Compensating Fund Managers for *Risk-Adjusted* Performance

Thomas S. Coleman
Æquilibrium Investments, Ltd.

Laurence B. Siegel
The Ford Foundation

Journal of Alternative Investments
Winter 1999

In contrast to traditional fund investment, managers of hedged asset pools have typically received two forms of compensation: a flat fee, and a performance fee consisting of a percentage of profits in excess of a benchmark.[1] The flat fee (typically 1% to 2% of assets per annum) is intended to compensate the manager for his or her costs, and the performance fee (typically 15% to 20% of profits) provides an incentive for the manager to provide the highest possible return per dollar invested.[2]

While various agency issues exist in regard to performance based compensation plans, two issues with regard to the standard performance fee structure stand out. The first centers around aligning investor preferences regarding risk and return with managers' incentives. In the case where an investor is risk averse, the standard fee does not align the manager's incentives with investor preferences since the standard fee rewards higher returns with no reference to volatility or risk. The second issue is asymmetry – the fee is positive only and can create an option-like transfer of expected value from the investor to the manager. For example, the manager can invest in a promising but risky strategy. If the strategy works, both investor and manager benefit. If the strategy fails, however, the investor loses his capital while the manager loses only fees.[3] This asymmetry may provide an incentive for the manager to add to the risk of the fund by "leveraging up" the underlying portfolio positions.

In this article, our goal is to explore a risk-adjusted performance fee structure which addresses incentive compatibility and helps reduce asymmetry, while at the same time being feasible and easy to implement.[4] The result is a practical performance fee structure that principally compensates the manager for *risk-adjusted* performance by adjusting for the volatility of returns. For a fund manager this can provide a credible way to offer a volatility target as well as a return target, since the manager's compensation is contingent on realized volatility. For risk-averse investors this opens up the opportunity to choose managers based on the manager's incentives to meet risk as well as return objectives.[5]

Thomas S. Coleman is director of Æquilibrium Investments Ltd. in London. (63 Curzon Street, London W1Y7PE, UK. Regulated by IMRO.) Laurence B. Siegel is director of investment policy research at The Ford Foundation. (320 East 43rd Street, New York, NY 10017)

WHAT IS RISK-ADJUSTED PERFORMANCE?

To measure the risk-adjusted performance of a portfolio we start with the Sharpe ratio (S), which is the portfolio's excess return above the riskless or benchmark rate divided by the standard deviation of portfolio returns:

$$S = \frac{r_p - r_f}{s_p} \tag{1}$$

where r_p is the arithmetic mean rate of return on portfolio *p*, r_f is the riskless or benchmark rate of return, and s_p is the standard deviation of returns on portfolio *p*.

For a given underlying strategy, this measure will always be the same no matter what the leverage of the fund, as long as leverage is added or removed by borrowing or lending at the riskless rate. For example, take a strategy with an excess return of 5% and a volatility of 6.667%. This will have a Sharpe ratio of 0.75. By levering 2:1 (borrowing $100 at the benchmark rate for every $100 in the strategy) the excess return can be doubled from 5% to 10%.[6] The volatility will also double to 13.33%, leaving the Sharpe ratio unchanged at 0.75.

The fundamental idea behind our fee structure is that a manager should be compensated for returns earned through skill or effort rather than earned simply by levering the portfolio, or at a minimum the incentive to lever should be controlled. This leads to the objective of paying a manager an incentive fee that depends in some manner on the risk-adjusted performance and which is invariant to leverage – that is, the manager should be paid more for increased (excess) return only if volatility rises by less than return. Since the Sharpe ratio as a measure of risk-adjusted performance of a portfolio does not change with leverage it is a natural point of departure.

In what follows it will also be useful to transform the Sharpe ratio into M^2, which is a measure of risk-adjusted return. Following Modigliani and Modigliani [1997]:[7]

$$M^2 = s_m S + r_f = (r_p - r_f) \cdot \frac{s_m}{s_p} + r_f \tag{2}$$

where s_m is the reference or benchmark standard deviation, for example the standard deviation of returns on the market portfolio or benchmark being used to evaluate portfolio p.

While M^2 may be viewed simply as a linear transformation of S and either could be used equally effectively to calculate the performance fee, the M^2 measure has some interesting properties that make it attractive:

- M^2 expresses risk-adjusted performance as a *return*, unlike the Sharpe ratio which is unitless and can be difficult to interpret.

- More precisely, M^2 is the annual rate of return that the portfolio would have earned if it had been leveraged up or down to have the same risk as the benchmark. This can easily be seen by rewriting Equation (2) as

$$M^2 = r_f + \frac{s_m}{s_p}\left(r_p - r_f\right) \qquad (3)$$

  That is, the reward for taking risk, $[r_p - r_f]$, in the original portfolio is multiplied by the leverage ratio, $[s_m / s_p]$.

- The functional form of M^2 is adaptable to risk measures other than standard deviation by substituting the desired risk measure for s_p and s_m in Equations (2) and (3) respectively.[8]

DEVELOPING A RISK-ADJUSTED FEE STRUCTURE

We now turn to specific risk-adjusted fee structures. The first is: fee strictly proportionate to risk-adjusted excess return. In other words, a fee which is strictly proportional to the Sharpe ratio:

$$FEE1 = \mathrm{f} \cdot \mathrm{S} = \mathrm{f} \cdot \left[\frac{\left(r_p - r_f\right)}{s_p}\right] \qquad (4)$$

where f is the "fee factor" that determines the actual monetary level of the fee. For purposes of illustration we choose the factor to be 0.012. We chose this so that for a portfolio volatility of 8% FEE1 is equal to 15% of the excess return (when $\sigma_p = 0.08$ then $f / \sigma_p = 0.012 / 0.08 = 0.15$).

This fee structure is attractive within a broad range of performance results because it provides the desired incentive for the manager to take risk only in the pursuit of higher risk-adjusted return. However, it is problematic because at low levels of risk, the fee can become prohibitively high in comparison to the actual return on the portfolio. Take for example a portfolio with a 6% annual return and 1% annualized standard deviation. (Assume that the riskless rate is 5%.) The Sharpe ratio would be (0.06-0.05)/0.01=1. Multiplying by 0.012 gives a performance fee of 1.2%, higher than the original excess return of 1%. Clearly such a fee would be impractical.

To solve this problem, we turn to a fee proportionate to risk-adjusted return for volatility above a threshold, but not to exceed a fixed percentage (say 15%) of actual excess return for low volatility. We cap the fee as a percent of excess return for low volatility, and set a volatility threshold. (We also impose a floor of zero, which incorporates the common condition that the manager does not pay the investor when performance is below the benchmark rate.) Below the threshold our risk-adjusted fee is the same as the standard performance fee (a fixed percentage of excess return). Above the threshold the risk-adjusted fee is proportional to the Sharpe ratio (as for FEE1). This removes incentives to generate higher returns through leverage when volatility is at or above the threshold, since the manager does not receive extra compensation if he "purchases" extra return through leverage. The manager's compensation will increase only if the risk-adjusted return increases (that is, if excess return increases faster than volatility) – and this cannot be accomplished through leverage alone.

This risk-adjusted fee also provides a setting where investors can use "homemade leverage" if they desire returns (and volatility) above or below what the manager targets. The manager can target a given level of volatility and manage the fund to achieve the highest attainable risk-adjusted return, while the investor can leverage or de-leverage to adjust the level of volatility to match his own preferences.

Mathematically, the fee is:

$$FEE\,2 = \begin{cases} \max\left(0, (\mathrm{p} \cdot s\,*) \cdot \mathrm{S}\right) = \max\left[0, (\mathrm{p} \cdot s\,*) \cdot \left(r_p - r_f\right) / s_p\right] & s_p > s\,* \\ \mathrm{p} \cdot \max\left\{0, \left(r_p - r_f\right)\right\} & s_p \le s\,* \end{cases} \qquad (5)$$

where $s*$ is the volatility threshold below which the standard performance fee is applied, p is the standard performance fee percentage applied, and $(\mathrm{p} \cdot s*) = \mathrm{f}$ is the "fee factor" from above applied to the Sharpe ratio.

For illustration purposes we have chosen 0.08 as the volatility threshold, 0.15 as the standard performance fee percentage, and thus 0.012 as the Sharpe ratio fee factor. Equation (5) shows that FEE2 is equal to the standard fee (15% of excess return) when the volatility is 8% or less, but depends on the Sharpe ratio and is an inverse function of volatility when the volatility is above 8%.

Equation (5) can also be expressed as a function of M^2 (using $s*$ as the benchmark volatility in defining M^2):

$$FEE\,2 = \begin{cases} \max\left[0, \mathrm{p} \cdot \left(r_p - r_f\right) \cdot s\,* / s_p\right] = \max\left[0, \mathrm{p} \cdot \left(M^2 - r_f\right)\right] & s_p > s\,* \\ \mathrm{p} \cdot \max\left\{0, \left(r_p - r_f\right)\right\} & s_p \le s\,* \end{cases} \qquad (5a)$$

$$FEE\,2 = \begin{cases} \max\left[0, 0.15 \cdot \left(M^2 - r_f\right)\right] & s_p > 8\% \\ 0.15 \times \max\left\{0, \left(r_p - r_f\right)\right\} & s_p \le 8\% \end{cases} \qquad (5b)$$

Expressed in this way we can see that, for higher levels of volatility, the fee is 15% of the excess return after adjusting to an equivalent 8% volatility portfolio – that is to say, 15% of "excess" M2 using an 8% benchmark volatility. As mentioned above this risk-adjusted fee is particularly suited to targeting a predetermined level of volatility for a fund. Expressing the fee in terms of M2 shows how (for volatility above the threshold) returns are adjusted to an equivalent 8% volatility, reducing or removing incentives to construct a high volatility portfolio.

COMPARISONS ACROSS FEE STRUCTURES

The structure represented by FEE2 provides a realistic baseline to which other fee structures can be compared. The traditional performance fee and FEE2 provide two bounds:

- The traditional performance fee, which rewards for return with no adjustment for risk, provides an incentive to increase return with no reference to volatility.
- FEE2, which rewards for risk-adjusted return, provides an incentive to increase return *only* if doing so increases the risk-adjusted return (that is, M^2 or the Sharpe ratio).[9]

Exhibit 1 shows how the traditional performance fee, consisting of a fixed percentage of return in excess of the riskless rate, compares with the risk-adjusted performance fee (FEE2) as realized fund volatility varies.

Exhibit 1
Comparison of Traditional Performance Fees with Fee for Risk-Adjusted Performance
(Fees Expressed as Percentage of Return in Excess of Riskless Rate)

Annualized Fund Volatility	Traditional Performance Fee	Risk-Adjusted Performance Fee
8%	15.0%	15.0%
9%	15.0%	13.3%
10%	15.0%	12.0%
11%	15.0%	10.9%
12%	15.0%	10.0%
13%	15.0%	9.2%
14%	15.0%	8.6%
16%	15.0%	7.5%
18%	15.0%	6.7%
20%	15.0%	6.0%
22%	15.0%	5.5%
24%	15.0%	5.0%
26%	15.0%	4.6%
30%	15.0%	4.0%
35%	15.0%	3.4%

As mentioned earlier, actual performance fees are usually calculated as a percentage of the return in excess of the riskless rate, as shown in Exhibit 1. For example, if the return on a fund were 35% – say, net asset value (NAV) going from \$100 to \$135 – and the percentage rate were fixed at 15%, then the return in excess of the riskless rate would be 30% and the fee would be 4.5% (30% × 0.15, or \$4.50). Under FEE2, however, the percentage is adjusted by the realized volatility. If the return is 35% and the volatility is 8%, the percentage rate under FEE2 would be 15% (the same as with the traditional schedule), and the fee would still be 4.5% or \$4.50. However, if the return were 35% and the volatility were 20%, the percentage rate under FEE2 would be 6%, so the fee paid would be 1.8% (30% × 0.06, or \$1.80 if the NAV went from \$100 to \$135).

Between the two bounds of a fixed percentage rate and FEE2 (which rewards based on the Sharpe ratio or M^2), one can set up a fee structure that provides some incentive to increase return even when the Sharpe ratio goes down, on the assumption that investors might prefer higher returns even with a lower Sharpe ratio.[10] The important point is that fees structured in this manner allow one to tailor manager incentives to investor goals.

In evaluating fee structures, it is important to analyze carefully the incentives created by the structures, for investors as well as for managers. Consider the puzzle posed by two funds (A and B) which, ex ante, undertake an identical strategy except that fund B uses twice the leverage (thus earning twice the excess return with twice the volatility). They both use the same fee structure: no management fee, performance fee 15% of M^2. Ex-ante they target (promise to investors) the volatility and return profile in Exhibit 2

Exhibit 2
Comparison of Two Hypothetical Funds Having
Same Underlying Strategy but Differing in Leverage

	Fund A	Fund B
Volatility	10%	20%
Excess return	10%	20%
Sharpe ratio	1.0	1.0
M^2	8%	8%
Incentive fee	1.2%	1.2%

The apparent puzzle is that all investors would prefer Fund B over Fund A as long as investors can leverage and deleverage on their own. An investor who puts half his portfolio in cash at the riskless rate and half in Fund B would earn the same risk and return as putting his whole portfolio in Fund A, but would pay only half the fee (0.6% instead of 1.2%).

In practice, it is unlikely that the investor could identify Funds A and B such that a portfolio of Fund B and cash is a close substitute for Fund A. This is because markets are incomplete (there is a limited number of managers, each pursues a different strategy, and most do not offer risk-adjusted performance fees), and because information and search costs are high. Moreover, if fees roughly reflect costs to the manager, so that economic profits tend toward zero in the long run, Funds A and B will not co-exist in the same economy indefinitely; either B will lose money and go out of business or A will be forced to lower its fees. The puzzle, however, demonstrates that under circumstances that are at least theoretically plausible, risk-adjusted performance fees can motivate investors to "game" the fees.

Our puzzle brings into focus several characteristics of the risk-adjusted performance fee structure. First, if the structure becomes widespread or (at a minimum) an important element of fee negotiation, investors will receive an additional "degree of freedom" in fee negotiation, at the expense of the manager community. Such a dynamic would tend to lower the overall level of fees over time. Second, our fee structure helps a manager target a desired level of volatility because it provides a financial incentive to hit it – that is, the fund manager will not have an incentive to increase volatility above the threshold level, unless it can be accomplished with increased risk-adjusted returns (a higher Sharpe ratio or higher M^2).

The risk-adjusted performance fee structure that we have presented does not, of course, solve all of the problems inherent in contracting between investors and hedge fund managers (or other types of investment managers). It does, however, remove one obvious source of incentive to take uncompensated risk, so that one is free to focus on other, more subtle problems.

SKILL OR LUCK?

While it is appropriate to pay managers a performance fee for achieving favorable risk-adjusted returns, such returns can conceivably result from luck as well as skill. There is no sure way to distinguish one from the other, but a simulation approach can provide some insight into the question.

We conducted an experiment in which we assumed that the return-generating process is random (that is, due purely to luck), with a mean annual return of 0% in excess of the riskless rate, and an annualized standard deviation of 6%. The expected Sharpe ratio is thus 0/0.06 = 0, corresponding to an M^2 of 5% (the riskless rate). In each run of the simulation, we constructed 52 weekly returns, representing a reasonable time-period length and measurement frequency for calculating hedge fund performance fees. After 10,000 runs, the distribution of realized Sharpe ratios and risk-adjusted returns (M^2s) is as shown in Exhibit 3. We also show a performance fee calculated as 15% of the M^2 in excess of the riskless rate, but not less than zero.

Exhibit 3
Distribution of Simulated Results of Random Return-Generating Process

	Sharpe ratio	M^2	Performance fee
Mean	0.002	5.02%	0%
Standard deviation	1.023		
5th percentile	-1.688	-8.5%	0.0%
10^{th} percentile	-1.306	-5.4%	0.0%
50^{th} percentile	0.004	5.03%	0.005%
90^{th} percentile	1.303	15.4%	1.6%
95^{th} percentile	1.679	18.4%	2.0%

Note: The risk-adjusted return, or M^2, is calculated assuming that the standard deviation of the benchmark is 8% and the riskless rate is 5%.

Note that *all* of the returns in the distribution shown in Exhibit 3 are achieved through luck. Nevertheless, a manager would be paid a performance fee in half of the cases, and would be paid rather well (a 2% performance fee or higher) one time in twenty. The wide distributions revealed in this simulation suggest that neither managers nor investors should read too much into high achieved returns: even when they win big in one year, managers should not feel too brilliant, and investors should not feel too confident that their managers will repeat the performance in the next period. It is possible for such results to be achieved purely by chance.[11]

High-water marks

Goetzmann, Ingersoll and Ross [1998], among others, note that the high-water mark feature of hedge fund performance fee contracts, while appearing to limit fees and save the investor money, can represent a large transfer of expected value from investor to manager. (A high-water mark is a provision that makes the payment of a performance fee in a given period conditional on exceeding the maximum achieved share value in previous periods, for which the manager has already been compensated.) The transfer of expected value occurs because the manager is motivated to take risk when he is below the high-water mark so as to maximize the likelihood of exceeding the mark and earning a performance fee. Such a motivation to take extra risk is costly to the investor for reasons we discussed earlier. While Goetzmann et al. do not numerically separate the option value of the high-water mark provision from the option value of the performance fee absent a high-water mark, they demonstrate that such value exists and that it may be substantial.

Incentive issues within a single performance-measurement period

Even if there are no multi-period incentive considerations resulting from high-water marks, managers still face incentives within a single performance-measurement period to behave in ways that are not optimal for the investor. First, the manager is motivated to take excessive risk when "behind" part-way through the year, having no chance to earn a performance fee except by making risky investments. Second, the manager may be tempted to "shut down" and invest in riskless strategies when ahead (that is, having achieved a high M^2 or Sharpe ratio) part-way through the year, so he can lock in the performance fee.

The first problem cannot be addressed without relaxing the zero constraint on the performance fee. Since few managers will work for potentially negative fees, the problem remains. The second problem can be at least partially addressed by skewing the fee structure a little so that the manager gets paid more for achieving a higher return even if not accompanied by a higher M^2.

Downside risk vs. standard deviation

Many investors say they care only about downside risk and would not want a manager penalized for upside volatility. Harlow [1991] and Sortino and van der Meer [1991], among others, have proposed downside-risk measures that are constructed by summing squared deviations below, but not above, some cutoff rate of return. (The cutoff can be the *ex post* mean of the return distribution, or it can be a fixed target such as 0%.) If downside risk is the component of risk to which investors are averse, then using a downside risk measure in equation (5b) to calculate performance fees is acceptable in principle. However, investors differ in their utility functions and consequently in the cutoff they would want used for designing their downside risk-based fee structure. (A downside risk-based performance fee would be highly sensitive to the choice of a cutoff.) Because all investors must agree to the same fee structure within a particular fund, it would be difficult to create a scheme agreeable to everyone (including the manager). Moreover, some observers have objected to the use of downside risk measures on the basis that upside volatility signals risk-taking on the part of the manager and should not be ignored.

CONCLUSION

This paper has presented a risk-adjusted performance fee structure that modifies standard hedge fund fees to better align the incentives of the manager and investors. The resulting fee provides incentives for the manager to control volatility by rewarding risk-adjusted returns, and allows the manager to credibly target a level of volatility. The specific form of the fee can be adjusted to alter the trade-off between compensating unadjusted (actual) versus risk-adjusted return. We do not claim that this fee structure solves all incentive problems in aligning investor goals with fund management incentives, but it does address some of the issues.

Bibliography

Chriss, Neal A., Peter R. Fanelli, and Leah Modigliani. "Risk adjusted return measures." *Global Equity and Derivative Markets,* Morgan Stanley Dean Witter (New York, NY), January 1998..

Goetzmann, William N., Jonathan Ingersoll, Jr., and Stephen A. Ross. "High water marks." Yale School of Management working paper, January 1998, http://papers.ssrn.com/paper.taf?ABSTRACT_ID=57933#PaperDownload

Harlow, V. Van. "Asset allocation in a downside risk framework." *Financial Analysts Journal,* September-October 1991.

Modigliani, Franco, and Leah Modigliani. "Risk-adjusted performance." *Journal of Portfolio Management,* Winter 1997.

Modigliani, Leah (1997a). "Are hedge funds worth the risk?" *Investment Strategy*, Morgan Stanley Dean Witter (New York, NY), December 12, 1997.

Modigliani, Leah (1997b). "Yes, you can eat risk-adjusted returns." *Investment Perspectives,* Morgan Stanley Dean Witter (New York, NY), February 1997.

Siegel, Laurence B., and Elizabeth R. Hilpman. "Market neutral investing: Industry background and prospects," Baring Asset Management (Boston, MA), 1998. Available from the authors.

Sortino, Frank, and Robert van der Meer. "Downside risk." *Journal of Portfolio Management,* Summer 1991.

Endnotes

The authors thank Leah Modigliani of Morgan Stanley Dean Witter, New York, for many valuable discussions about measuring risk-adjusted performance and about evaluating hedge fund and other manager returns using this measure. Richard Spurgin, of Clark University, Worcester, Mass., refereed the paper for the Journal of Alternative Investments and provided the basis for the section on incentives faced by the investor, and contributed many other valuable thoughts. Elizabeth Hilpman of GAM Global Asset Management, New York, has engaged us in conversations about every aspect of the hedge-fund and investment-management businesses over the years, and has graciously allowed us to adapt some of the material in Siegel and Hilpman [1998] for use in endnote 2 of this article. Andrea Pignataro of Salomon Smith Barney helped clarify the importance of investor preferences in choosing the type of fee structure. Clinton Stevenson of the Ford Foundation read several drafts of the paper and provided many helpful suggestions. In addition, we thank the rest of the "Hedge Fund Research Group" (Cliff Asness, Peter Bernstein, Ted Caldwell, and Robert Shultz, as well as Hilpman) for their intellectual and moral support. Any remaining errors or omissions are the responsibility of the authors.

---

[1] Originally the benchmark was zero return – that is, managers collected performance fees even on returns below the Treasury bill rate, as long as the returns were positive. Competitive forces have caused the benchmark, in many cases, to be raised to the "riskless" rate (the Treasury bill rate or LIBOR), and, sometimes, a fixed premium (200 to 500 basis points per annum) above that rate.

[2] This paper assumes that (some) managers should or, at any rate, will get performance fees and seeks only to determine how much, but a brief justification for paying such fees at all is likely to be helpful to some readers.

Principal-agent argument. Some argue that a flat fee is incentive enough for managers to provide the highest achievable rate of return. However, a performance fee is required to align investor and manager interests if the asset class or strategy in question has limited capacity. Management firms should be assumed to be profit (not rate-of-return) maximizers. A low-fee product can be profitable through asset gathering, but that is bad for the client in a limited-capacity strategy. As assets in the strategy become large, the effect the strategy is intended to exploit goes away and the rate of return (per unit of capital invested) is diluted.

Alternatively, the client can pay a performance-based fee, which can become quite high as a percentage of assets under management if the performance is good. This structure motivates the manager to limit, rather than expand, assets under management. If the manager's own money is also in the strategy, that is an even stronger motivator for high returns and limited asset gathering – but we would be reluctant to propose that a manager ignore the basic principles of prudence and diversification in personal investing in order to achieve a more perfect alignment of his interest with those of his clients. We urge clients, moreover, not to expect managers to put all their money in the strategy they manage.

Information asymmetry argument. Performance fees are also used to separate good managers from bad. Because investors have a difficult time identifying good managers, and all managers represent themselves as good, performance fees serve as a mechanism to enable investors to pay minimally for those bad managers who will inevitably be selected from time to time, while rewarding good managers appropriately. A side benefit of the performance fee is that bad managers will drop out of the opportunity set (because they are failing to earn adequate fees) more quickly than with a flat fee.

[3] The cost to the manager is not zero: If capital is withdrawn the manager will lose future as well as current fees, and if the fund closes the manager may suffer reputational costs. Nonetheless, the monetary cost to the investor likely outweighs the monetary cost to the manager, and reputational costs are difficult to measure and will differ across managers. Our motivation in devising a risk-adjusted fee structure is to provide an explicit contract which reduces or eliminates asymmetry rather than relying on unmeasured and unknowable incentives.

[4] By "feasible" we mean the fee must depend only on variables, such as realized return and volatility, that can be easily measured and independently verified. We have therefore not used variables such as leverage, which in practice is almost impossible to define rigorously. Because the fee structure under discussion is based on readily measured variables, it is easy to implement and indeed has been implemented for Æquilibrium Absolute Return Fund Plc., a designated investment company authorized by the Central Bank of Ireland and managed by Æquilibrium Investments Limited, of which one of us (Coleman) is a director.

[5] For certain investors, particularly hedge fund investors, the assumption of risk neutrality may be a good approximation. In this case the standard fee structure may be appropriate. Where the assumption of risk neutrality is not reasonable, as we believe will be the case for many investors, our fee will permit the rewarding of managers who manage risk as well as produce return. Comparison between fee structures, however, should be made after adjusting the absolute level of the fee to make them equal in expectation.

[6] The extra \$100 will earn 5% plus the riskless rate and the borrowing will cost the riskless rate, leaving an increase in return of 5%.

[7] The M^2 designation refers to the authors' names, M and M, and adds a touch of whimsy to this otherwise sober topic. Modigliani and Modigliani [1997] used the designation RAP (risk-adjusted performance) to refer to their statistic. In other work, Morgan Stanley Dean Witter used the more euphonious M^2 and it stuck.

[8] M^2 calculated using a risk measure other than standard deviation will retain the properties described above if the risk measure is chosen "such that an increase in leverage increases the risk and return of a portfolio in

the same proportion" (Modigliani and Modigliani [1997]). Semivariance below the mean of a distribution meets this criterion, as does beta, while semivariance below a fixed target (other than the mean) does not.

[9] For example, under *FEE1*, and in most cases under *FEE2*, the manager will be indifferent between an excess return of 10% with 10% volatility and an excess return of 20% with 20% volatility.

[10] For example, the fee structure could make the manager indifferent between an excess return of 10% with 10% volatility and an excess return of 20% with volatility of 22%.

[11] To make sure the results were not somehow distorted by the weekly simulation technique, we repeated the experiment using monthly returns instead. The results were almost identical. As one would expect, the distributions tighten when calendar time (as opposed to frequency of measurement) is lengthened; the chance of earning superior returns by chance is much less over five years than it is in one year. We are not suggesting, of course, that hedge fund managers wait five years before collecting a performance fee.