateral Inefficiency)**.** *Suppose the return shock ϵ is normally distributed, and consider any contract between the manager and the informed investor alone in which the manager's compensation is nonnegative and the investor's payoff is nondecreasing in the fund's realized return. Then:*

(i) If the contract implements first-best investment, the manager's expected rent is zero.

(ii) If the manager's expected rent is positive, investment is distorted on both margins: there is a funding cutoff $\alpha_I > \alpha_M$, so that managers who should be funded are not, and every fund that is launched is strictly smaller than first best.

(iii) An increase in return volatility σ_ϵ strictly contracts the set of attainable payoffs wherever the manager earns positive rents, leaving the zero-rent, first-best point unchanged.

Appendix C proves these results in a more general setting, for any return distribution with a nondecreasing hazard rate. Normality is assumed here only for concreteness: it guarantees both the unbounded upside that delivers (i) and the downside risk that delivers (ii).

Part (i) is stronger than the screening logic above suggests — it says that at the first best the manager earns not merely small rents but exactly zero. The reason is worth spelling out. Limited liability means the manager's compensation is a nonnegative, nondecreasing claim on the fund's return, so it is option-like. To preserve truth-telling by the investor, the manager's rents can grow with skill only to the extent that his claim is sensitive to the return, as it is this sensitivity the makes it costly for the investor to under-report. But at the efficient threshold α_M the fund generates no surplus, so the claim must be worthless there; and since returns are unbounded, the only worthless claim is literally the zero security. That security provides no deterrence at all, and compensation can grow only in proportion to the deterrence it already provides — so, starting from zero, it never lifts off. Part (iii) has the same source. Noisier returns make a claim of given value less sensitive to the fund's performance, so the same expected payment to the manager buys less exposure to skill, and more distortion is needed to deliver the same rent.

Uninformed capital breaks the link. Because the fund's size is $q_0 + k\,q_i$ rather than q_i, the manager can shrink the informed investor's stake to whatever the incentive constraint

requires while holding the fund at its efficient size, and his rents come from fees on the whole fund rather than from her stake alone. In the optimal delegated contract the informed stake falls to zero as $\alpha \to \alpha_M$, while the fund is still launched at $q_0 = \alpha_M/2\beta$.

It is worth being clear that this is not merely a matter of letting the informed investor buy into the uninformed pool on the same terms. Suppose the manager allowed her to supply additional capital up to the cap $\bar{q}(q_i)$ at the uninformed terms. Since the uninformed break even, expected payoffs would be unchanged — but the contract would no longer be incentive compatible. Once she holds any part of the uninformed allocation, she gains by distorting her commitment downward, say to $q_i(\alpha - \eta)$: the loss on her informed position is second order, while the manager infers a lower α and sets the uninformed fee below break-even, so she earns a first-order gain on the rest of her holding.[26] The separation is therefore not merely that uninformed capital exists, but that the informed investor is excluded from it: her report moves the terms offered to others, and she captures none of that movement. It is the fact that uninformed investors free-ride on the information that imposes the extra cost on the informed investor that then disciplines truthful reporting. By distorting her commitment downward, she would allow uninformed investors to earn positive rents, a cost that is absent when uninformed investors do not participate.

Figure 4 compares the two. It shows, for each contracting framework, the highest possible expected payoff for the manager as a function of the expected payoff to the investor (not including the ex ante investigation cost). Because delegation achieves the first best for any division of surplus, its frontier is a straight line; the bilateral frontier lies strictly below it except at the single point where the manager earns nothing. Two features stand out. At the bilateral contract most favorable to the manager, he receives 1.27 against the 1.75 that delegation would deliver at the same payoff to the informed investor — roughly a third more — and about a fifth of first-best surplus is destroyed. And the informed investor cannot be pushed to zero: she retains a strictly positive payoff even when the manager has all the bargaining power, because there is a minimum information rent she can always extract. The larger part of the loss comes from the extensive margin: as reported in Appendix C, funds that are never launched account for several times the surplus lost to undersized funds. The bottom line of this figure is that unless the investigation cost is so large that all the rents must be given to the investor, the manager will always strictly prefer the delegated contract.

[26]The uninformed fee would fall to $\hat{f}_u = \hat{\alpha}/2 - r_0$, so the expected pre-carry return on the uninformed allocation would rise to $\alpha - \beta\hat{\alpha}/(2\beta) - \hat{f}_u = r_0 + \eta$, and the uninformed would earn $\mathcal{R}(r_0 + \eta) > 0$ per dollar.

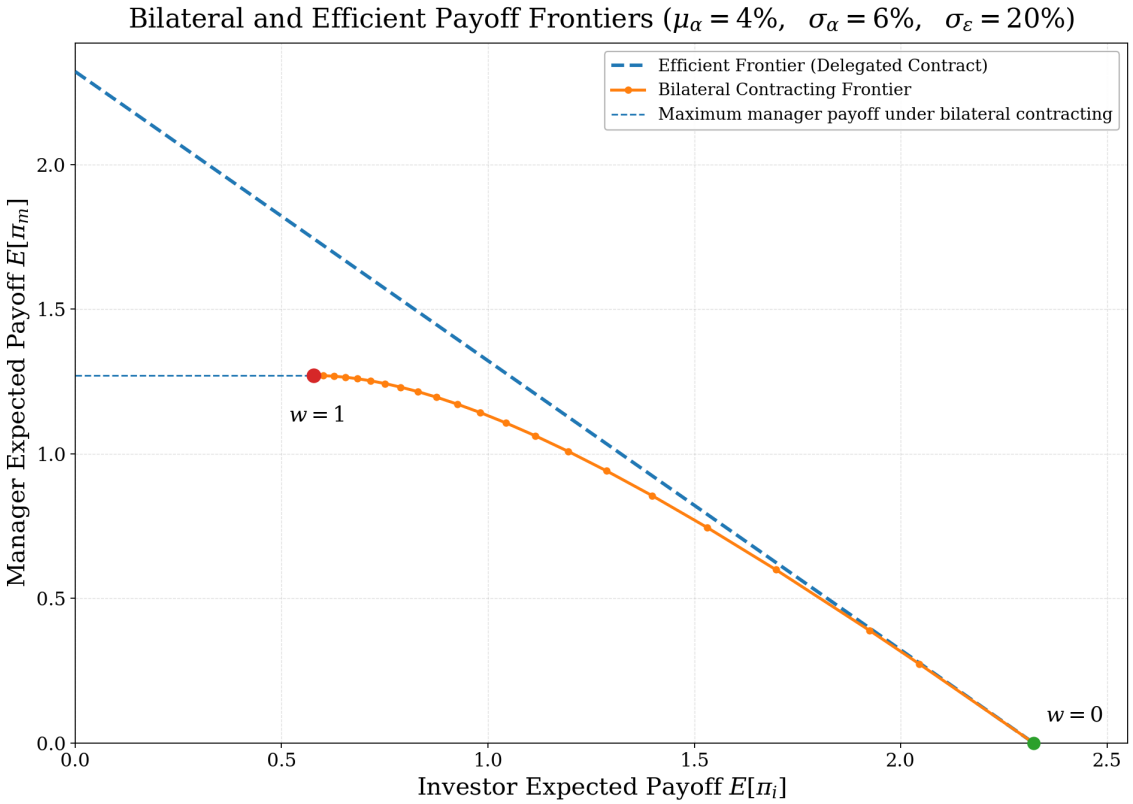


Figure 4: Bilateral and Efficient Payoff Frontiers. The dashed blue line is the efficient frontier, along which total surplus equals its first-best value $E[S^{FB}(\alpha)] = 2.322$; by Proposition 4 the delegated contract attains every point on it. Points to the left of an investor payoff of $c^* = 1.125$ are implemented by the standard contract with $k \geq 1$; delivering more than c^* to the informed investor requires $k < 1$ and the fee schedule of Section 3.3. The orange curve is the bilateral payoff frontier, traced by the weight w on manager rents. At $w = 0$ investment is efficient and the manager earns zero rents (Proposition 9), so the two frontiers meet; for every $w > 0$ the bilateral frontier lies strictly inside, reflecting both the exclusion of types above α_M and underinvestment in the funds that are launched (Proposition 10). At $w = 1$ the manager's expected payoff is maximized at 1.271, leaving the informed investor 0.577; the horizontal dashed segment marks this maximum. The figure uses $\mu_\alpha = 4\%$, $\sigma_\alpha = 6\%$, and $\sigma_\epsilon = 20\%$.

5.2 Why Not Certification?

The informed investor in our model performs two functions: she produces information, and she supplies capital. A natural question is whether these functions could be unbundled. Suppose the manager instead hires a *certifier*, an agent with the same information technology as the informed investor, who investigates at cost c_T and reports α. The manager then raises the entire fund from uninformed investors at the certified terms. The certifier supplies no capital, and like the manager she is protected by limited liability: her compensation must be nonnegative.

Proposition 7 (No Certification)**.** *Suppose the return shock ϵ is normally distributed and the certifier's compensation is a nonnegative claim on the fund's cash flows, so that she receives nothing when no fund is launched. Then no compensation scheme induces the certifier to acquire information.*

The proof is the zero-security argument of Proposition 6(i) in a new role. Efficiency requires that bad news end the process, so a truthful bad-news report pays nothing. To keep bad-news types from faking good news, every launch claim must also be worthless to them, and with unbounded returns the only worthless claim is the zero security — which pays every type nothing. A certifier paid from the fund cannot be paid for saying no, and an agent who cannot be paid for saying no will not pay to find out whether the answer is no.

Could the manager fix this with a retainer, a fee paid up front or upon a no-launch report? Committed cash for bad news solves one problem but creates another. The certifier's temptation is no longer to fake good news; it is to stop investigating. If she skips the investigation, she can simply pick the report that pays best on average and collect. Here the option-like shape of her compensation, which limited liability makes unavoidable, works against the arrangement in a way it never does for an investor. An option is worth more when the outcome is riskier, and to an uninformed certifier the manager's skill is part of the risk: not knowing α makes her claim more valuable, not less. Deterring her therefore requires paying an informed certifier a premium on top of the investigation cost itself, and the premium is largest exactly where information is subtle. When skill is dispersed, a certifier who does the work can identify reports far from the mean that an uninformed one would not dare make, and the premium is modest. When managers are hard to tell apart, the informed and uninformed certifier hold nearly the same lottery, so the reward for genuine knowledge must be enormous to be detectable. The total cost of certification — investigation plus premium — then quickly exceeds the entire surplus the fund can generate. In our calibration, certification becomes infeasible once investigation costs reach at most a fifth of first-best surplus, whereas an informed investor remains viable at any cost below the full surplus.

The informed investor escapes both problems at once, and for the same reason: her stake. Bad news pays her automatically, because she keeps her capital rather than committing it to a losing fund, so no one must fund a reward for saying no. Skipping the investigation holds no temptation either, since an uninvestigated commitment is her own money at risk rather than a free option on someone else's fund. Skin in the game is not an incidental feature of how information is produced in this market; it is what makes information cheap to buy. The natural industry structure ties information production to capital provision because capital is the one payment that rewards bad news and punishes ignorance without anyone having to write those contingencies down.

6 Discussion

The central implication of the model is that the delegated money management industry is optimally bifurcated. In public markets, where underlying assets are liquid and performance is informative, investors can learn about managerial ability from public signals over time and reallocate capital in response. In that environment, it is inefficient to incur the cost of ex ante due diligence, and the mutual fund contract, in which investors face common terms and managers are compensated through a fee based on assets under management, implements the efficient allocation. In private markets, by contrast, the illiquidity and opacity of the underlying investments make realized performance a less effective mechanism for revealing skill and adjusting capital. As a result, some investors optimally acquire information in advance, and once they do, efficiency requires a different contractual structure: investors cannot all face the same terms, and the optimal contract includes both management fees and carry, together with limits on the scale of fundraising. The contractual differences between public and private delegated management are therefore not incidental institutional details, but the equilibrium response to different information environments.

6.1 Alpha, Persistence, and Investor Heterogeneity

As in the competitive benchmark for mutual funds, our model implies that managers extract the rents once the costs of information have been accounted for. However, that implication does not mean that all investors in private markets earn zero alpha. While uninformed investors earn zero alpha, informed investors earn positive ex post alpha because they incur the cost of acquiring and validating information about managerial skill. One might naively infer that the positive alpha reflects information rents, but in fact these rents are spent acquiring the information in the first place. This result mirrors its counterpart in securities

markets: in the Grossman and Stiglitz (1980) equilibrium, informed traders earn excess returns that exactly compensate them for the cost of acquiring information. Just as in that equilibrium, informed investors in our model are competitively supplied and face the same investigation cost, so the positive ex post alpha exactly compensates them for that cost and no informed investor outperforms on a net basis.

A simple extension of our model changes this conclusion in a natural way. If investors differ in their costs of acquiring information, those with the lowest costs earn true excess returns, while only the marginal informed investor breaks even. Moreover, investors with lower investigation costs always choose to become informed, bifurcating the investor population and producing the striking finding of Begenau and Siriwardane (2024): fee discounts vary across investors rather than across funds.

More generally, when the capacity to conduct effective due diligence is itself scarce, no investor is marginal, and the rents associated with managerial ability will be shared between managers and informed investors, presumably as the outcome of a bargaining process. This equilibrium can be implemented in our setting by folding the outcome of that bargaining into the effective cost of inducing informed participation. That is, one can interpret C, the cost of acquiring information, as also including any information-acquisition rents. Importantly, this does not alter the efficiency result. The same contractual structure still implements the efficient allocation, with the additional rents accommodated through different values of the carry, participation factor, or hurdle rate rather than through a distortion in investment, provided the total effective cost of informed participation remains below the available surplus. When the effective cost exceeds c^* — the largest rent the manager can deliver through fee discounts on the informed investor's stake without displacing uninformed capital — the accommodation takes the form derived in Section 3.3: the fund tilts toward informed capital and the management fee is waived beyond the solo-investor threshold, with efficiency still intact. The model can therefore account not only for positive ex post alpha, but also for persistent outperformance by a subset of investors with a durable advantage in due diligence and manager selection.

More broadly, the model suggests a different interpretation of alpha in private markets. Positive alpha need not reflect a breakdown of competition or an unexplained anomaly. Instead, it can arise because illiquidity and opacity limit the ability of capital markets to aggregate information and reallocate capital efficiently through prices and flows alone. In that environment, informed investors can earn positive ex post alpha even though the overall equilibrium remains competitive.

6.2 Managerial Information About Own Skill

Both for simplicity and for consistency with Berk and Green (2004), we assume that neither the manager nor investors initially know the fund's quality. We use fund quality broadly. It reflects not only the manager's skill, but also the value and scalable capacity of the particular investment strategy in the prevailing market environment. A manager may understand his own capabilities and investment process without knowing how valuable the strategy will prove to be, how rapidly its returns will decline as capital expands, whether similar strategies are becoming crowded, or how changes in the macroeconomic environment will affect its performance. Investors that evaluate many managers may be better positioned to assess these questions. They can compare the manager with a broad cross-section of alternatives, observe how capital is flowing across strategies, and bring information about market conditions that the manager may not possess.

Managers may also assess their own strategies imperfectly. Familiarity with the strategy does not necessarily provide an objective assessment of its value, particularly when managers are overconfident about their ability or the capacity of their investment opportunities. Consistent with this interpretation, Berk, van Binsbergen and Liu (2017) find that a mutual fund manager's value added increases when fund executives take capital away. Because the manager could have indexed the marginal capital, this finding suggests that managers had overestimated the amount of capital they could deploy productively. The evidence therefore speaks especially to uncertainty about the scalability of the strategy, rather than simply to whether managers understand their own innate skill. Braun et al. (2026) find related evidence on capital allocation within private equity firms.

This informational assumption is what gives delegation an essential role in our model. Suppose instead that the manager already knew the fund's quality. If he could credibly communicate that information, outside due diligence would be unnecessary. If he could not communicate it but an investor could verify it through due diligence, then once the investor became informed both parties would know the fund's quality. A single informed investor could then contract bilaterally with the manager and finance the efficient level of investment. Our delegated contract would remain optimal and would continue to implement the first-best allocation, but it would no longer dominate bilateral contracting, which could achieve the same outcome. Because the investor is not capital constrained in our setting, uninformed capital would therefore serve no essential role. Thus, within our framework, delegation is uniquely valuable because investor diligence produces information about the strategy that the manager does not already possess. The prevalence of sophisticated investor due diligence, anchor investors, and multi-investor funds is consistent with this interpretation.

Finally, a manager who is privately informed about the fund's quality might attempt to

signal that information. The option-based contract cannot achieve separation on its own because it exposes the manager only to the fund's upside. A low-quality manager can therefore mimic a high-quality manager without risking his own wealth. A separating signal instead requires a commitment whose expected cost differs with fund quality, such as meaningful exposure of the manager's wealth to the fund's downside, as in the signaling-by-retention logic of Leland and Pyle (1977). Carry clawbacks provide only limited exposure of this kind because they generally require the manager to return previously accrued compensation rather than bear losses beyond that compensation. A sufficiently large and unhedged GP capital commitment could, in principle, serve as a signal, particularly if poor fund quality makes that commitment more costly in expectation. Such commitments need not fully resolve the informational problem, however. They may be too small relative to fund size, may be financed or hedged, and may not reveal the strategy's scalable capacity or its value in the prevailing market environment. The continued importance of investor due diligence even when managers invest alongside their funds is therefore consistent with managerial signaling being incomplete rather than eliminating the uncertainty that motivates our analysis.

6.3 Open Questions on Contract Design and Implementation

Our analysis identifies the contractual structure needed to implement the efficient allocation in private markets, but it does not fully pin down the exact combination of contract terms. In particular, as shown in Section 3.1 and in the option-contract extension that follows, there is an indeterminacy among the participation factor, the carry, and the hurdle rate. This indeterminacy is not about the division of rents. Rather, it reflects the fact that different combinations of these instruments can deliver the same equilibrium payoffs while implementing the same efficient allocation. Unmodeled considerations may therefore determine which contractual form is selected in practice. For example, if investors face liquidity risk or portfolio-concentration concerns, managers may prefer a higher participation factor and lower carry, since this allows the informed investor to be compensated while holding a smaller share of the fund. Conversely, stronger managerial incentive concerns, which we abstract from, may favor a higher carry and a lower participation factor. Tax considerations may also matter, since carry and management fees are taxed differently.

A second open question concerns the credibility of informed participation. In the baseline model, the informed investor takes only a very small position near the investment threshold. In practice, however, such a small investment may be neither economical nor credible as a signal of manager quality. As we show in Appendix B, once fixed costs of allocating capital are introduced, the informed investor optimally makes a meaningful minimum investment if she

invests at all. The logic is straightforward: when becoming informed is costly, a small position does not credibly signal that the investor has incurred those costs. A larger investment makes the signal more believable because it shows that the investor has committed meaningful capital on the basis of her information. This extension suggests that the substantial initial positions often taken by cornerstone investors may reflect not only scale, but also the need to make their information production and endorsement credible.

7 Conclusion

This paper extends the competitive paradigm of delegated money management from mutual funds to private markets. Under the assumptions of competitive capital markets and rational expectations, the same underlying logic continues to apply: after the costs of information have been accounted for, the equilibrium allocation is efficient and managers extract the rents. What differs across sectors is not the presence or absence of competition, but the way information is produced and transmitted.

The crucial distinction is that public-market investments are liquid and transparent, whereas private-market investments are illiquid and opaque. In public markets, investors can learn about managerial ability over time from realized performance and can reallocate capital in response, so it is inefficient to incur the cost of ex ante due diligence. In private markets, by contrast, illiquidity and opacity make realized performance a less effective mechanism for revealing skill and adjusting capital. As a result, some investors optimally acquire information in advance, and the efficient contractual response is different. The model therefore explains why private markets feature management fees together with carry, why different investors receive different terms, why fund-size limits are binding in equilibrium, and why the largest funds concede fee discounts to their anchor investors. It also reinterprets alpha in private markets: positive ex post alpha for informed investors is not an anomaly or a failure of competition, but the competitive return to information production in an environment where prices and flows cannot transmit it.

Finally, the model highlights an unexpected role for delegation itself. The manager and the informed investor together possess all the information needed to identify the efficient outcome, and uninformed investors are not needed to provide scarce capital. Yet bilateral contracting between the manager and the informed investor is generally inefficient, because the only way to deliver rents to the manager is to distort investment. Delegation changes the problem fundamentally: by bringing in uninformed investors on different terms, the manager can break the link between the informed investor's stake and the total scale of the fund, eliminate the bilateral information-rent distortion, and implement the first best. In this sense,

uninformed investors are essential not because they provide information or scarce capital, but because they make efficiency implementable. More broadly, our results show that when illiquidity and opacity prevent markets from aggregating information through prices and flows alone, both the optimal contract and the optimal structure of delegation must adjust to restore efficiency.

Appendix

A Option to Shut Down

This appendix characterizes the efficient operating and shutdown policy. At reset dates $0, T, 2T, \dots$ the fund's capital $q \geq 0$ can be reset and the fund may instead be shut down permanently (yielding payoff 0 thereafter). Between reset dates, capital is held fixed. If the fund continues, the operating cost $\delta > 0$ is incurred each period and beliefs about α evolve as data are observed.[27]

Full information. If α is learned at $t = 0$, the efficient policy operates the fund only if the optimized per-period surplus is nonnegative. Under the quadratic technology, this is equivalent to $\alpha \geq \alpha_M$, where

$$\alpha_M \equiv 2\sqrt{\beta\delta}.$$

The full-information value is

$$V_0^{FI} = \frac{1}{r} E\left[\left(\frac{\alpha^2}{4\beta} - \delta\right) \mathbf{1}_{\{\alpha > \alpha_M\}}\right].$$

Comparative statics in T

Let $V_0(T)$ denote the maximal expected discounted surplus when capital and the shutdown decision can be changed only at reset dates $\{0, T, 2T, \dots\}$ (so shutdown is not allowed within a block). Next we consider lengthening the illiquidity period by an integer multiple.

Proposition 8. ***(Value declines with a longer illiquidity interval.)*** *For any integer $m \geq 1$,*

$$V_0(T) \;\geq\; V_0(mT).$$

Proof. Fix any admissible policy under interval mT. Construct a feasible policy under interval T that makes the same shutdown/continue decisions at dates $0, mT, 2mT, \dots$ and, conditional on continuing, holds the same capital level fixed over each of the m finer sub-intervals within every coarse block. This policy is feasible because it simply ignores intermediate information and does not change decisions at intermediate reset dates. It generates identical cash flows, hence the same expected discounted payoff. Taking suprema over admissible policies yields $V_0(T) \geq V_0(mT)$. □

[27] We allow learning even if $q = 0$ (e.g., the fund can record trades without executing them). This simplifies notation; alternatively one could impose a minimum investment level.

Remark 1. Proposition 8 gives monotonicity along divisibility chains because $\{0, mT, 2mT, \dots\} \subseteq \{0, T, 2T, \dots\}$. A standard way to obtain monotonicity for all real T is to model T as a *capital lockup*: capital may be reset at stopping times $\tau_0 = 0 < \tau_1 < \dots$ satisfying $\tau_{k+1} - \tau_k \geq T$ (with shutdown allowed only at reset times). Larger T then imposes a tighter restriction on admissible policies, implying $V_0(T)$ is weakly decreasing in T without any distributional assumptions.

Value of information as T becomes large. Full information at $t = 0$ is a refinement of the information available under learning, so it weakly increases the maximal payoff. Proposition 8 implies that $V_0(T)$ weakly decreases as the reset interval is coarsened along multiples, so the value of acquiring information upfront is (weakly) increasing in T.

In the limiting case in which capital is never reset after $t = 0$ (formally, along $mT \to \infty$), the efficient policy either shuts down immediately or operates forever based on the initial posterior mean $\hat{\alpha}_0 \equiv E[\alpha \mid I_0]$. The corresponding *per-period* no-information surplus is

$$\left(\frac{\hat{\alpha}_0^2}{4\beta} - \delta\right) \mathbf{1}_{\{\hat{\alpha}_0 > \alpha_M\}}.$$

Thus acquiring information at cost C is efficient if and only if

$$rC \;<\; E\left[\left(\frac{\alpha^2}{4\beta} - \delta\right) \mathbf{1}_{\{\alpha > \alpha_M\}}\right] - \left(\frac{\hat{\alpha}_0^2}{4\beta} - \delta\right) \mathbf{1}_{\{\hat{\alpha}_0 > \alpha_M\}}. \tag{57}$$

When $\delta = 0$, condition (57) coincides with (5).

Computing $V_0(T)$ in the normal case

To determine the value of information when $T < \infty$, we must solve numerically a dynamic program for the optimal investment and continuation decision. Assume that at reset dates $t = nT$ the posterior satisfies $\alpha \mid I_{nT} \sim \mathcal{N}(\mu_n, s_n^2)$, where

$$\mu_n \equiv E[\alpha \mid I_{nT}], \qquad s_n^2 \equiv \mathrm{Var}(\alpha \mid I_{nT}).$$

Under Gaussian updating with per-period signal noise σ_ε,

$$s_{n+1}^2 = \left(\frac{1}{s_n^2} + \frac{T}{\sigma_\varepsilon^2}\right)^{-1}, \qquad \mu_{n+1} = \mu_n + \eta_n, \;\; \eta_n \sim \mathcal{N}\big(0,\; s_n^2 - s_{n+1}^2\big).$$

Define the optimized discounted payoff from operating for one reset interval as a function of the posterior mean,

$$g_T(\mu) \equiv \left(\frac{1-e^{-rT}}{r}\right)\left(\frac{\mu_+^2}{4\beta} - \delta\right), \qquad \mu_+ \equiv \max\{\mu, 0\},$$

and let $\gamma \equiv e^{-rT}$. The one-step recursion is

$$V(\mu_n, s_n^2) = \max\left\{0,\ g_T(\mu_n) + \gamma\, E\big[V(\mu_{n+1}, s_{n+1}^2) \mid \mu_n, s_n^2\big]\right\}.$$

Because s_n evolves deterministically in time, we can compute $V_0(T) = V(\mu_0, \sigma_0^2)$ by backward induction on a grid for μ, evaluating the conditional expectation with Gauss–Hermite quadrature.

B Extension: Minimum Informed Investment

In our optimal contract, the informed investor's stake, given by (14), approaches zero if α is close to $\alpha_I = \alpha_M$. In reality, a de minimus informed investment may be unrealistic — as well as unconvincing to uninformed investors. Indeed, it is likely that there are additional costs to the informed investor from deploying capital, and potentially from monitoring the manager to be sure the investment strategy is correctly implemented. There is also a potential opportunity cost of not being able to invest profitably elsewhere.

To incorporate this idea into the model, suppose that there is a fixed cost δ_i borne by the informed investor upon investing in the fund, and let δ_m be the fixed cost of the manager, so that the total cost of running the fund is $\delta = \delta_m + \delta_i$ and the total available surplus (31) is unchanged.

Let us continue to define $\alpha_M \equiv 2\sqrt{\beta\delta}$ to be the efficient investment threshold as in (7). Then, in the linear model, the optimal contract in this case will adjust as follows:

$$f_i = \frac{\alpha_M}{2} - \sqrt{\frac{\beta\delta_i k}{1-s}} \quad \text{and} \quad q_0 = \frac{f_i}{\beta}, \tag{58}$$

The informed investor will then invest according to

$$q_i(\alpha) = \left(\frac{\alpha/2 - f_i}{\beta k}\right)\ \mathbf{1}_{\{\alpha > \alpha_M\}}. \tag{59}$$

Then, at the investment threshold α_M, the minimum informed investment is

$$q_i(\alpha_M) = \sqrt{\frac{\delta_i}{\beta k(1-s)}}. \tag{60}$$

Note that we can use equation (60) to calibrate δ_i based on the minimum investment level for an "anchor" investor.

With these changes, the investment cap continues to provide the efficient level of capital,

$$\bar{q}(q_i(\alpha)) = q_0 + kq_i(\alpha) = \frac{f_i}{\beta} + \frac{\alpha/2 - f_i}{\beta} = \frac{\alpha}{2}, \tag{61}$$

and the minimum rent for the informed investor is equal to

$$\pi_i(\alpha_M) = q_i(\alpha_M)(1-s)(\alpha_M/2 - f_i) = \delta_i, \tag{62}$$

which also implies the manager earns $\alpha_M^2/(4\beta) - \delta_i = \delta_m$, as required.

Finally, the co-investment factor k must be set so that informed investor breaks even:

$$E\left[(\pi_i(\alpha) - \delta_i)^+\right] = c_T. \tag{63}$$

Compared to our earlier setting, compensating the informed investor for their opportunity cost will lead to a somewhat lower co-investment factor.

C Bilateral Contracting with an Informed Investor

This appendix studies bilateral contracting between a manager and a single informed investor. The investor acquires information about the manager's skill and supplies all of the capital; no uninformed investors participate. The main result is that any bilateral contract that gives the manager positive expected rents is necessarily inefficient, distorting investment relative to the first-best.

C.1 Environment and Admissible Contracts

Timing and Commitment. The manager commits to a menu of bilateral contracts before the investor acquires information. After observing the menu, the investor may incur the information cost c_T to investigate the manager. Investigation privately reveals the manager's skill α to the investor.

After observing α, the investor may either decline to finance the fund or select a contract from the manager's menu. Each contract specifies an investment scale q and a transfer to the manager contingent on the fund's realized payoff. If the investor declines, no investment occurs and both parties receive zero operating payoffs. If the investor accepts a contract, the investor supplies all of the capital, the fund operates, its payoff is realized, and the contractual transfer is made.

The manager can commit to the menu, but the investor cannot be committed to participate before acquiring information. This timing reflects the fact that an investor can investigate a manager before committing capital and can decline to invest after receiving unfavorable information. Thus, after learning α, the investor retains the option to reject the menu and receive an operating payoff of zero. Consequently, participation must hold conditional on the investor's information rather than only on an ex ante basis.[28]

Participation and Payoff Conventions. As before, let $\pi_i(\alpha)$ and $\pi_m(\alpha)$ denote the investor's and manager's expected operating payoffs conditional on the investor having learned α, where expectations are taken over the return shock. The investor can decline to finance the fund after observing α, so interim participation requires

$$\pi_i(\alpha) \geq 0 \qquad \text{for every } \alpha. \tag{64}$$

As in the main model, the investor must also be willing to acquire information. Because the prior mean of skill satisifes $\hat{\alpha}_0 \equiv E[\alpha] \leq \alpha_M$, the investor's optimal payoff without information is zero, and information acquisition requires $E[\pi_i(\alpha)] \geq c_T$. This condition determines whether the investor enters the bilateral contracting problem but does not affect the contract conditional on α having been learned, so we do not need to explicitly impose it below.

Fund Payoff. The investment technology is the same as in the main text. If $q = 0$, the fund is not operated and both parties receive zero operating payoffs. If $q > 0$, the fund earns the return

$$R = \alpha - \beta q + \epsilon, \tag{65}$$

[28]Because the investor observes α before choosing whether to participate, the appropriate individual-rationality condition is interim rather than ex ante. As is standard in such environments, we normalize the investor's outside option to zero and exclude unconditional transfers made before information is acquired. Such transfers would also be difficult to sustain in an extension with free entry by observationally similar potential managers, because they would attract low-skill entrants seeking the payment without generating corresponding investment surplus.

where ϵ is independent of α, has mean zero, and has a nondecreasing hazard rate on an interval support. We also impose two boundary conditions that play distinct roles:

- *Downside risk:* negative returns are possible for every skill and investment level.
- *Unbounded upside:* the support of ϵ is unbounded above.

Downside risk is essentially a no arbitrage condition that ensures that the manager cannot be given all of the marginal return exposure while preserving investor participation. Unbounded upside ensures that even a low-skill manager can realize any sufficiently high return, a property used later in establishing the impossibility of positive manager rents in an efficient contract. Both conditions are satisfied by the normal distribution used in our numerical examples.

We denote the fund's terminal value, net of the fixed operating cost δ, by

$$Y = q(1 + R) - \delta = q\big(1 + \alpha - \beta q + \epsilon\big) - \delta. \tag{66}$$

Because the investor supplies the initial capital q, the realized operating surplus shared by the investor and manager is $Y - q$.

Admissible Contracts. Each item in the manager's menu specifies an investment scale $q \geq 0$ and a transfer schedule $\tau(Y, q)$ from the investor to the manager. If the investor selects the contract, the manager receives $\tau(Y, q)$ and the investor receives the residual payoff

$$Y - q - \tau(Y, q). \tag{67}$$

We impose manager limited liability and require that neither party's payoff decline when the realized fund payoff increases. Thus, the manager cannot be required to make payments to the investor, and both the manager's transfer and the investor's residual claim must be nondecreasing in Y. We therefore restrict attention to transfer schedules satisfying

$$\tau(Y, q) \geq 0 \qquad \text{and} \qquad 0 \leq \tau_Y(Y, q) \leq 1, \tag{68}$$

where τ_Y denotes the derivative of the transfer with respect to Y.

First-Best Benchmark. For $q > 0$, total operating surplus is

$$S(q, \alpha) = E[Y - q \mid \alpha] = q\alpha - \beta q^2 - \delta, \tag{69}$$

with $S(0,\alpha) = 0$. The first-best investment rule is therefore

$$q^{FB}(\alpha) = \frac{\alpha}{2\beta}\mathbf{1}_{\{\alpha \geq \alpha_M\}}, \qquad \alpha_M = 2\sqrt{\beta\delta}. \tag{70}$$

Thus, α_M is the lowest skill level for which operating the fund generates nonnegative expected surplus. The resulting first-best surplus is

$$S^{FB}(\alpha) = \left(\frac{\alpha^2}{4\beta} - \delta\right)\mathbf{1}_{\{\alpha \geq \alpha_M\}}. \tag{71}$$

C.2 Incentive Compatibility and Information Rents

We next characterize the restrictions imposed by the investor's private information and the implications for the design of the manager's security. To do so, it is convenient to express the manager's payment as a compensation rate per dollar invested, contingent on the realized return.

Direct Mechanism in Return Space. By the revelation principle, we may represent the manager's menu as a direct mechanism indexed by the investor's report $\hat{\alpha}$. A report $\hat{\alpha}$ specifies an investment level $q(\hat{\alpha})$ and, when $q(\hat{\alpha}) > 0$, a per-dollar transfer schedule $t(R;\hat{\alpha})$ defined by

$$t(R;\hat{\alpha}) \equiv \frac{\tau\left(q(\hat{\alpha})(1+R) - \delta, q(\hat{\alpha})\right)}{q(\hat{\alpha})}. \tag{72}$$

Thus, the manager receives $q(\hat{\alpha})t(R;\hat{\alpha})$. For $q(\hat{\alpha}) > 0$, the admissibility conditions on τ are equivalent to

$$t(R;\hat{\alpha}) \geq 0 \qquad \text{and} \qquad 0 \leq t_R(R;\hat{\alpha}) \leq 1. \tag{73}$$

Incentive Compatibility. If the investor observes α but instead reports $\hat{\alpha}$, the realized return is $R = \alpha - \beta q(\hat{\alpha}) + \epsilon$. For $q(\hat{\alpha}) > 0$, the investor's expected operating payoff from this report is

$$\pi_i(\alpha;\hat{\alpha}) = q(\hat{\alpha})E\left[R - t(R;\hat{\alpha})\right] - \delta, \tag{74}$$

where the expectation is taken over ϵ and R is evaluated using the true skill α. A report for which $q(\hat{\alpha}) = 0$ yields an operating payoff of zero. Incentive compatibility requires truthful reporting to solve the investor's reporting problem:

$$\pi_i(\alpha) = \max_{\hat{\alpha}} \pi_i(\alpha;\hat{\alpha}) = \pi_i(\alpha;\alpha). \tag{75}$$

For types with $q(\alpha) > 0$, the envelope theorem implies, at points of differentiability,

$$\pi_i'(\alpha) = \left.\frac{\partial \pi_i(\alpha; \hat{\alpha})}{\partial \alpha}\right|_{\hat{\alpha}=\alpha} = q(\alpha) E\left[1 - t_R(R; \alpha)\right], \tag{76}$$

where the expectation is evaluated under truthful reporting. Define the "Delta" of the manager's security by

$$\Delta(\alpha) \equiv E\left[t_R(R; \alpha)\right] \in [0, 1]. \tag{77}$$

Because skill shifts the realized return one-for-one, $\Delta(\alpha)$ is both the expected Delta of the manager's security with respect to the fund return and, holding the selected contract fixed, the sensitivity of the manager's expected compensation rate to skill.

The envelope condition can therefore be written as

$$\pi_i'(\alpha) = q(\alpha)\big(1 - \Delta(\alpha)\big). \tag{78}$$

Thus, a marginal increase in skill raises the fund's expected payoff by $q(\alpha)$, of which the manager's security captures the fraction $\Delta(\alpha)$ and the investor retains the fraction $1 - \Delta(\alpha)$.

Investor Information Rents. Recall that α has interval support $[\underline{\alpha}, \overline{\alpha}]$, possibly unbounded, with prior mean $\hat{\alpha}_0 \leq \alpha_M < \overline{\alpha}$, so the investor can profitably operate the fund only by first acquiring information. Because $q(\alpha) \geq 0$ and $\Delta(\alpha) \leq 1$, the envelope condition (78) implies that the investor's payoff is nondecreasing in skill. The set of participating types is therefore an upper interval. Let α_I denote the lowest participating type. Because $\alpha_I \geq \alpha_M > \underline{\alpha}$ is interior, incentive compatibility and continuity imply

$$\pi_i(\alpha_I) = 0. \tag{79}$$

Indeed, if $\pi_i(\alpha_I) > 0$, a type just below α_I could select the contract intended for α_I and obtain a positive payoff, contradicting the definition of α_I. Hence, for $\alpha \geq \alpha_I$, integrating the envelope condition gives

$$\pi_i(\alpha) = \int_{\alpha_I}^{\alpha} q(a)\big(1 - \Delta(a)\big)\, da. \tag{80}$$

C.3 Security Design

Equation (80) shows how the security affects the cost of incentive compatibility. Starting from any given contract, a higher Delta reduces the rate at which the informed investor's payoff must increase with skill to induce truthful reporting. Thus, if we replace the manager's

security by another with the same expected payment but a higher Delta, the current allocation is unchanged, while less additional information rent must be given to higher-skill investors. Continuing the contract upward therefore requires lower investor rents and leaves more surplus available for manager compensation. This makes it useful to characterize, for any given expected payment to the manager, the highest Delta that an admissible security can provide.

Consider any admissible security t with expected payment $m = E[t]$. The next result shows that, for a given m, a call option delivers the highest possible Delta.

Lemma 2 (Security Frontier)**.** *Suppose ϵ has an absolutely continuous distribution F_ϵ with a nondecreasing hazard rate. For each $m \geq 0$, let $\kappa_0(m)$ be the centered strike of a call option on the return shock define by*

$$m = E\left[(\epsilon - \kappa_0(m))^+\right]. \tag{81}$$

Then every admissible security with expected payment m satisfies

$$\Delta \leq 1 - F_\epsilon\left(\kappa_0(m)\right) \equiv \overline{\Delta}(m). \tag{82}$$

The bound is attained by the call option on R with strike $E[R] + \kappa_0(m)$. The frontier $\overline{\Delta}(m)$ is increasing and concave, with $\overline{\Delta}(0) = 0$.

Proof. Let $\bar{R} = E[R]$ and rewrite the security as a function of the return shock $t_0(\epsilon) \equiv t(\bar{R}+\epsilon)$ so that

$$\Delta = \int t_0'(x) f_\epsilon(x)\, dx.$$

To maximize Delta while holding m fixed, note that we can rewrite m as

$$m = \int t_0(x)\, dF_\epsilon(x) = c + \int t_0'(x)\left[1 - F_\epsilon(x)\right]\, dx.$$

where the second equality follows from integration by parts and $c \geq 0$ is the minimum security payoff. Thus, adding slope at x has a compensation cost per unit of Delta equal to the inverse hazard rate

$$\frac{1 - F_\epsilon(x)}{f_\epsilon(x)}.$$

Because the hazard rate is nondecreasing, this cost is nonincreasing in x. Hence compensation is optimally used to assign slope first to the highest realizations of ϵ. Since admissibility requires $t_0' \in [0, 1]$, the maximizing security has slope one above a cutoff and zero below it, until the slope has reached one everywhere. This is precisely a call option, with centered strike $\kappa_0(m)$ determined by (81), and hence $\overline{\Delta}(m)$ is given by (82), which proves the bound.

Finally,

$$\frac{d\kappa_0}{dm} = -\frac{1}{[1 - F_\epsilon(\kappa_0)]}, \qquad \frac{d\overline{\Delta}}{dm} = \frac{f_\epsilon(\kappa_0)}{1 - F_\epsilon(\kappa_0)}. \tag{83}$$

Thus $\overline{\Delta}(m)$ is increasing. As m rises, $\kappa_0(m)$ falls, so the nondecreasing-hazard-rate assumption implies that $d\overline{\Delta}/dm$ falls. Hence $\overline{\Delta}(m)$ is concave. As $m \downarrow 0$, $\kappa_0(m)$ approaches the upper endpoint of the support of ϵ and therefore $\overline{\Delta}(m) \downarrow 0$. □

The next result shows that as long as there is downside risk, the manager's security must provide downside protection, implying $\Delta < 1$. The informed investor must therefore retain some exposure to returns, so truth-telling requires information rents that continue to rise with skill.

Corollary 1 (Strict Delta Bound)**.** *Suppose returns have downside risk. Then every feasible bilateral contract satisfies* $\Delta(\alpha) < 1$.

Proof. Suppose instead that $\Delta = E[t_R(R;\alpha)] = 1$ for some α. Because admissibility requires $0 \leq t_R(R;\alpha) \leq 1$, this implies that $t_R(R;\alpha) = 1$ almost surely. Hence, on the support of returns,

$$t(R;\alpha) = R - \kappa$$

for some constant κ. Limited liability, $t(R;\alpha) \geq 0$, and downside risk implies

$$\kappa \leq \inf \operatorname{supp}(R \mid \alpha) < 0.$$

The investor's expected payoff from the contract is therefore

$$q(\alpha)E[R - t(R;\alpha)] - \delta = q(\alpha)\kappa - \delta < 0,$$

contradicting participation. □

Thus, downside risk prevents the manager from capturing the entire marginal return to skill. Even a frontier security leaves the informed investor with strictly positive exposure $1 - \Delta(\alpha)$.

C.4 First-Best Investment and the No-Rent Result

The security frontier identifies the greatest possible exposure to skill that can be transferred to the manager for any given expected payment. It therefore provides the most effective available instrument for limiting the informed investor's information rents. We next ask

whether this instrument is sufficient to implement first-best investment while also providing rents to the manager.

The answer is no. At the efficient operating threshold α_M there are no rents to divide, so the manager's compensation begins at zero. The key incentive constraint is easiest to see using the centered strike, which is simply the strike on the return shock ϵ. Suppose an investor understates α by one percentage point. Relative to the contract intended for that lower type, his realized return then looks as though the shock ϵ were one percentage point higher. Thus, if the strike on ϵ rose by more than one percentage point as α fell, underreporting would more than offset the apparent improvement in the shock and reduce the manager's compensation. Because a nearby deviation from first-best investment entails only a second-order efficiency loss, such a deviation would be profitable. Incentive compatibility therefore requires the centered strike to rise by no more than one-for-one as α falls.

Whether compensation can nevertheless rise above zero depends on the upside of the return distribution. With bounded upside, a zero-valued option at α_M can have a finite strike at the upper endpoint $\overline{\epsilon}$ of the shock distribution. As α rises, that strike can gradually fall into the support and the option can acquire value. With unbounded upside, by contrast, zero compensation requires an infinite strike on ϵ. Since incentive compatibility allows that strike to fall only at a bounded rate as α rises, it cannot move from infinity to a finite level over any finite range of α. Manager compensation therefore remains zero for every type.

Proposition 9 (First-Best Strike Bound and No Manager Rents)**.** *Let $\overline{\epsilon} \leq \infty$ be the upper endpoint of the support of the return shock, and let $\kappa_0(m)$ denote the centered strike of the frontier call option defined in Lemma 2. Then any incentive-compatible bilateral contract that implements first-best investment satisfies, for every $\alpha \geq \alpha_M$,*

$$\kappa_0(m(\alpha)) \geq \overline{\epsilon} - (\alpha - \alpha_M). \tag{84}$$

If $\overline{\epsilon} < \infty$, it follows that

$$m(\alpha) \leq E\left[(\epsilon - \overline{\epsilon} + (\alpha - \alpha_M))^+\right]. \tag{85}$$

In particular, with unbounded upside the manager receives the zero security and earns zero expected payoff for every type.

Proof. At the efficient operating threshold α_M, first-best investment generates zero expected surplus. Investor participation and manager limited liability therefore imply

$$\pi_i(\alpha_M) = \pi_m(\alpha_M) = 0, \qquad m(\alpha_M) = 0.$$

Under first-best investment, the envelope theorem implies

$$\frac{d}{d\alpha}S^{FB}(\alpha) = q^{FB}(\alpha),$$

whereas incentive compatibility implies

$$\pi_i'(\alpha) = q(\alpha)\big(1 - \Delta(\alpha)\big).$$

Because $S^{FB} = \pi_i + \pi_m$ under first-best investment, it follows that

$$\pi_m'(\alpha) = q^{FB}(\alpha)\Delta(\alpha).$$

Using $\pi_m(\alpha) = q(\alpha)m(\alpha)$ and $q^{FB}(\alpha) = \alpha/(2\beta)$ gives

$$m'(\alpha) = \Delta(\alpha) - \frac{m(\alpha)}{\alpha}. \tag{86}$$

The centered strike depends on α only through $m(\alpha)$. For $m(\alpha) > 0$,

$$m(\alpha) = E\left[(\epsilon - \kappa_0(m(\alpha)))^+\right],$$

and hence

$$m'(\alpha) = -\overline{\Delta}\big(m(\alpha)\big)\frac{d}{d\alpha}\kappa_0(m(\alpha)).$$

Combining this expression with (86) gives

$$\frac{d}{d\alpha}\kappa_0(m(\alpha)) = -\frac{\Delta(\alpha)}{\overline{\Delta}\big(m(\alpha)\big)} + \frac{m(\alpha)}{\alpha\,\overline{\Delta}\big(m(\alpha)\big)}.$$

Since $\Delta(\alpha) \leq \overline{\Delta}\big(m(\alpha)\big)$, we obtain

$$\frac{d}{d\alpha}\kappa_0(m(\alpha)) \geq -1. \tag{87}$$

Thus, as α rises, the centered strike can fall by at most one-for-one with α.

When upside is bounded, $m(\alpha_M) = 0$ implies $\kappa_0(0) = \bar{\epsilon}$. Hence (87) implies

$$\kappa_0(m(\alpha)) \geq \bar{\epsilon} - (\alpha - \alpha_M),$$

which proves (84). Since the value of a call is decreasing in its strike,

$$m(\alpha) \leq E\left[(\epsilon - \overline{\epsilon} + (\alpha - \alpha_M))^+\right].$$

When $\overline{\epsilon} = \infty$, zero compensation at α_M requires $\kappa_0(0) = \infty$. Since (87) allows the strike to fall by at most one-for-one with α, it cannot fall from infinity to a finite value over any finite increase in α. Hence $\kappa_0(m(\alpha)) = \infty$ and $m(\alpha) = 0$ for every participating type. By limited liability, the manager therefore receives the zero security. □

Thus, when upside is unbounded, first-best investment cannot support any manager compensation. When upside is bounded, manager compensation can be supported only by shock realizations within $\alpha - \alpha_M$ of the upper endpoint of the return distribution. Particularly near the operating threshold, first-best investment can therefore generate only very limited manager rents. Providing the manager with compensation beyond this upper-tail bound requires departing from first-best investment. Transferring substantial surplus to the manager therefore requires reduced investment, exclusion of otherwise efficient types, or both.

C.5 Investment Distortions and the Bilateral Payoff Frontier

The preceding results reveal a sharp contrast between bilateral contracting and delegation. Under delegation, financing can be divided between informed and uninformed investors, allowing the contract to compensate the manager and elicit information while preserving first-best total investment. Under bilateral contracting, by contrast, the informed investor supplies all of the capital and selects investment from the manager's menu. Investment then also governs the growth of the investor's information rents; in particular, at first-best investment, reducing investment lowers the continuation rents of higher-skill investors at no first-order surplus cost.

The security frontier identifies the most effective way to limit these rents for any given expected payment to the manager. Downside risk nevertheless implies that every feasible security has $\Delta(\alpha) < 1$ (Corollary 1), so the information-rent wedge cannot be eliminated. Compensating the manager must therefore ultimately require distortions in investment and participation.

We now characterize the tradeoff between total surplus and manager rents. For $w \in [0, 1]$, consider the weighted objective

$$W_w = E\left[(1 - w)\pi_i(\alpha) + \pi_m(\alpha)\right] = E\left[S(q(\alpha), \alpha) - w\pi_i(\alpha)\right]. \tag{88}$$

The weight on the manager is normalized to one, so increasing w shifts relative weight from

the informed investor toward the manager. When $w = 0$, the objective is total surplus. When $w = 1$, it is the manager's expected payoff. Varying w therefore traces the supported bilateral payoff frontier.

The distortion result follows from a simple local argument. At first-best investment, a marginal reduction in investment has no first-order effect on operating surplus. It does, however, slow the growth of the informed investor's information rents and thereby reduce the rents of all higher types. Once the manager receives positive Pareto weight, first-best investment therefore cannot be optimal.

Proposition 10 (Downward Investment Distortion)**.** *Fix $w > 0$ and suppose returns have downside risk. Then, at any regular optimum—meaning one for which the schedules $q(\alpha)$, $\pi_i(\alpha)$, and $m(\alpha)$ are continuous on the participating region and continuously differentiable on its interior—the participation cutoff satisfies*

$$\alpha_I > \alpha_M, \qquad q(\alpha_I) = \frac{\alpha_M}{2\beta}, \qquad m(\alpha_I) = \alpha_I - \alpha_M. \tag{89}$$

Moreover, for every participating type below the upper endpoint of the skill distribution,

$$q(\alpha) < q^{FB}(\alpha) = \frac{\alpha}{2\beta}. \tag{90}$$

Thus, investment is strictly below first best for almost every participating type, and some types for whom investment would be efficient are excluded.

Proof. At a regular optimum, the security frontier binds for every participating type below the upper endpoint of the skill distribution. To see why, suppose that $\Delta(\alpha) < \overline{\Delta}\big(m(\alpha)\big)$ at some such type. Replace the security locally by a frontier security with the same expected payment, and follow through by lowering the promised utilities of higher types according to the incentive-compatibility envelope, while leaving their investment unchanged. This leaves total operating surplus unchanged but reduces investor rents for a positive mass of types. Because

$$W_w = E\left[S(q(\alpha), \alpha) - w\pi_i(\alpha)\right]$$

and $w > 0$, the adjustment strictly raises the objective. Hence $\Delta(\alpha) = \overline{\Delta}\big(m(\alpha)\big)$.

Suppose that a participating type α below the top of the distribution receives investment at or above the first-best level. Consider a small reduction in investment over a short interval around α, with the remaining contract terms adjusted to preserve incentive compatibility.

The local incentive-compatibility condition is

$$\pi_i' = q\left[1 - \overline{\Delta}(m)\right], \qquad m = \alpha - \beta q - \frac{\pi_i + \delta}{q}.$$

Holding current investor utility fixed, the effect of investment on the growth of investor rents is

$$\frac{\partial \pi_i'}{\partial q} = 1 - \overline{\Delta}(m) + \overline{\Delta}'(m)\left[m - \big(\alpha - 2\beta q\big)\right].$$

If $q \geq q^{FB}(\alpha)$, then

$$\alpha - 2\beta q \leq 0.$$

Because $m \geq 0$, $\overline{\Delta}'(m) \geq 0$, and downside risk implies $\overline{\Delta}(m) < 1$, it follows that

$$\frac{\partial \pi_i'}{\partial q} > 0.$$

Reducing investment therefore lowers the growth of investor rents and hence lowers the rents of all higher types; the reduction propagates because a lower π_i raises m and hence $\overline{\Delta}(m)$ at every higher type (the propagation factor is computed explicitly in the proof of Corollary 2).

At the same time, the first-order effect on operating surplus is

$$dS = \big(\alpha - 2\beta q\big)dq \geq 0,$$

since both factors are nonpositive, while the reduction in investor rents is first order and strictly negative. Because $w > 0$ and there are higher types, this perturbation strictly raises W_w, contradicting optimality. Therefore

$$q(\alpha) < q^{FB}(\alpha)$$

for every participating type below the upper endpoint of the skill distribution. The same argument applies one-sidedly at the participation cutoff.

Participation and manager limited liability require

$$S\big(q(\alpha_I), \alpha_I\big) = \pi_i(\alpha_I) + \pi_m(\alpha_I) \geq 0.$$

If $\alpha_I < \alpha_M$, even first-best investment generates negative operating surplus. If $\alpha_I = \alpha_M$, operating surplus is zero only at $q^{FB}(\alpha_M)$, whereas (90) requires strictly lower investment.

Both cases contradict participation. Hence

$$\alpha_I > \alpha_M.$$

It remains to characterize investment and compensation at the cutoff. The investor's participation constraint binds, so

$$\pi_i(\alpha_I) = 0$$

and therefore

$$m(\alpha_I) = \alpha_I - \beta q(\alpha_I) - \frac{\delta}{q(\alpha_I)}. \tag{91}$$

At the cutoff, there are two marginal ways to reduce the information rents of higher types: exclude the marginal type altogether, or retain that type and reduce his investment slightly. Because the cutoff investor earns zero, exclusion sacrifices $S\big(q(\alpha_I), \alpha_I\big) = q(\alpha_I)m(\alpha_I)$, or $m(\alpha_I)$ per unit of investment eliminated. A marginal reduction in investment instead sacrifices $\alpha_I - 2\beta q(\alpha_I)$ per unit. After the remaining contract terms are adjusted optimally, both perturbations act on the same continuation-rent margin. At an interior optimum, their marginal costs must therefore coincide:[29]

$$m(\alpha_I) = \alpha_I - 2\beta q(\alpha_I). \tag{92}$$

Combining (91) and (92) yields

$$\beta q(\alpha_I) = \frac{\delta}{q(\alpha_I)},$$

and hence

$$q(\alpha_I) = \sqrt{\frac{\delta}{\beta}} = \frac{\alpha_M}{2\beta}.$$

[29]Formally, let $\Lambda(\alpha_I) > 0$ denote the shadow value of reducing continuation information rents (constructed explicitly in (94) below). The one-sided investment condition at the cutoff is

$$\big(\alpha_I - 2\beta q(\alpha_I)\big) f(\alpha_I) = \Lambda(\alpha_I)\left[1 - \overline{\Delta} + \overline{\Delta}'(m)\big(m - \alpha_I + 2\beta q\big)\right],$$

while the free-boundary condition is

$$m(\alpha_I) f(\alpha_I) = \Lambda(\alpha_I)\big(1 - \overline{\Delta}\big).$$

Subtracting gives

$$\big[m(\alpha_I) - \alpha_I + 2\beta q(\alpha_I)\big]\big[f(\alpha_I) + \Lambda(\alpha_I)\overline{\Delta}'(m(\alpha_I))\big] = 0,$$

which yields (92).

Substituting back gives

$$m(\alpha_I) = \alpha_I - \alpha_M,$$

which completes the proof. □

The preceding proposition establishes that bilateral contracting excludes some types for whom investment would be efficient. The next result provides a quantitative lower bound on the extent of this exclusion.

Corollary 2 (Lower Bound on the Participation Cutoff)**.** *At any regular optimum with $w > 0$, the participation cutoff satisfies*

$$\alpha_I \geq \alpha_M + w\left[1 - \overline{\Delta}(\alpha_I - \alpha_M)\right] \frac{1 - F(\alpha_I)}{f(\alpha_I)}. \tag{93}$$

Let $\underline{\alpha}_I$ denote the smallest value satisfying (93)*. Then*

$$\alpha_I \geq \underline{\alpha}_I > \alpha_M.$$

Proof. Let $\Lambda(\alpha)$ denote the gain in the objective from reducing $\pi_i(\alpha)$ by one unit, allowing the resulting reduction in investor rents to propagate to all higher types through the incentive-compatibility condition. To compute this propagation, note that

$$m(\alpha) = \alpha - \beta q(\alpha) - \frac{\pi_i(\alpha) + \delta}{q(\alpha)},$$

so a reduction in π_i raises m. Holding the allocation schedule fixed, let $z(s)$ denote the reduction in $\pi_i(s)$ generated by a one-unit reduction in $\pi_i(\alpha)$, so that $z(\alpha) = 1$. From the payoff identity, this reduction raises the manager's compensation rate at type s by $z(s)/q(s)$. Because the security frontier binds, local incentive compatibility therefore implies

$$z'(s) = \overline{\Delta}'(m(s))z(s).$$

Consequently,

$$z(s) = \exp\left\{\int_\alpha^s \overline{\Delta}'(m(u))\, du\right\}.$$

Because a one-unit reduction in investor rents raises the weighted objective by w, the total gain from reducing $\pi_i(\alpha)$ by one unit is

$$\Lambda(\alpha) = w\int_\alpha^{\overline{\alpha}} f(s) \exp\left\{\int_\alpha^s \overline{\Delta}'(m(u))\, du\right\} ds. \tag{94}$$

In other words, reducing rents at type α reduces rents for every higher type. The associated increase in the manager's compensation raises the sensitivity attainable on the security frontier, which further reduces the growth of investor rents and amplifies the initial saving.

Because $\overline{\Delta}' \geq 0$,

$$\Lambda(\alpha) \geq w\left[1 - F(\alpha)\right]. \tag{95}$$

The free-boundary condition at the participation cutoff is

$$m(\alpha_I) f(\alpha_I) = \Lambda(\alpha_I)\left[1 - \overline{\Delta}(m(\alpha_I))\right]. \tag{96}$$

Using (95) and the cutoff condition $m(\alpha_I) = \alpha_I - \alpha_M$ gives

$$\alpha_I - \alpha_M \geq w\left[1 - \overline{\Delta}(\alpha_I - \alpha_M)\right] \frac{1 - F(\alpha_I)}{f(\alpha_I)},$$

which is (93).

Since α_I satisfies this inequality, it must be at least its smallest solution $\underline{\alpha}_I$. Moreover, the inequality cannot hold at α_M, because $\overline{\Delta}(0) = 0$ and $[1 - F(\alpha_M)]/f(\alpha_M) > 0$. Hence $\underline{\alpha}_I > \alpha_M$. □

Both forms of distortion—the exclusion of low participating types and underinvestment among those who participate—reduce total surplus. The next corollary shows that this aggregate inefficiency weakly increases with the weight placed on reducing investor rents.

Corollary 3 (Aggregate Inefficiency)**.** *Expected total surplus* $E\left[S\big(q_w(\alpha), \alpha\big)\right]$ *is weakly decreasing in* w*. Equivalently, aggregate inefficiency is weakly increasing in the relative weight placed on manager rents.*

Proof. Consider $0 \leq w_1 < w_2 \leq 1$, and let

$$\Sigma_j = E\left[S\big(q_{w_j}(\alpha), \alpha\big)\right], \qquad I_j = E\left[\pi_{i,w_j}(\alpha)\right], \qquad j = 1, 2.$$

Optimality at the two weights implies

$$\Sigma_1 - w_1 I_1 \geq \Sigma_2 - w_1 I_2$$

and

$$\Sigma_2 - w_2 I_2 \geq \Sigma_1 - w_2 I_1.$$

Adding the two inequalities gives $I_1 \geq I_2$. The first inequality then implies

$$\Sigma_1 - \Sigma_2 \geq w_1(I_1 - I_2) \geq 0.$$

Thus expected total surplus is weakly decreasing, and aggregate inefficiency is weakly increasing, in w. □

A second comparative static concerns the return shock. Greater idiosyncratic volatility reduces the exposure to skill that can be delivered to the manager for a fixed expected payment. It therefore lowers the security frontier and contracts the supported bilateral payoff frontier.

Corollary 4 (Volatility Shrinks the Bilateral Frontier)**.** *Let $\epsilon = \sigma_\epsilon z$, where z has mean zero, unit variance, an absolutely continuous distribution with positive density on its support, and a nondecreasing hazard rate. Suppose returns have downside risk. Then*

$$\overline{\Delta}_{\sigma_\epsilon}(m) = \overline{\Delta}_1\left(\frac{m}{\sigma_\epsilon}\right),$$

and the security frontier is strictly decreasing in σ_ϵ for every $m > 0$. Therefore, for every $w \in (0, 1]$, the contract frontier

$$V(w; \sigma_\epsilon) \equiv \max W_w$$

is strictly decreasing in the return volatility σ_ϵ.

Proof. Let

$$h(k) \equiv E\left[(z-k)^+\right].$$

A frontier call with normalized strike k has expected payment and sensitivity

$$m = \sigma_\epsilon h(k), \qquad \Delta = 1 - F_z(k).$$

It follows that

$$\overline{\Delta}_{\sigma_\epsilon}(m) = \overline{\Delta}_1\left(\frac{m}{\sigma_\epsilon}\right).$$

Because $\overline{\Delta}_1$ is strictly increasing on the relevant region,

$$\frac{\partial}{\partial \sigma_\epsilon}\overline{\Delta}_{\sigma_\epsilon}(m) = -\frac{m}{\sigma_\epsilon^2}\overline{\Delta}_1'\left(\frac{m}{\sigma_\epsilon}\right) < 0$$

for every $m > 0$.

Now fix

$$0 < \sigma_\epsilon < \sigma'_\epsilon$$

and let a regular optimal menu at volatility σ'_ϵ be given. Because $w > 0$, the security frontier binds. Thus, the security intended for type $\hat{\alpha}$ can be written as

$$t_{\sigma'_\epsilon}(R; \hat{\alpha}) = \left(R - \overline{R}(\hat{\alpha}) - \sigma'_\epsilon k(\hat{\alpha})\right)^+,$$

where

$$\overline{R}(\hat{\alpha}) = \hat{\alpha} - \beta q(\hat{\alpha}).$$

At the lower volatility, replace this security by

$$t_{\sigma_\epsilon}(R; \hat{\alpha}) = (\sigma'_\epsilon - \sigma_\epsilon)\, h\big(k(\hat{\alpha})\big) + \left(R - \overline{R}(\hat{\alpha}) - \sigma_\epsilon k(\hat{\alpha})\right)^+. \tag{97}$$

For the intended type, the expected payment is unchanged:

$$(\sigma'_\epsilon - \sigma_\epsilon)\, h(k) + \sigma_\epsilon h(k) = \sigma'_\epsilon h(k) = m.$$

Investment and both parties' truthful expected payoffs are therefore unchanged.

To verify global incentive compatibility, let

$$d = \alpha - \hat{\alpha}.$$

Under volatility σ, the payoff advantage to type α from taking the contract intended for $\hat{\alpha}$, relative to the intended type's payoff from that contract, is

$$q(\hat{\alpha}) G_\sigma(d, k(\hat{\alpha})),$$

where

$$G_\sigma(d, k) = \int_0^d F_z\left(k - \frac{u}{\sigma}\right) du.$$

If $\sigma_\epsilon < \sigma'_\epsilon$, then

$$G_{\sigma_\epsilon}(d, k) \leq G_{\sigma'_\epsilon}(d, k)$$

for every d. For $d > 0$, the integrand is lower at the lower volatility. For $d < 0$, the integrand is higher, but the direction of integration is reversed. Thus, every deviation is weakly less attractive at the lower volatility, while all truthful payoffs are unchanged. The transformed

menu is therefore incentive compatible and delivers the same weighted objective. Hence

$$V(w;\sigma_\epsilon) \geq V(w;\sigma'_\epsilon).$$

The inequality is strict. At a regular optimum with $w > 0$, Proposition 10 implies

$$m(\alpha_I) = \alpha_I - \alpha_M > 0.$$

By continuity, $m(\alpha) > 0$ on an interval of participating types. For these types, the transformed security has sensitivity

$$\overline{\Delta}_{\sigma'_\epsilon}(m) < \overline{\Delta}_{\sigma_\epsilon}(m),$$

so the lower-volatility security frontier is strictly slack on a set of positive measure. By the same frontier-replacement argument that establishes frontier binding at a regular optimum, this slack permits a feasible increase in the security's sensitivity and a corresponding reduction in continuation rents, strictly raising W_w. Consequently,

$$V(w;\sigma_\epsilon) > V(w;\sigma'_\epsilon)$$

for every $w \in (0, 1]$. □

The limiting case provides a useful benchmark. For every fixed finite expected payment m,

$$\lim_{\sigma_\epsilon \to \infty} \overline{\Delta}_{\sigma_\epsilon}(m) = \overline{\Delta}_1(0) = 0.$$

Thus, as idiosyncratic return risk becomes arbitrarily large, no finite expected payment can provide the manager with exposure to skill. Because admissible securities satisfy $t_R \geq 0$, $\Delta = E[t_R] = 0$ implies that the security is constant almost surely. Hence the bilateral contracting problem converges to the fixed-payment benchmark, $\Delta = 0$.

The opposite limit corresponds to nearly perfect screening. For every fixed $m > 0$,

$$\overline{\Delta}_{\sigma_\epsilon}(m) = \overline{\Delta}_1\left(\frac{m}{\sigma_\epsilon}\right) \longrightarrow 1$$

as $\sigma_\epsilon \downarrow 0$. Thus, as return noise vanishes, the information-rent wedge disappears. Under regularity, $\alpha_I \to \alpha_M$ and $q(\alpha) \to q^{FB}(\alpha)$, so expected surplus converges to first best. For $w > 0$, investor rents also vanish, and the manager captures the first-best surplus in the limit.

C.6 Computing the Bilateral Payoff Frontier

Although endogenous security design prevents a closed-form characterization of the optimal investment schedule, the supported bilateral payoff frontier can be computed efficiently from the first-order conditions. We first solve the locally incentive-compatible problem and then verify that the resulting menu satisfies sufficient conditions for global incentive compatibility.

The endpoint $w = 0$ is available analytically: investment is first best, $\alpha_I = \alpha_M$, and the manager receives zero rents. We therefore describe the computation for $w > 0$. At a regular optimum, the security frontier binds, so

$$\Delta(\alpha) = \overline{\Delta}\big(m(\alpha)\big).$$

Local incentive compatibility then requires

$$\pi_i'(\alpha) = q(\alpha)\left[1 - \overline{\Delta}\big(m(\alpha)\big)\right], \qquad \pi_i(\alpha_I) = 0, \tag{98}$$

where the manager's expected compensation rate is recovered from the investor's payoff identity:

$$m(\alpha) = \alpha - \beta q(\alpha) - \frac{\pi_i(\alpha) + \delta}{q(\alpha)}. \tag{99}$$

Recall from (94) that $\Lambda(\alpha)$ is the gain in the weighted objective from reducing $\pi_i(\alpha)$ by one unit, including the resulting reduction in the rents of all higher types. Given $m(\cdot)$, it can be computed from (94), or equivalently from

$$\Lambda'(\alpha) = -wf(\alpha) - \overline{\Delta}'(m(\alpha))\Lambda(\alpha), \qquad \Lambda(\overline{\alpha}) = 0, \tag{100}$$

with the usual limiting interpretation when the support of skill is unbounded.

At an interior participating type, the investment first-order condition is

$$S_q(q, \alpha) f(\alpha) = \Lambda(\alpha) \frac{\partial}{\partial q} \left\{ q \left[1 - \overline{\Delta}(m)\right] \right\},$$

where the derivative with respect to q holds the state π_i fixed. The left-hand side is the marginal contribution of investment to operating surplus. The right-hand side is the shadow cost of the resulting marginal increase in the growth of investor rents.

For notational convenience, denote the marginal operating surplus by

$$\Gamma(\alpha) \equiv S_q(q(\alpha), \alpha).$$

Using $qm = S - \pi_i$, we have, holding π_i fixed,

$$q\frac{\partial m}{\partial q} = \Gamma - m.$$

The first-order condition therefore becomes

$$\Gamma(\alpha)f(\alpha) = \Lambda(\alpha)\left\{1 - \overline{\Delta}(m(\alpha)) + \overline{\Delta}'(m(\alpha))\left[m(\alpha) - \Gamma(\alpha)\right]\right\}. \tag{101}$$

For computational purposes, define

$$r(\alpha) \equiv \frac{\Lambda(\alpha)}{f(\alpha)}.$$

Solving (101) for marginal operating surplus gives

$$\Gamma(\alpha) = \frac{r(\alpha)\left\{1 - \overline{\Delta}(m(\alpha)) + \overline{\Delta}'(m(\alpha))m(\alpha)\right\}}{1 + r(\alpha)\overline{\Delta}'(m(\alpha))}. \tag{102}$$

For the quadratic specification,

$$\Gamma(\alpha) = \alpha - 2\beta q(\alpha) = 2\beta\left[q^{FB}(\alpha) - q(\alpha)\right].$$

Thus, $\Gamma(\alpha)$ measures the downward investment distortion in marginal-surplus units, and investment is recovered pointwise from

$$q(\alpha) = \frac{\alpha - \Gamma(\alpha)}{2\beta}. \tag{103}$$

At the participation cutoff,

$$q(\alpha_I) = \frac{\alpha_M}{2\beta}, \tag{104}$$

or equivalently,

$$\Gamma(\alpha_I) = \alpha_I - \alpha_M. \tag{105}$$

Together with $\pi_i(\alpha_I) = 0$, this also implies

$$m(\alpha_I) = \alpha_I - \alpha_M.$$

The fixed-point formulation also makes the economics transparent. Compensation determines the attainable security sensitivity; security sensitivity determines the propagation and shadow cost of investor rents; that shadow cost determines the marginal-surplus wedge

Γ; and Γ determines the extent of underinvestment.

Together, the shadow-value equation (100), the pointwise investment rule (102)–(103), the cutoff condition (105), and the local incentive-compatibility and payoff identities (98)–(99) yield a vectorized fixed-point algorithm. Given a current compensation schedule $m^{(n)}(\cdot)$, we first compute $\Lambda^{(n)}(\cdot)$ from (100) and then compute $\Gamma^{(n)}(\cdot)$ from (102). The updated cutoff $\widetilde{\alpha}_I^{(n+1)}$ is the unique solution to

$$\Gamma^{(n)}(\alpha) = \alpha - \alpha_M,$$

and investment is updated pointwise according to

$$\widetilde{q}^{(n+1)}(\alpha) = \frac{\alpha - \Gamma^{(n)}(\alpha)}{2\beta}, \qquad \alpha \geq \widetilde{\alpha}_I^{(n+1)}.$$

Holding the preceding compensation schedule fixed in the local incentive-compatibility equation, the updated investor payoff is

$$\widetilde{\pi}_i^{(n+1)}(\alpha) = \int_{\widetilde{\alpha}_I^{(n+1)}}^{\alpha} \widetilde{q}^{(n+1)}(s) \left[1 - \overline{\Delta}\big(m^{(n)}(s)\big)\right] ds.$$

The implied compensation schedule is then

$$\widetilde{m}^{(n+1)}(\alpha) = \alpha - \beta \widetilde{q}^{(n+1)}(\alpha) - \frac{\widetilde{\pi}_i^{(n+1)}(\alpha) + \delta}{\widetilde{q}^{(n+1)}(\alpha)}.$$

Thus, one iteration of the fixed-point map is

$$m^{(n)} \longrightarrow \Lambda^{(n)} \longrightarrow \Gamma^{(n)} \longrightarrow \left(\widetilde{\alpha}_I^{(n+1)}, \widetilde{q}^{(n+1)}\right) \longrightarrow \widetilde{\pi}_i^{(n+1)} \longrightarrow \widetilde{m}^{(n+1)}.$$

The equation for Λ is evaluated using a reverse cumulative integral, while the equation for π_i is evaluated using a forward cumulative integral. Each iteration is therefore linear in the number of grid points.

After convergence, we evaluate

$$W_w = \int_{\alpha_I}^{\overline{\alpha}} \left[S\big(q(\alpha), \alpha\big) - w\pi_i(\alpha)\right] f(\alpha)\, d\alpha. \tag{106}$$

We verify that

$$q(\alpha) > 0, \qquad m(\alpha) \geq 0, \qquad \pi_i(\alpha) \geq 0$$

throughout the participating region. We use homotopy in w, taking the solution at each value of w as the initial condition for the next. If multiple regular fixed points are obtained,

we retain the feasible, globally incentive-compatible solution with the highest value of W_w. Repeating the calculation across values of w traces the supported bilateral payoff frontier.

Global incentive compatibility. The preceding calculation imposes local incentive compatibility. To determine when it also guarantees global incentive compatibility, let $\kappa_0(\widehat{\alpha})$ denote the centered strike of the frontier call offered to reported type $\widehat{\alpha}$:

$$m(\widehat{\alpha}) = E\left[(\epsilon - \kappa_0(\widehat{\alpha}))^+\right].$$

If the investor learns α but she selects the contract intended for $\widehat{\alpha}$, her expected payoff is

$$\pi_i(\alpha;\widehat{\alpha}) = q(\widehat{\alpha})\left[\alpha - \beta q(\widehat{\alpha}) - E\left[(\alpha - \widehat{\alpha} + \epsilon - \kappa_0(\widehat{\alpha}))^+\right]\right] - \delta. \tag{107}$$

The sensitivity of this payoff to true skill is

$$\frac{\partial \pi_i(\alpha;\widehat{\alpha})}{\partial \alpha} = q(\widehat{\alpha})\Pr\left(\epsilon \le \widehat{\alpha} + \kappa_0(\widehat{\alpha}) - \alpha\right).$$

Contracts intended for higher reported types therefore become relatively more attractive to higher-skill investors provided

$$q'(\widehat{\alpha}) \ge 0, \qquad 1 + \kappa_0'(\widehat{\alpha}) \ge 0. \tag{108}$$

Indeed,

$$\begin{aligned}\frac{\partial^2 \pi_i(\alpha;\widehat{\alpha})}{\partial \alpha\, \partial \widehat{\alpha}} &= q'(\widehat{\alpha}) F_\epsilon\left(\widehat{\alpha} + \kappa_0(\widehat{\alpha}) - \alpha\right) \\ &\quad + q(\widehat{\alpha}) f_\epsilon\left(\widehat{\alpha} + \kappa_0(\widehat{\alpha}) - \alpha\right)\left[1 + \kappa_0'(\widehat{\alpha})\right],\end{aligned} \tag{109}$$

which is nonnegative under (108). Economically, the first condition says that contracts intended for higher types entail weakly greater investment. The second says that, holding the investor's true skill fixed, the manager's option becomes weakly less likely to be exercised as the intended type rises.

The payoff identity and the envelope condition ensure that truthful reporting satisfies the local first-order condition. Under single crossing, a type α then prefers higher reports when $\widehat{\alpha} < \alpha$ and lower reports when $\widehat{\alpha} > \alpha$. Truthful reporting is therefore globally optimal. Because the cutoff type earns zero and the payoff from every fixed menu item is increasing in true skill, all types below α_I also prefer not to participate.

The strike condition in (108) can be checked directly from the equations used in the

computation. Since

$$\frac{dm}{d\kappa_0} = -\overline{\Delta}(m),$$

differentiating the investor's payoff identity and using (98) gives

$$m'(\alpha) = \overline{\Delta}\big(m(\alpha)\big) + \frac{q'(\alpha)}{q(\alpha)}\left[\alpha - 2\beta q(\alpha) - m(\alpha)\right].$$

Hence, wherever $m(\alpha) > 0$,

$$1 + \kappa_0'(\alpha) = \frac{q'(\alpha)}{q(\alpha)\overline{\Delta}(m(\alpha))}\left[m(\alpha) - \big(\alpha - 2\beta q(\alpha)\big)\right]. \tag{110}$$

It is therefore sufficient to verify at the computed solution that

$$q'(\alpha) \geq 0, \qquad m(\alpha) \geq \alpha - 2\beta q(\alpha) = 2\beta\left[q^{FB}(\alpha) - q(\alpha)\right]. \tag{111}$$

The second inequality has a useful economic interpretation. It places a lower bound on the manager's expected compensation rate in terms of the investment distortion. When investment rises with the intended type, a marginal increase in investment creates operating surplus $\alpha - 2\beta q$. If this gain exceeds the amount m captured by the manager, then the investor captures too much of the incremental surplus for single crossing to be assured.

Together with the local incentive-compatibility and participation conditions, (111) guarantees that the resulting menu is globally incentive compatible. If the inequalities fail, the local solution may still be implementable, but global incentive compatibility must then be verified directly.

C.7 Numerical Illustration of the Bilateral Frontier

For the figures and table below, we assume

$$\alpha \sim N(\mu_\alpha, \sigma_\alpha^2), \qquad \mu_\alpha = 4\%, \qquad \sigma_\alpha \in \{2\%, 6\%\}.$$

We set $\beta = 0.04\%$ and $\delta = 1$, so that the first-best operating threshold is

$$\alpha_M = 2\sqrt{\beta\delta} = 4\%.$$

For idiosyncratic return risk, we consider $\epsilon \sim N(0, \sigma_\epsilon^2)$ with $\sigma_\epsilon \in \{5\%, 20\%\}$, as well as the fixed-payment benchmark $\Delta = 0$, in which the manager's compensation is independent of the realized return — this benchmark is the limiting case as σ_ϵ becomes large.

Unless otherwise noted, the results below correspond to $w = 1$, so that the bilateral contract maximizes the manager's expected payoff. The final figure instead varies w to trace the bilateral Pareto frontier.

Table 3: Bilateral Contracting: Expected Payoffs and Efficiency Losses ($w = 1$, $\sigma_\alpha = 2\%$). The first-best row is the delegated management benchmark. In each case, the cutoff is α_I, $E[\pi_m]$ is the maximum attainable manager payoff, $E[\pi_i]$ is the corresponding investor payoff. Expected total surplus $E[S]$ is shown as a percentage of first best. Exclusion loss is the percent of first-best surplus forgone from types between α_M and α_I. Scale loss is the percent of surplus forgone because investment is below the first-best level.

Scenario	Allocation			Payoffs		Efficiency (% of first best)		
	α_I (%)	$q(\alpha_I)$	$\Pr(\alpha > \alpha_I)$	$E[\pi_m]$	$E[\pi_i]$	$E[S]$	Excl. loss	Scale loss
First best (delegation benchmark)	4.00	50.0	50.00	0.5239	0.0000	100.00	0.00	0.00
Bilateral ($\sigma_\epsilon = 5\%$)	5.28	50.0	26.11	0.3093	0.1170	81.36	15.57	3.07
Bilateral ($\sigma_\epsilon = 20\%$)	5.44	50.0	23.52	0.2481	0.1547	76.87	19.51	3.62
Bilateral ($\Delta = 0$)	5.50	50.0	22.61	0.2215	0.1727	75.23	21.03	3.74

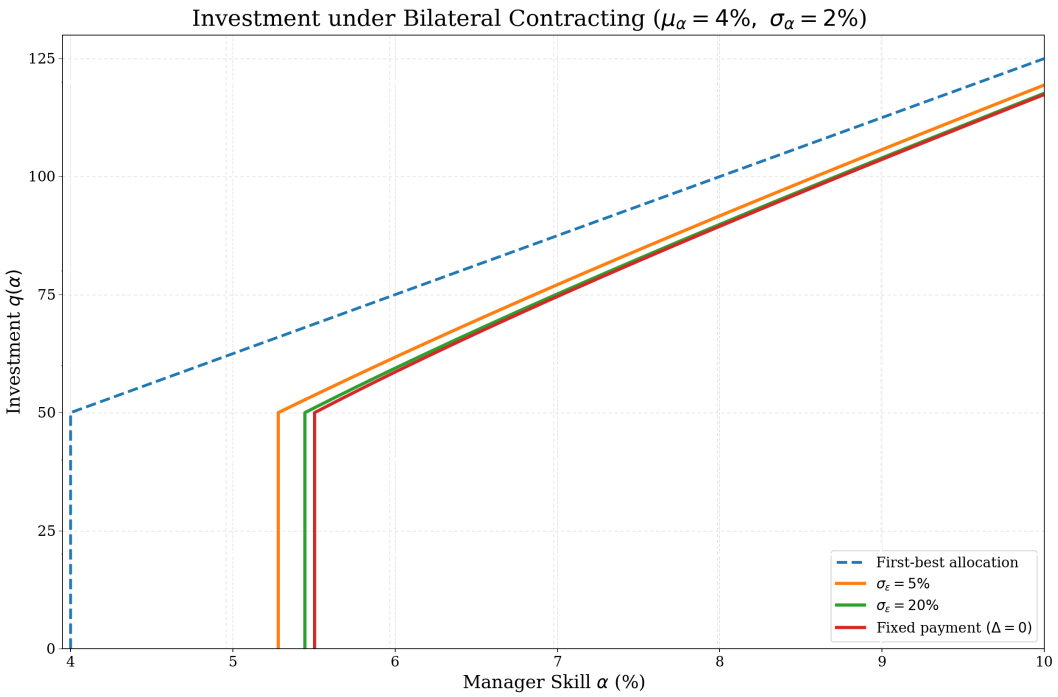


Figure 5: Investment under Bilateral Contracting. The dashed line represents the first-best investment rule $q^{\text{FB}}(\alpha) = \alpha/(2\beta)$. The colored lines show the manager-optimal bilateral allocations under different levels of idiosyncratic return risk and under the fixed-payment benchmark $\Delta = 0$. The colored vertical segments mark the corresponding participation cutoffs α_I, while the dashed vertical line marks the first-best threshold α_M. Greater idiosyncratic risk reduces the Δ attainable from a given expected payment, exacerbating the optimal distortion for the manager. The optimal bilateral contract therefore excludes more types and reduces investment further below the first best. The figure uses $w = 1$ and $\sigma_\alpha = 2\%$.

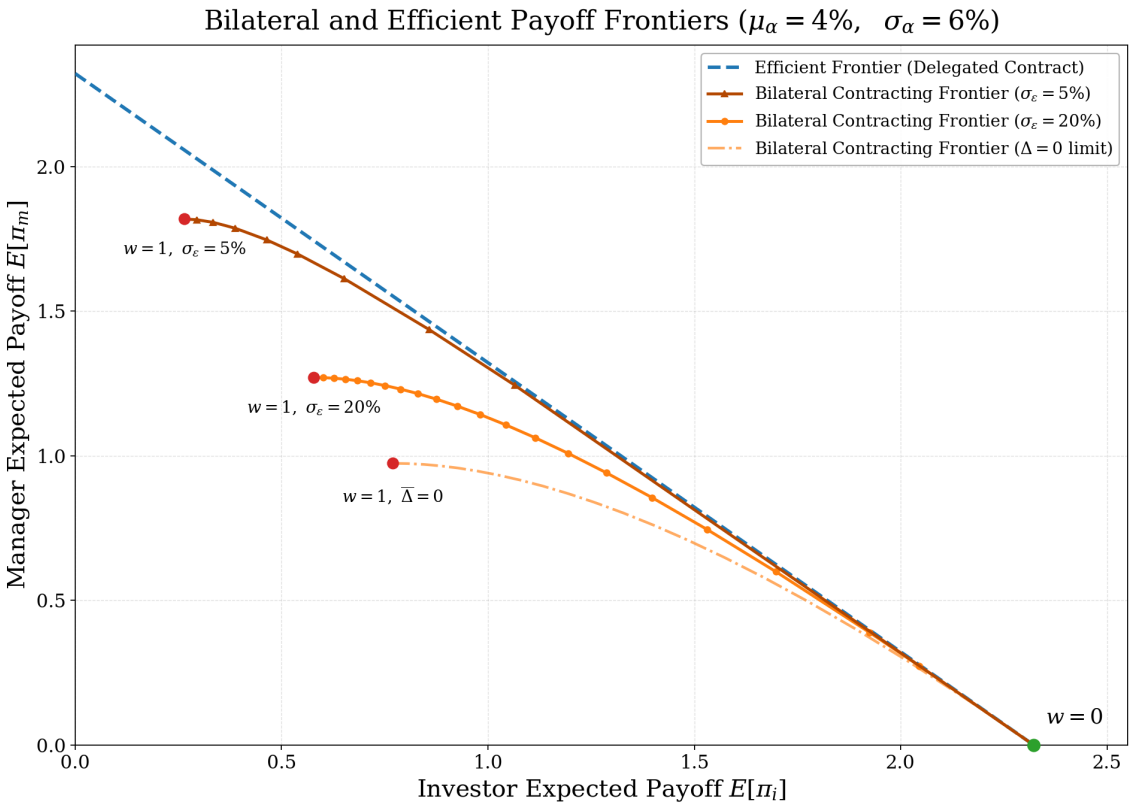


Figure 6: Bilateral and Efficient Payoff Frontiers. The dashed blue line is the efficient frontier, along which total surplus equals its first-best value $E[S^{FB}(\alpha)] = 2.322$; the delegated contract attains every point on it — with $k \geq 1$ for investor payoffs up to $c^* = 1.125$, and with $k < 1$ and the fee schedule of Section 3.3 beyond (Proposition 4). The remaining curves are bilateral payoff frontiers, each traced by the weight w on manager rents, for return volatilities $\sigma_\epsilon = 5\%$ and 20%, and in the limit as return volatility grows without bound, so that the fund's return becomes uninformative for compensation and the manager can be paid only with fixed payments ($\overline{\Delta} = 0$). All three meet the efficient frontier at the same zero-rent point $w = 0$, where investment is first best and the manager earns nothing (Proposition 9); everywhere else they lie strictly inside, reflecting exclusion of types above α_M and underinvestment in the funds that operate (Proposition 10). By Corollary 4, the bilateral frontier contracts monotonically as σ_ϵ rises, pivoting on the efficient endpoint: the frontier for any volatility between 5% and 20% lies between the first two curves, and for any higher volatility between the second and the $\overline{\Delta} = 0$ limit. Noisier returns thus shrink the manager's attainable rents — his maximum expected payoff (the labeled $w = 1$ endpoints) falls from 1.819 at $\sigma_\epsilon = 5\%$ to 1.271 at 20% and toward 0.974 in the limit — while raising the minimum rent the informed investor retains, from 0.264 to 0.577 and toward 0.769. The figure uses $\mu_\alpha = 4\%$ and $\sigma_\alpha = 6\%$.

References

Agarwal, Vikas, Vyacheslav Fos, and Wei Jiang, "Inferring Reporting-Related Biases in Hedge Fund Databases from Hedge Fund Equity Holdings," *Management Science*, 2013, *59* (6), 1271–1289.

Axelson, Ulf, Per Strömberg, and Michael S. Weisbach, "Why Are Buyouts Levered? The Financial Structure of Private Equity Funds," *The Journal of Finance*, 2009, *64* (4), 1549–1582.

Barth, Daniel, Juha Joenväärä, Mikko Kauppila, and Russ Wermers, "A bias-free assessment of the hedge fund industry: A new evaluation of total assets, alphas, and the flow-performance relation," *OFR WP*, 2023, pp. 20–01.

Begenau, Juliane and Emil N. Siriwardane, "Fee Variation in Private Equity," *The Journal of Finance*, 2024, *79* (2), 1199–1247.

Berk, Jonathan B. and Richard C. Green, "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 2004, *112* (6), 1269–1295.

_ , Jules H. van Binsbergen, and Binying Liu, "Matching Capital and Labor," *Journal of Finance*, 2017, (Forthcoming).

Binfarè, Matteo, Gregory Brown, Robert Harris, and Christian Lundblad, "How does human capital affect investing? Evidence from university endowments," *Review of Finance*, 2023, *27* (1), 143–188.

Braun, Reiner, Nils Dorau, Tim Jenkinson, and Daniel Urban, "Size, returns, and value: Do private equity firms allocate capital according to manager skill?," *The Journal of Finance*, 2026, *81* (3), 1661–1700.

Brown, Gregory W., Oleg R. Gredil, and Preetesh Kantak, "Finding Fortune: How Do Institutional Investors Pick Asset Managers?," *Review of Financial Studies*, 2023, *36* (8), 3071–3121.

Brown, Keith C, Lorenzo Garlappi, and Cristian Tiu, "Asset allocation and portfolio performance: Evidence from university endowment funds," *Journal of Financial Markets*, 2010, *13* (2), 268–294.

Cavagnaro, Daniel R, Berk A. Sensoy, Yingdi Wang, and Michael S Weisbach, "Measuring Institutional Investors' Skill at Making Private Equity Investments," *The Journal of Finance*, 2019, *74* (6), 3089–3134.

Edelman, Daniel, William Fung, David A Hsieh, and Narayan Y Naik, "Funds of hedge funds: performance, risk and capital formation 2005 to 2010," *Financial Markets and Portfolio Management*, 2012, *26* (1), 87–108.

Fu, Xiaoyong Jack and Lucian A Taylor, "Due Diligence and the Allocation of Venture Capital," Technical Report, National Bureau of Economic Research 2025.

Fung, William, David A Hsieh, Narayan Y Naik, and Tarun Ramadorai, "Hedge funds: Performance, risk, and capital formation," *The Journal of Finance*, 2008, *63* (4), 1777–1803.

Glode, Vincent and Richard C. Green, "Information spillovers and performance persistence for hedge funds," *Journal of Financial Economics*, 2011, *101* (1), 1–17.

Gompers, Paul and Josh Lerner, "An analysis of compensation in the U.S. venture capital partnership," *Journal of Financial Economics*, 1999, *51* (1), 3–44.

Goyal, Amit, Sunil Wahal, and M. Deniz Yavuz, "Picking Partners: Manager Selection in Private Markets," *Journal of Financial Economics*, 2026, *178*, 104250.

Grossman, Sanford J. and Joseph E. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *The American Economic Review*, jun 1980, *70* (3), 393–408.

Harris, Robert S, Tim Jenkinson, Steven N Kaplan, and Ruediger Stucke, "Has persistence persisted in private equity? Evidence from buyout and venture capital funds," *Journal of Corporate Finance*, 2023, *81*, 102361.

Hochberg, Yael V, Alexander Ljungqvist, and Annette Vissing-Jørgensen, "Informational holdup and performance persistence in venture capital," *The Review of Financial Studies*, 2014, *27* (1), 102–152.

Howell, Sabrina T., Dean Parker, and Ting Xu, "Tyranny of the Personal Network: The Limits of Arm's Length Fundraising in Venture Capital," NBER Working Paper 33080, National Bureau of Economic Research 2024.

Jagannathan, Ravi, Alexey Malakhov, and Dmitry Novikov, "Do hot hands exist among hedge fund managers? An empirical evaluation," *The Journal of Finance*, 2010, *65* (1), 217–255.

Kaplan, Steven N.. and Antoinette Schoar, "Private Equity Performance: Returns, Persistence, and Capital Flows," *The Journal of Finance*, 2005, *60* (4), 1791–1823.

Kaplan, Steven N and Berk A Sensoy, "Private equity performance: A survey," *Annual Review of Financial Economics*, 2015, *7* (1), 597–614.

Korteweg, Arthur, "Risk and return in private equity," in "Handbook of the economics of corporate finance," Vol. 1, Elsevier, 2023, pp. 343–418.

Kosowski, Robert, Narayan Y. Naik, and Melvyn Teo, "Do hedge funds deliver alpha? A Bayesian and bootstrap analysis," *Journal of Financial Economics*, 2007, *84* (1), 229–264.

Leland, Hayne E and David H Pyle, "Informational asymmetries, financial structure, and financial intermediation," *The journal of Finance*, 1977, *32* (2), 371–387.

Lerner, Josh and Antoinette Schoar, "The illiquidity puzzle: theory and evidence from private equity," *Journal of Financial Economics*, 2004, *72* (1), 3–40.

_ , _ , and Jialan Wang, "Secrets of the Academy: The Drivers of University Endowment Success," *Journal of Economic Perspectives*, September 2008, *22* (3), 207–22.

_ , _ , and Wan Wongsunwai, "Smart institutions, foolish choices: The limited partner performance puzzle," *The Journal of Finance*, 2007, *62* (2), 731–764.

Mailath, G. J., "Incentive Compatibility in Signaling Games with a Continuum of Types," *Econometrica*, November 1987, *55* (6), 1349–1365.

Marquez, Robert, Vikram Nanda, and M. Deniz Yavuz, "Private Equity Fund Returns and Performance Persistence," *Review of Finance*, 8 2015, *19* (5), 1783–1823.

Martinez, Jose Vicente and Yiming Qian, "Selectivity, Favoritism, and Performance: The Role of Investment Consultants in Private Equity," *Review of Corporate Finance Studies*, 2026, p. cfag014.

Maurin, Vincent, David T. Robinson, and Per Strömberg, "A Theory of Liquidity in Private Equity," *Manage. Sci.*, October 2023, *69* (10), 5740–5771.

Metrick, Andrew and Ayako Yasuda, "The Economics of Private Equity Funds," *Review of Financial Studies*, 2010, *23* (6), 2303–2341.

Mittal, Vrinda, "Desperate Capital Breeds Productivity Loss: Evidence from Public Pension Investments in Private Equity," *Available at SSRN 4283853*, 2026.

MJ Hudson LP Unit, "Private Equity Fund Terms Research: Part I of III: Economics," Technical Report, MJ Hudson 2018.

Pan, Colin and Joe Mason, "Fair Fees and Alignment—Building Long-Term Partnerships with Best-In-Class Asset Managers," Technical Report, Partners Capital 2019.

Pegoraro, Stefano, "Incentives and Performance with Optimal Money Management Contracts," *Journal of Political Economy*, 2023, *131* (10), 2920–2968.

_ , Sophie Shive, and Rafael Zambrana, "Democratizing Illiquid Assets: Liquidity Transformation and Performance in Interval Funds," *Available at SSRN 5365061*, 2025.

Rainey, Brenda, Hugh MacArthur, Lindsay Creedon, and Aiyu Nicholson, "Private Equity's Reality Check: The GP Outlook for 2026," *Bain & Company Insights*, 3 2026.

Sorensen, Morten, Neng Wang, and Jinqiang Yang, "Valuing Private Equity," *The Review of Financial Studies*, 2014, *27* (7), 1977–2021.