Incentive Fees, Alignment of Interests, and Rational Decision-Making in the Context of Special Opportunity Investments

Richard Hermann[1]

Original draft: December 29, 2010 · This revision: July 11, 2026

Abstract

That an incentive fee behaves like a call option, and is therefore worth more when the underlying is more volatile, is well established. This paper develops a consequence of that fact that has received less attention: a *cross-sectional* one. Because a concentrated, single-name position carries far more volatility than the diversified book it is drawn from, a flat carry—the same twenty percent charged on the fund—is worth close to twice as much on the isolated bet. We call this wedge the *concentration premium* in incentive fees, and we argue it is largest precisely where it is least examined: in "special opportunity" investments held outside a diversified fund—separately managed accounts, side pockets, and above all single-firm, deal-by-deal vehicles such as independent-sponsor SPVs, where profits are never netted across a portfolio. This channel is distinct from the time-series erosion documented by Ben-David, Birru, and Rossi (2020), whose effective-versus-nominal wedge arises from contract asymmetry and the absence of netting across a fund's life; here the wedge arises cross-sectionally, from the volatility of the single position itself. At the illustrated volatilities the premium is a hurdle: the isolated position must out-earn the fund by roughly eighty basis points a year merely to pay for its own fee structure, before any compensation for undiversified risk. We state the fairness of a flat carry as two conditions—sufficient incremental expected return, and the investor's ability to verify it ex ante—and argue the market fails to reprice because errors on the two sides of the table reinforce one another: probability weighting of the big payout and inside-view forecasting on the demand side, and self-attributed overconfidence, fed by track records too noisy to reveal skill, on the supply side. We close with fee-design implications robust to that failure: carry scaled to position volatility, cumulative cross-investment measurement, clawbacks, and high-water marks. The measurement remedies echo Litvak (2009) and Metrick and Yasuda (2010); the volatility-scaling prescription and the single-firm application are, to our knowledge, new.

Keywords: incentive fees, carried interest, deal-by-deal carry, co-investment, concentration, volatility, behavioral finance.

JEL Classification: G11, G12, G23, G24, H30.

[1]An earlier version of this paper was drafted and privately circulated to a hedge fund manager in December 2010. This revision preserves the original argument, sharpens its central mechanism, and situates the paper within the literature that has developed in the intervening years; the author claims no priority over that literature.

1. Introduction

"Special opportunity investments" are one-off investments offered to investors on the strength of an unusually high anticipated return. Alongside that attractive expected return, they usually carry risk attributes or constraints—concentration and illiquidity chief among them—that prevent them from being aggregated with the holdings of a diversified fund. They are therefore held, and charged for, on their own: through separately managed accounts, side pockets, standalone special-purpose funds, and, in private markets, single-firm, deal-by-deal vehicles such as independent-sponsor SPVs and co-investments.

This paper isolates one feature of the fees charged on these investments. That an incentive fee is economically a call option written by the investor to the manager, and that its value rises with volatility, is well established; managers thus have a structural incentive to court volatility, and the fee is worth more the more volatile the book. Our point is the *cross-sectional* corollary. A special opportunity is, almost by definition, far more volatile than the diversified fund from which its terms are borrowed—it is a single name, undiversified, its idiosyncratic risk not netted against anything. Applied to that isolated position, a flat twenty-percent carry is worth close to twice what the same headline rate is worth on the fund. We call this wedge the *concentration premium* in incentive fees. It is a feature of the fee's option value, not of any misconduct, and it is invisible so long as one looks only at the headline rate.

The paper sits among three strands of prior work, and it is worth saying at the outset what is not new. First, the option characterization of performance fees and the resulting volatility incentive are standard (Goetzmann, Ingersoll, and Ross, 2003; Carpenter, 2000; Anson, 2001). Second, and closest to this paper, Ben-David, Birru, and Rossi (2020) show empirically that the effective incentive fee investors pay is roughly 2.5 times the nominal rate—on the order of fifty percent rather than twenty—because gains on which fees are collected are later offset by losses that are never clawed back. That wedge is real and large, but it is a *time-series* phenomenon: it arises from contract asymmetry and the absence of netting across a fund's life. The concentration premium is orthogonal to it—a *cross-sectional* wedge that arises from the volatility of the single position, and that would exist even in a world with perfect intertemporal netting. Third, the measurement remedies we reach—cumulative, cross-investment accounting and clawbacks—are anticipated by the private-equity contracting literature, which has long studied the value of distribution and timing rules (Litvak, 2009) and the present value of carried interest under alternative waterfall and clawback schemes (Metrick and Yasuda, 2010). What we add to that literature is the prescription that the carry *rate* itself, not merely its timing, should fall as the volatility of the underlying position rises, and the observation that the single-firm, deal-by-deal segment is where the un-netted concentration premium is largest.

Private markets in fact display the two poles of this argument side by side. Co-investments offered alongside a fund are commonly struck at reduced or zero carry—the market, in effect, already discounts the concentrated single-asset bet, if for relationship reasons rather than the one

developed here. Independent-sponsor and other fundless, deal-by-deal structures do the opposite: they charge a full promote, frequently in the ten-to-thirty-percent range, on a single private company, with no fund and therefore no cross-deal netting. It is in the second setting that the concentration premium is neither diversified away nor negotiated down, and it is that setting the paper is ultimately about.

Our objectives are two. First, to make the concentration premium explicit, so that managers and investors price single-name and deal-by-deal carry with the volatility of the position in view rather than by anchoring to the fund's headline rate. Second, to improve long-term value creation where the interests of managers and investors overlap, by proposing fee terms—volatility-scaled rates, cumulative measurement, clawbacks, and high-water marks—that align the fee with the economics of the bet.

1.1 A Motivating Example

Consider a publicly traded company that agrees to sell a private placement to a hedge fund. One term of the placement restricts the fund from selling the security for one year. If the fund's existing structure offers investors quarterly liquidity, there is a mismatch between the security's one-year lockup and the fund's quarterly terms that may preclude the fund from participating.

In situations like this, funds often offer the opportunity to investors "outside" the fund. The three most common vehicles are (i) separately managed accounts, (ii) side pockets, and (iii) standalone special-purpose funds. Under two of these three structures it is customary to charge separate fees, typically a management fee of 0% to 1.5% plus a 20% performance fee[2] on the appreciation of the investment. The management fee is comparable to, or slightly lower than, the fee on a typical investment in the flagship fund.

2. Incentive Fees as Call Options

A body of prior work models a manager's incentive fee as a European call option on the value of the fund's assets.[3] The strike price is the net asset value at the time the investor invests. Like most models this one is imperfect—it assumes, for instance, that asset prices are normally distributed, whereas financial returns are observably leptokurtic (they exhibit "fat tails"). As an approximation, however, it isolates several important drivers of the expected value of an incentive fee.

2.1 Volatility Raises the Expected Value of the Fee

[2] A management fee in the range of 0% to 1.5% together with a 20% performance (carried-interest) fee is the customary structure for such arrangements; the management fee is typically comparable to, or slightly below, the management fee charged on a diversified investment in the manager's flagship fund.

[3] The characterization of a performance fee as a call option written by the investor to the manager, and the valuation of high-water-mark provisions in that framework, is developed in Goetzmann, Ingersoll, and Ross (2003); see also Anson (2001), Agarwal and Naik (2004), and, on the risk-taking incentives created by option-like compensation, Carpenter (2000).

Treating the incentive fee as a European call, its value can be written with the Black–Scholes formula:

$$C(S_t,\ t)\ =\ S_t\,N(d_1)\ -\ K\,e^{-r(T-t)}\,N(d_2)$$

$$d_1\ =\ \frac{\ln\!\left(\frac{S_t}{K}\right)+\left(r+\frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}$$

$$d_2\ =\ d_1\ -\ \sigma\sqrt{T-t}$$

where:

- $C(S_t, t)$ = the price of the European call option;
- S_t = the current price of the underlying asset;
- K = the strike price;
- r = the risk-free interest rate;
- $T - t$ = the time to maturity;
- σ = the volatility of the underlying asset;
- $N(\cdot)$ = the standard normal cumulative distribution function;
- ln = the natural logarithm; and
- e = Euler's number.

All else equal, a call on a more volatile asset is worth more to its holder (the option has positive vega).[4] It follows that a 20% incentive fee is expected to deliver a materially larger payoff to the manager when the underlying position is more volatile.

The table below values a call on two otherwise-identical assets—one with 10% volatility, the other with 20%—at the money, over a one-year horizon. The value of the higher-volatility call is very nearly twice that of the lower-volatility call, and so, therefore, is the expected incentive fee it generates:

	Lower-volatility asset	**Higher-volatility asset**
Underlying price = strike (S = K)	$100	$100
Volatility (σ)	10%	20%
Time to maturity (T)	1 year	1 year
Risk-free rate (r)	0%	0%
Value of the call per $100 invested	**$3.99**	**$7.97**

[4]The observation that a call option is more valuable on a more volatile underlying is a direct property of the Black–Scholes–Merton model (Black and Scholes, 1973; Merton, 1973): call value is strictly increasing in σ (positive vega).

	Lower-volatility asset	Higher-volatility asset
Expected incentive fee at 20% carry	$0.80	$1.59
Expected fee relative to low-vol case	1.00×	**2.00×**

Note. Values are Black–Scholes call prices for S = K = $100, T = 1 year, and r = 0, expressed per $100 invested; the expected incentive fee is 20% of the call value. Setting r = 0 is the natural convention when performance is measured on a relative or excess basis, and it makes the near-linearity of at-the-money option value in σ transparent: doubling volatility roughly doubles the fee. A positive risk-free rate narrows the ratio only modestly (to about 1.8× at r = 2%).

A natural objection is that only historical volatility is known ex ante.[5] For a special opportunity in a publicly traded single name, however, the market supplies a forward-looking estimate directly: the implied volatility embedded in the security's exchange-traded options is precisely the market's ex ante view of the volatility that will govern the fee. The argument does not depend on backward-looking data.

Because special opportunity investments are often concentrated single-stock positions, it is logically uncomfortable that an investor would accept a 20% incentive fee on a diversified portfolio and, at the same time, a 20% incentive fee on a single stock—when the latter also exposes the investor to greater risk,[6] not merely a greater expected fee. On realistic volatilities, it would not be unreasonable for an investor to bear more than double the incentive fee on a special opportunity investment purely on account of the volatility difference.

2.2 The Countervailing Consideration: Expected Return

The volatility argument is incomplete on its own, and honesty requires stating the counterpoint plainly. An incentive fee's expected value rises with volatility *holding the drift of the underlying constant*. But special opportunities are marketed precisely on the premise of a higher expected return—a higher drift, or genuine alpha. The economically correct question is therefore not whether the fee is higher (it is), but whether the incremental expected return compensates the investor for both the incremental risk *and* the incremental fee.

This reframing sharpens rather than dissolves the concern. If a manager can reliably identify investments with limited downside and substantial upside—return distributions materially more attractive than the normal case—then a flat 20% carry may be fair, or even a bargain, and the higher fee is simply the price of access to real skill. Absent demonstrable edge, however, the volatility-driven fee premium is a wealth transfer from investor to manager, dressed in the

[5] An immediate objection is that only historical volatility is directly observed, whereas the incentive fee's value depends on volatility over the life of the investment. For special opportunity investments in publicly traded single names this objection is largely answered by the market itself: the implied volatility embedded in the security's exchange-traded options is a forward-looking estimate of exactly the volatility that governs the value of the fee. The argument therefore need not rest on backward-looking data.

[6] A single-name position also exposes the investor to idiosyncratic risk that would otherwise be diversified away inside the fund; the investor bears more risk and, as shown below, pays a higher expected fee for the privilege.

language of alignment. The remainder of the paper asks which of these descriptions fits observed behavior, and why investors so rarely price the distinction.

2.3 The Break-Even Hurdle and a Fairness Proposition

The wedge in the table can be restated as a hurdle. At the illustrated parameters, the expected incentive fee on the concentrated position exceeds the expected fee on the fund by $0.79 per $100 invested over the year. Relative to leaving the same capital in the fund at the same headline rate, the special opportunity must therefore be expected to out-earn the fund by roughly eighty basis points, gross, merely to leave the investor's expected fee bill unchanged. That is the fee hurdle alone; it says nothing yet about risk.

The risk hurdle comes on top, and it is the larger term. The concentrated position exposes the investor to idiosyncratic risk the fund would diversify away, and estimates of the cost of capital for holders of undiversified single-company positions run to a multiple—on the order of two to four times—of the rate appropriate for a diversified investor.[7] The premium a special opportunity must credibly promise, before a flat carry is fair, is therefore far larger than the fee hurdle by itself suggests.

This yields a compact statement of when the customary terms are defensible. *Proposition: a flat, fund-level carry on an isolated position is fair only if (i) the position's expected gross return exceeds the fund's by at least the incremental expected fee plus the premium its undiversified risk requires, and (ii) the investor can verify condition (i) ex ante, from evidence available at the time of commitment.* Where both conditions hold, the fee is the price of genuine skill and the arrangement needs no apology; where either fails, the concentration premium is a transfer. The next section asks why the market does not price this distinction—why neither side ever seems to renegotiate the headline rate—and Section 4's remedies are precisely those that remain sensible when condition (ii) fails, as we will argue it generally does.

3. Why the Concentration Premium Is Not Competed Away

If the flat carry systematically overcharges concentrated bets, the natural question is why bargaining has not repaired it. Fees are negotiated by sophisticated parties, and a persistent mispricing needs a mechanism. We argue the premium survives because errors on the two sides of the table reinforce one another: the investor overvalues and over-forecasts the position's right tail; the manager sincerely believes in it; and the verification that would discipline both—condition (ii) of the proposition—is generally unavailable.

3.1 The Draw of the Big Payout

[7] Kerins, F., J. K. Smith, and R. Smith (2004), "Opportunity Cost of Capital for Venture Capital Investors and Entrepreneurs," Journal of Financial and Quantitative Analysis 39(2), 385–405, estimate that the opportunity cost of capital for the holder of an undiversified position in a single early-stage company is on the order of two to four times that of a well-diversified investor in the same asset.

Begin with what sells the investment—and with a precise statement of the claim, because it is easily caricatured. The claim is not that investors miscalculate the expected value: no one believes each dollar committed will return some enormous multiple, and nobody prices the modal outcome as a home run. The claim is that the *prospect* of a large payoff is crucial to the draw of these investments, and that this prospect enters the decision with a weight out of proportion to its probability. Formally, this is probability weighting: under cumulative prospect theory, decision weights on small-probability extreme gains exceed the probabilities themselves, so the right tail of a positively skewed payoff looms larger in choice than it does in expectation.[8] The bias is directional, and it needs no assumption of innumeracy to operate.

Nor does the evidence for it rest on lotteries. In public equity markets—at win probabilities orders of magnitude more favorable than any lottery's—stocks with lottery-like payoffs—high skewness, extreme past maximum returns—are systematically overpriced and subsequently underperform.[9] The mechanism therefore operates at exactly the odds a special opportunity offers: a meaningful chance at a multiple, not an astronomically remote one. The same taste shows up in the willingness to hold insurance and lottery exposure simultaneously,[10] and the pull strengthens when the bet is chosen rather than assigned[11]—and a special opportunity is always opted into.

The draw and the overcharge are the same object. The incentive fee is a call option, and a call option's value is concentrated in the right tail; the fatter the tail that attracts the investor, the more the flat twenty percent is worth to the manager. The feature that seduces is, dollar for dollar, the feature charged for most heavily—while the charge is felt least, because it falls due only in the states the investor is happiest to imagine. Probability weighting explains why investors overvalue the tail; the next two subsections describe the beliefs that inflate its perceived probability.

3.2 Narrow Framing and the Inside View

Under a traditional utility function defined over consumption or total wealth, an agent evaluates a new gamble by first combining it with the risks he already faces and then judging whether the combination is attractive. Narrow framing, documented in experimental settings, is the tendency

---

[8] Tversky, A., and D. Kahneman (1992), "Advances in Prospect Theory: Cumulative Representation of Uncertainty," Journal of Risk and Uncertainty 5(4), 297–323. Under the cumulative-prospect-theory weighting function, decision weights on small-probability extreme gains exceed the probabilities themselves. For the magnitude of the tilt in its limiting case, see Lockwood, B. B., H. Allcott, D. Taubinsky, and A. Sial (2025), "What Drives Demand for State-Run Lotteries? Evidence and Welfare Implications," Review of Economic Studies 92(4), 2578–2623: lottery sales respond strongly to the jackpot but are unresponsive to variation in smaller prizes, and the structural decision weight players attach to the jackpot is estimated at more than two hundred times its objective probability.

[9] On the overpricing and subsequent underperformance of positively skewed, "lottery-like" securities, see Kumar (2009), Barberis and Huang (2008), and Bali, Cakici, and Whitelaw (2011).

[10] The simultaneous purchase of insurance and lottery tickets—risk aversion and risk seeking in the same individual—is the classic puzzle rationalized by the doubly-inflected utility function of Friedman and Savage (1948).

[11] Langer, E. J. (1975), "The Illusion of Control," Journal of Personality and Social Psychology 32(2), 311–328. Subjects permitted to choose their own lottery ticket demanded a substantially higher price to relinquish it than subjects assigned a ticket at random.

to instead evaluate a new gamble in isolation—as if utility flowed directly from the outcome of that single gamble rather than from its contribution to overall wealth.[12]

If investors are already inclined to think about investments in isolation, fee structures that treat each investment as an isolated bet will reinforce the bias rather than mitigate it. The very fact that these investments are discussed in isolation, incentivized in isolation, and measured in isolation deepens their perceived separateness.

To see the magnitude of the resulting inconsistency, consider the following proposal. From the manager's entire portfolio, the investor and manager agree to single out one stock for which the manager may charge a separate fee. This stock carries equal or higher risk and equal or lower liquidity than the rest of the portfolio; its expected return is no higher than the rest of the portfolio; and the performance fee on its profits equals the fee charged on the remainder of the manager's diversified fund, irrespective of how the rest of the portfolio performs. Framed this way—with the higher-expected-return justification deliberately removed, so as to isolate the volatility effect—the proposal is decidedly unattractive.[13] That the identical economic arrangement provokes a different reaction under a different frame is itself evidence that framing shapes the investor's perception.

A close cousin of narrow framing operates on the forecast itself. Kahneman and Lovallo distinguish the "inside view"—building a prediction from the particulars of the case at hand—from the "outside view," which asks what happened, on average, to comparable cases.[14] Forecasters overwhelmingly take the inside view, and the very label "special opportunity" is an instruction to do so: each deal is presented as sui generis—this company, this catalyst, this discount—which is precisely an invitation to abandon the reference class. Yet the reference class exists, and its record is sobering. For the private placement in our own motivating example the evidence is direct: firms issuing PIPEs tend to be those unable to raise capital through conventional channels, and they significantly underperform in the two years after issuance; the sophisticated investors who profit do so through discounts, repricing rights, and hedges that an outside investor in a side pocket rarely enjoys.[15] Narrow framing isolates the bet from the portfolio; the inside view isolates it from its base rate. The word "special" does the work in both.

3.3 Anchoring on the Headline Rate

[12]Barberis, Huang, and Thaler (2006), "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing," American Economic Review 96(4), 1069–1090.

[13]The evidence here is anecdotal, drawn from informal conversations with a small and non-representative sample of investors; it is offered only to illustrate the framing effect, not as empirical support.

[14]Kahneman, D., and D. Lovallo (1993), "Timid Choices and Bold Forecasts: A Cognitive Perspective on Risk Taking," Management Science 39(1), 17–31.

[15]Brophy, Ouimet, and Sialm (2009). Firms issuing PIPEs tend to have poor fundamentals and pronounced informational asymmetries, and to be constrained from raising equity through conventional channels; they significantly underperform in the two years following issuance. The hedge funds that profit from PIPEs do so through price discounts, repricing rights, and short hedges.

Anchoring is a well-documented bias in which estimates are pulled toward an initial value:

> *In many situations, people make estimates by starting from an initial value that is adjusted to yield the final answer… adjustments are typically insufficient.*[16]

That managers select the identical rate—20% of profits—as the fee on a diversified portfolio suggests investors are anchored to that number, applying the heuristic that any investment by the manager is fairly priced at 20% of profits, without adjusting for the difference in volatility.

3.4 Extrapolating Skill from Noisy Track Records

The beliefs above concern the deal; a second family concerns the manager. Special opportunities are characteristically offered by managers coming off strong runs, and the investor's willingness to pay fund-level carry on a concentrated bet rests heavily on the inference that past success reveals skill. The founding evidence for that inference—persistence in private equity fund performance—has aged badly where it matters most.

At the deal level, persistence among buyout managers has declined to statistical insignificance for post-2001 vintages, and fund-level estimates show the probability that a top-quartile buyout firm repeats in the top quartile falling from roughly forty percent in early vintages to the mid-twenties—indistinguishable from chance—as the asset class matured.[17]

Even where some persistence survives, the signal is faint. Decomposing performance into skill and luck, the noise is large enough that a limited partner cannot statistically distinguish a skilled top-quartile firm from a lucky one on the strength of one or two funds; identifying skill with reasonable confidence requires observing more funds than most relationships ever produce.[18] The verification demanded by condition (ii) of our proposition is, in practice, unavailable at precisely the moment the special opportunity is offered.

Investors nonetheless act on the extrapolated belief, and pay for it: institutional limited partners systematically re-commit to underperforming sponsors, and capital raised in the wake of strong performance goes on to underperform.[19] The belief is strongest exactly when prospective returns are weakest.

3.5 The Supply Side: Overconfident Sellers and Adverse Selection

---

[16] Tversky, A., and D. Kahneman (1974), "Judgment under Uncertainty: Heuristics and Biases," Science 185(4157), 1124–1131, at p. 1128.

[17] Braun, Jenkinson, and Stoff (2017) find deal-level persistence among post-2001 buyout vintages statistically indistinguishable from chance; Harris, Jenkinson, Kaplan, and Stucke (2020) document the probability that a top-quartile buyout firm repeats in the top quartile falling from roughly 0.42 in pre-2001 vintages to roughly 0.26 thereafter.

[18] Korteweg and Sorensen (2017). Although the estimated spread in expected returns between top- and bottom-quartile private equity firms is substantial, performance is noisy enough that identifying the firms with top-quartile expected performance requires observing an impractically large number of past funds.

[19] Lerner, Schoar, and Wongsunwai (2007); Kaplan and Schoar (2005), who also document that capital raised in the wake of strong industry performance goes on to underperform.

Might there nonetheless be a rational account of the seller's side? Where the buyer is himself an agent deploying other people's money—a consultant, say—he may derive a marketing benefit from "buying expensive lottery tickets," because his own end-clients perceive positive utility in such investments for the very reasons of Section 3.1. It is then plausible that the marketing benefit to the agent outweighs the modestly lower returns the investments are expected to deliver.

A behavioral account may nonetheless dominate, and it is no longer mere assertion: success manufactures overconfidence. Self-attribution bias leads decision-makers to credit wins to skill and losses to circumstance, so a run of good outcomes mechanically inflates confidence; the pattern is documented among frequent corporate acquirers, whose early success predicts more, and worse, subsequent deals.[20] A manager launching a special-opportunity vehicle after a hot streak is therefore launching it at the peak of his own sincerity.

The supply of these vehicles is also adversely selected. Empirically they have a "last days of Rome" quality: they are launched disproportionately by managers late in a strong run—managers who are, for that very reason, most exposed to subsequent mean reversion and to the behavioral hazards of success, hubris included. The structure compounds the problem once an investment sours, because a discrete, separately branded bet is harder to unwind and to concede than a line item inside a diversified book; prospect theory predicts exactly this reluctance, since decision-makers become risk-seeking in the domain of losses.

Put the two sides together and the puzzle of Section 2 resolves. The premium is not competed away because both parties sincerely believe the same wrong number: the manager, made confident by his own past success, genuinely believes the alpha; the investor, extrapolating from the same record and forecasting from the inside view, over-infers it; and the one party with an incentive to reprice—the investor—cannot verify what he would need to. The mispricing is sustained by symmetric, sincere error rather than by deceit, which is why disclosure alone does not repair it, and why the fee terms of the next section are designed to be robust to it.

4. Desirable Characteristics of an Improved Fee Structure

If the investor cannot verify skill ex ante, the structure of the fee must carry the burden. Fee structures that encourage investors to think of discrete projects in terms of their combined utility—even when that utility is distributed across different time horizons—should deliver:

1. a greater perception of fairness;
2. better alignment of interests; and
3. greater long-term value to investors and managers where their interests overlap.

[20] Gervais and Odean (2001) model self-attribution formally: because decision-makers credit successes to skill and losses to circumstance, a run of good outcomes mechanically inflates confidence. Billett and Qian (2008) document the pattern in frequent acquirers, whose early deal success predicts more frequent and worse subsequent acquisitions.

One behavioral consideration cuts against these objectives and is worth flagging. The gambler's ruin describes the incentive to add further bets in order to "make up" prior losses.[21] Such mental accounting can distort decision-making when earlier special opportunity investments have lost money, prompting investors to double down rather than to appraise each new opportunity on its own merits.

5. Conclusion

We noted earlier that the return distribution of a special opportunity investment may depart from normality. Consider the extreme, hypothetical case of a manager able to identify investments with no downside risk and significant upside. In that limiting case the expected fee to the manager under a single cumulative fee structure equals the expected fee under a separate fee for each discrete investment—so cumulative measurement costs a genuinely skilled manager nothing in expectation, while protecting the investor when skill is absent. With that in mind, we offer the following suggestions to improve alignment, perceived and actual fairness, and, in turn, long-term value creation for manager and investor alike:

1. Price incentive fees on higher-volatility investments lower, to offset the effect of volatility on the expected fee.
2. Charge incentive fees on cumulative results—across all investments made since inception—rather than investment by investment.
3. Use clawbacks where practicable.
4. Where a clawback cannot practically be implemented, adopt, at a minimum, a high-water mark.

It seems logical that a fund with a genuine research edge in identifying successful special opportunities would want to encourage participation in as many future opportunities as possible—that is the long-term value-creation case—and would therefore be amenable to the suggestions above, practical accounting considerations aside. For firms without such an edge, by contrast, selling expensive lottery tickets may be a rational course of action, fraught with agency issues and, as often as not, sustained by the manager's own overconfidence. Caveat emptor.

[21] This is a form of mental accounting: prior losses on earlier special opportunity investments are treated as a separate "account" to be recouped, distorting the appraisal of the next opportunity on its own merits.

References

Agarwal, V., and N. Y. Naik (2004). Risks and portfolio decisions involving hedge funds. *Review of Financial Studies* 17(1), 63–98.

Anson, M. J. P. (2001). Hedge fund incentive fees and the "free option." *The Journal of Alternative Investments* 4(2), 43–48.

Bali, T. G., N. Cakici, and R. F. Whitelaw (2011). Maxing out: Stocks as lotteries and the cross-section of expected returns. *Journal of Financial Economics* 99(2), 427–446.

Barberis, N., and M. Huang (2008). Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review* 98(5), 2066–2100.

Barberis, N., M. Huang, and R. H. Thaler (2006). Individual preferences, monetary gambles, and stock market participation: A case for narrow framing. *American Economic Review* 96(4), 1069–1090.

Ben-David, I., J. Birru, and A. Rossi (2020). The performance of hedge fund performance fees. *NBER Working Paper* No. 27454 (rev. 2023).

Billett, M. T., and Y. Qian (2008). Are overconfident CEOs born or made? Evidence of self-attribution bias from frequent acquirers. *Management Science* 54(6), 1037–1051.

Black, F., and M. Scholes (1973). The pricing of options and corporate liabilities. *Journal of Political Economy* 81(3), 637–654.

Braun, R., T. Jenkinson, and I. Stoff (2017). How persistent is private equity performance? Evidence from deal-level data. *Journal of Financial Economics* 123(2), 273–291.

Brophy, D. J., P. P. Ouimet, and C. Sialm (2009). Hedge funds as investors of last resort? *Review of Financial Studies* 22(2), 541–574.

Carpenter, J. N. (2000). Does option compensation increase managerial risk appetite? *The Journal of Finance* 55(5), 2311–2331.

Friedman, M., and L. J. Savage (1948). The utility analysis of choices involving risk. *Journal of Political Economy* 56(4), 279–304.

Fung, W., and D. A. Hsieh (2002). Hedge-fund benchmarks: Information content and biases. *Financial Analysts Journal* 58(1), 22–34.

Gervais, S., and T. Odean (2001). Learning to be overconfident. *Review of Financial Studies* 14(1), 1–27.

Glosten, L. R., and R. Jagannathan (1994). A contingent claim approach to performance evaluation. *Journal of Empirical Finance* 1(2), 133–160.

Goetzmann, W. N., J. E. Ingersoll, Jr., and S. A. Ross (2003). High-water marks and hedge fund management contracts. *The Journal of Finance* 58(4), 1685–1718.

Grinblatt, M., and S. Titman (1989). Adverse risk incentives and the design of performance-based contracts. *Management Science* 35(7), 807–822.

Harris, R. S., T. Jenkinson, S. N. Kaplan, and R. Stucke (2020). Has persistence persisted in private equity? Evidence from buyout and venture capital funds. *NBER Working Paper* No. 28109.

Ineichen, A. M. (2003). *Absolute Returns: The Risk and Opportunities of Hedge Fund Investing*. Hoboken, NJ: John Wiley & Sons.

Kahneman, D., and D. Lovallo (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science* 39(1), 17–31.

Kahneman, D., and A. Tversky (1979). Prospect theory: An analysis of decision under risk. *Econometrica* 47(2), 263–291.

Kaplan, S. N., and A. Schoar (2005). Private equity performance: Returns, persistence, and capital flows. *The Journal of Finance* 60(4), 1791–1823.

Kerins, F., J. K. Smith, and R. Smith (2004). Opportunity cost of capital for venture capital investors and entrepreneurs. *Journal of Financial and Quantitative Analysis* 39(2), 385–405.

Korteweg, A., and M. Sorensen (2017). Skill and luck in private equity performance. *Journal of Financial Economics* 124(3), 535–562.

Kritzman, M., and D. Rich (1998). Risk containment for investors with multivariate utility functions. *The Journal of Derivatives* 5(3), 28–44.

Kumar, A. (2009). Who gambles in the stock market? *The Journal of Finance* 64(4), 1889–1933.

Langer, E. J. (1975). The illusion of control. *Journal of Personality and Social Psychology* 32(2), 311–328.

Lerner, J., A. Schoar, and W. Wongsunwai (2007). Smart institutions, foolish choices: The limited partner performance puzzle. *The Journal of Finance* 62(2), 731–764.

Litvak, K. (2009). Venture capital limited partnership agreements: Understanding compensation arrangements. *University of Chicago Law Review* 76(1), 161–218.

Lockwood, B. B., H. Allcott, D. Taubinsky, and A. Sial (2025). What drives demand for state-run lotteries? Evidence and welfare implications. *Review of Economic Studies* 92(4), 2578–2623.

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science* 4(1), 141–183.

Metrick, A., and A. Yasuda (2010). The economics of private equity funds. *Review of Financial Studies* 23(6), 2303–2341.

Tversky, A., and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. *Science* 185(4157), 1124–1131.

Tversky, A., and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty* 5(4), 297–323.