

Incentive Contracts and Hedge Fund Management

by

James E. Hodder

and

Jens Carsten Jackwerth

May 12, 2004

C:\Research\Paper21 Hedge Fund\Paper16.doc

James Hodder is from the University of Wisconsin-Madison, Finance Department, School of Business, 975 University Avenue, Madison, WI 53706, Tel: 608-262-8774, Fax: 608-265-4195, jhodder@bus.wisc.edu.

Jens Jackwerth is from the University of Konstanz, Department of Economics, PO Box D-134, 78457 Konstanz, Germany, Tel.: +49-(0)7531-88-2196, Fax: +49-(0)7531-88-3120, jens.jackwerth@uni-konstanz.de.

We would like to thank Günter Franke, Stewart Hodges, J. C. Hugonnier, Kostas Iordanidis, Pierre Mella-Barral, Antonio Mello, Mark Rubinstein, Paolo Sodini, Fabio Trojani, and seminar participants at Humboldt University, Stockholm School of Economics, University of Konstanz, University Svizzera Italiana, and University of Zurich for helpful comments on an earlier paper entitled "Pricing Derivatives on a Controlled Stochastic Process: A Simplified Approach".

Incentive Contracts and Hedge Fund Management

Abstract

This paper investigates dynamically optimal risk-taking by an expected-utility maximizing manager of a hedge fund. We examine the effects of variations on a compensation structure that includes a percentage management fee, a performance incentive for exceeding a specified high-water mark, and managerial ownership of fund shares. In our basic model, there is an exogenous liquidation barrier where the fund is shut down due to poor performance. We also consider extensions where the manager can voluntarily choose to shut down the fund as well as to enhance the fund's Sharpe Ratio through additional effort. We find managerial risk-taking which differs considerably from the optimal risk-taking for a fund investor with the same utility function. In some portions of the state space, the manager takes extreme risks. In another area, she pursues a lock-in style strategy. Indeed, the manager's optimal behavior even results in a bimodal return distribution. We find that seemingly minor changes in the compensation structure can have major implications for risk-taking. Additionally, we are able to compare results from our more general model with those from several recent papers that turn out to be focused on differing parts of the larger picture.

Incentive Contracts and Hedge Fund Management

Hedge funds have grown rapidly with assets under management ballooning from around \$50 billion in 1990 to \$600 billion in 2002.¹ As they have come to play a larger role in financial markets, there has been increasing attention focused on their management and investment practices. In that vein, we analyze how risk-taking by a hedge fund manager is influenced by her compensation structure. We have a single risk-averse manager who controls the allocation of fund assets between a risky investment and a riskless one. The manager's compensation can potentially include both a proportional management fee and an incentive fee based on exceeding a "high-water mark." We also consider the possibility that the manager has her own capital invested in the fund. In practice, a fund that performs poorly is frequently shut down and liquidated. We also include this influence on fund management via incorporating an exogenous liquidation boundary into the model as well as considering an endogenous shutdown decision by the manager.

Recognizing that a manager will control the hedge fund's investments, altering them through time, means the fund's value follows a controlled stochastic process. Since the manager's compensation is a payoff whose value depends on fund performance, we are effectively valuing a derivative on a controlled stochastic process. This generates significant challenges relative to more standard derivative valuation situations. We use a discrete-time framework to model the rebalancing decisions and develop a numerical procedure for determining the manager's sequence of optimal investment decisions. As discussed in the next section, that approach enhances realism and provides great modeling flexibility albeit at the cost

¹ From: "An Invitation from the SEC", *Economist*, vol. 367, No. 8324, May 17th, 2003, p. 63.

of losing the analytical tractability of a continuous-time model. Moreover, the basic approach developed here can be applied to valuing derivatives in other situations where a return process is controlled by a utility maximizing manager.

There is an important analogy with Merton (1969) who examines the optimal investment strategy for an expected utility maximizing individual who exercises continuous-time control over his own investment portfolio.² In Merton's model with constant relative risk aversion, the optimal proportion of wealth invested in the risky asset is a constant through time. Although optimally controlled, the associated wealth process evolves just like a standard geometric Brownian motion. There are circumstances where our hedge fund manager will follow the same strategy in discrete time. Those circumstances effectively amount to owning a proportional share of the fund with no other incentives or disincentives to influence the manager's behavior.

Importantly, an outside investor in the hedge fund with the same utility function as the manager would also find this solution optimal and desire that a constant proportion of the fund's capital be allocated to the risky investment. As we show, the manager's optimal strategy frequently differs substantially from that simple rule. This results in a striking contrast between the manager's optimal behavior and what our stereotypical outside investor would prefer.

Typically, hedge funds earn incentive fees for performance exceeding a high-water mark. This is analogous to a call option with the high-water mark corresponding to the strike price. As we shall see, that structure generates dramatic risk-taking below the high-water mark as the manager tries to assure that her incentive option will finish in-the-money. At performance levels modestly above the high-water mark, she reverses that strategy and opts for very low risk positions to "lock in" the option payoff. From the perspective of our outside investor, this is very perverse behavior.

² Merton's work in turn is based on Markowitz's (1959) dynamic programming approach and Mossin's (1968) implementation of that idea in discrete time.

Both managerial share ownership and the use of a liquidation boundary can play important roles in reducing the manager's risk-taking at modest distances below the high-water mark. How these aspects of the compensation structure interact is both interesting and important for thinking about incentives which do a better job of aligning the manager's interests with outside investors'. In that regard, we find that seemingly slight adjustments in the compensation structure can have enormous effects on managerial risk-taking. For example, even a relatively small penalty for hitting the lower boundary can eliminate risk-taking in the lower portions of the state space.

Several recent papers examine effects of incentive compensation on the optimal dynamic investment strategies of money managers. Carpenter (2000) and Basak, Pavlova, and Shapiro (2003) focus directly on this issue.³ Goetzmann, Ingersoll, and Ross (2003) focus primarily on valuing claims (including management fees) on a hedge fund's assets. Most of that paper assumes the fund follows a constant investment policy; however, one section briefly explores some limited managerial control of fund risk. These three papers all generate analytic solutions using equivalent martingale frameworks in continuous time. However, they generate seemingly conflicting results regarding the manager's optimal risk-taking behavior.

Although we pursue a different tack and use a numerical approach to determine the manager's optimal investment strategy, we are able to shed light on the differing results in the above papers by relating them to our own model. Perhaps not surprisingly, it turns out that these papers have (sometimes rather subtle) differences in how they model the manager's compensation structure. Again, some seemingly minor differences (e.g. continuous vs. discrete resetting of the high-water mark) have dramatic impacts on optimal risk-taking by the manager.

³ There are also related papers by Basak, Shapiro, and Teplá (2003), who investigate risk-taking when there is benchmarking, and by Ross (2004), who decomposes risk-taking according to three underlying causes.

It also appears that some simplifying assumptions used to generate analytic solutions result in leaving out important aspects of the problem. For example, Carpenter as well as Basak, Pavlova, and Shapiro ignore the possibility of the fund being shut down in response to poor performance. As we shall see, this possibility has important implications for managerial behavior.

In the next section, we present the basic model and briefly describe the solution methodology (more details are in the Appendix). Section II provides numerical results for a standard set of parameters. We actually begin our discussion with a simplified version of the model which is analogous to a discrete-time version of Mossin (1968) and Merton (1969). This allows us to build intuition as we add pieces of the compensation structure and examine the effects on managerial behavior.

Section III describes two extensions of our model including one where the manager can voluntarily choose to shut down the fund in order to pursue outside opportunities and/or avoid costs of continued operations. This is an American-style option which can be easily accommodated by our solution procedure. Both in practice and in our model, this is a realistic possibility if fund value is well below the high-water mark so that the manager's incentive option has low value. Our second extension is to allow the manager to enhance the fund's return distribution by exerting extra effort. The manager suffers a disutility from increased effort, and we investigate her optimal strategy for balancing the costs and benefits of effort exertion.

In Section IV, we compare our results with those from Carpenter (2000), Basak, Pavlova, and Shapiro (2003), and Goetzmann, Ingersoll, and Ross (2003). This is a useful exercise which allows us to see that these papers are effectively looking at different parts of a larger picture. It also helps our understanding of how different pieces of the compensation structure interact to influence risk-taking in various regions of the state space. Section V provides concluding comments.

I. The Basic Model and Solution Methodology

In modeling our hedge fund manager's problem, we attempt to introduce considerable realism while still retaining tractability. We will first address the stochastic process for the fund's value. Next, we discuss her compensation conditional on both upside performance and the possibility of fund liquidation at a lower boundary. Finally, we show how the manager optimally controls the fund value process to maximize her expected utility. Our approach utilizes a numerical procedure, with details on the implementation available in the Appendix.

A. The Stochastic Process for Fund Value

Assume that a single manager controls the allocation of fund value X between a riskless and a risky investment. The proportion of the fund value allocated to the risky investment is denoted by κ . We allow the manager to control κ , which is short for $\kappa(X,t)$. Think of the risky investment as a proprietary technology that can be utilized by the fund manager but is not fully understood by outside investors (and hence not replicable by them). The risky investment grows at a constant rate of μ and has a standard deviation of σ . The riskless investment simply grows at the constant rate r .⁴

The typical and mathematically convenient assumption is to model the fund value in continuous time as driven by a geometric Brownian motion for the risky investment. However, that approach inhibits modeling some important aspects of fund management. As a practical matter, many hedge funds are voluntarily shut down or forced to liquidate due to poor performance. We address this latter possibility by having a lower (liquidation) boundary.

⁴ The parameters μ , σ and r can be deterministic functions of (X,t) without generating much additional insight about managerial risk-taking.

However, in a continuous-time setting, the manager can always avoid liquidation since (by design) there is sufficient time to get out of any risky investment before hitting the lower boundary.

Related issues are human limitations as well as markets being closed which constrain trading frequency. This is in addition to the practical issue of transaction costs (which we do not model) that would make continuous-time trading financially unrealistic. Clearly, continuous-time trading is a simplifying assumption that greatly enhances analytical tractability. There is, however, a trade-off regarding both realism and modeling flexibility. We have opted to use a discrete-time framework where the manager can only change the risky investment proportion at discrete points in time. If the fund value is in the vicinity of the lower boundary, the manager can no longer pursue a risky strategy and avoid the risk of liquidation.⁵

For a given proportion allocated to the risky investment κ , we assume that the log returns on the fund value X are normally distributed⁶ over each discrete time step of length Δt with mean $\mu_{\kappa, \Delta t} = [\kappa\mu + (1 - \kappa)r - \frac{1}{2}\kappa^2\sigma^2]\Delta t$ and volatility $\sigma_{\kappa, \Delta t} = \kappa\sigma\sqrt{\Delta t}$. Most of the analysis in the paper uses time steps approximately equal to a trading day. However, we have also conducted runs with time steps of approximately 15 minutes ($1/32^{\text{nd}}$ of a trading day). That seems close to the maximum practical trading frequency, and our qualitative results were unchanged.

In order to proceed, we discretize the log fund values onto a grid structure (more details are provided in the Appendix). That grid has equal time increments as well as equal steps in log

⁵ Our approach also provides considerable flexibility in modeling, as well as the ability to solve free boundary problems such as the optimal endogenous liquidation decision of the manager which we analyze below.

⁶ For simplicity, we assume normality for log returns; however, our approach can accommodate alternative return distributions such as might be generated by a portfolio including option positions with their highly skewed returns.

X.⁷ To insure that a strategy of being fully invested in the riskless asset ($\kappa = 0$) will always end up on a grid point, we have points for the log fund value increase at the riskfree rate as time passes. From each grid point, we allow a multinomial forward move to a relatively large number of subsequent grid points (e.g. 121) at the next time step. We structure potential forward moves to land on grid points and calculate the associated probabilities by using the discrete normal distribution with a specified value for the control parameter kappa.

B. The Manager's Compensation Structure

We assume the manager has no outside wealth but rather owns a fraction of the fund. Frequently, a hedge fund manager has a substantial personal investment in the fund. Fung and Hsieh (1999, p. 316) suggest that this “inhibits excessive risk taking.” For much of our analysis, we will assume the manager owns $a = 10\%$ of the fund. That level of ownership, or more, is certainly plausible for a medium-sized hedge fund. A large fund with assets exceeding a billion dollars would likely have a substantially smaller percentage but still a non-trivial managerial ownership stake. On the remaining $(1-a)$ of fund assets, the manager earns a management fee at a rate of $b = 2\%$ annually plus an incentive fee of $c = 20\%$ on the amount by which the terminal fund value X_T exceeds the “high-water mark”. Such a fee structure is typical for a hedge fund.⁸

We use a high-water mark that is indexed so that it grows at the riskless interest rate during the evaluation period (a fairly common structure). Letting H_0 denote the high water mark at the beginning of an evaluation period with length T years, we have $H_0 e^{rT}$ at the period's end.

⁷ To economize on notation, we assume the fund value X and the time t are always multiples of $\Delta(\log X)$ and Δt without the use of indices.

⁸ See for example, Fung and Hsieh (1999) for a description of incentive fees as well as a variety of additional background information on hedge funds.

The manager is compensated based on the fund's performance if the fund is not liquidated prior to time T . Since the manager has no further personal wealth (or other income)⁹, her wealth at T equals her compensation and is equivalent to a fractional share plus a fractional call option (incentive option) struck at the high-water mark H_0e^{rT} :

$$W_T = aX_T + (1-a)bTX_T + (1-a)c(X_T - H_0e^{rT})^+ \quad (1)$$

A realistic complication is that if the fund performs poorly, it may be liquidated. The simplest approach is to have a prespecified lower boundary. Our basic valuation procedure uses this approach with the fund being liquidated if its value falls to 50% of the current high-water mark.¹⁰ Using Φ_t to denote the level of the liquidation boundary at time t , we set $\Phi_t = 0.5H_0e^{rt}$.

Now consider the manager's compensation if the fund value hits the lower (liquidation) boundary at time τ , with $0 \leq \tau \leq T$, and it is immediately liquidated. For the moment, we assume no dead weight cost to liquidation but do recognize that, in a discrete-time setting, the fund value may cross the barrier and have $X_\tau < \Phi_\tau$. Our base case assumption will be that the manager recovers her personal investment aX_τ plus a prorated portion of the management fee $\tau(1-a)b\Phi_\tau$. This total is reinvested until T at the riskless interest rate. This last step is because the manager's utility is defined in terms of time T wealth. This results in:

⁹ This assumption can readily be relaxed.

¹⁰ Apparently such liquidation boundaries are sometimes contractual and sometimes based on an unwritten understanding between the fund management and outside investors. Goetzmann, Ingersoll, and Ross (2003) also use a prespecified liquidation boundary based on the high-water mark.

$$W_T = aX_\tau e^{r(T-\tau)} + 0.5(1-a)b\tau H_0 e^{rT} \quad \text{for } 0 \leq \tau \leq T \quad (2)$$

where this value depends on when the fund reaches the boundary and by how much it crosses that boundary. Note, however, that once the boundary has been reached or crossed, we know X_τ and τ so the terminal payoff in (2) is certain. An obvious alternative to (2), which we will also consider, is that the manager receives a smaller amount due either to some liquidation costs or an explicit penalty built into the fee structure. In any case, we will refer to the payment the manager receives if the fund hits the liquidation boundary as her severance compensation.

As we shall see shortly, the lower (liquidation) boundary plays an important role in determining the manager's optimal portfolio allocations over time. Failure to consider such a boundary when modeling managerial behavior leads to very different and potentially seriously misleading results.

C. The Optimization of Expected Utility

We assume the manager seeks to maximize expected utility of terminal wealth W_T and has a utility function that exhibits constant relative risk aversion γ (an assumption that can readily be relaxed):

$$U(W_T) = \frac{W_T^{1-\gamma}}{1-\gamma} \quad (3)$$

For each terminal fund value, we calculate the manager's wealth and the associated utility. We then step backwards in time to $T-\Delta t$. At each possible fund value within that time step, we calculate the expected utilities for investment proportions (κ) in our discrete choice set (κ can

be zero or lie at specified steps between 0.2 and 10, details on that set are in the Appendix). We choose the highest of those expected utilities as the optimal indirect utility for that fund value and denote its value as $J_{X,T-\Delta t}$. We record the optimal indirect utilities and the associated optimal kappas for each fund value within that time step and then loop backward in time, repeating this process through all time steps. This generates the indirect utility surface and optimal kappa values for our entire grid. Formally:

$$J_{X,T} = U_{X,T}; \quad J_{X,t} = \max_{\kappa} E_{\kappa}[J_{X,t+\Delta t}] \quad (4)$$

where t takes the values $T - \Delta t, \dots, 2\Delta t, \Delta t, 0$ one after another.

II. Some Illustrative Results

We will frequently refer to a standard set of parameters as displayed in Table 1, which we will use as our reference case. The horizon is three months with portfolio revisions in 60 time steps, roughly once per trading day. For this reference case, the starting fund value of 1 equals the current high-water mark. We can think of the risky investment as a typical trading strategy employed by a hedge fund (e.g. convergence trades or macro bets). On an unlevered basis, we assume that the risky investment has a mean return of 7% and a volatility of 5%. The riskless asset yields 5%. This combination of mean returns and volatility would be consistent with a market-neutral strategy and implies a Sharpe Ratio of 0.40, which seems reasonable in light of the results reported in Brown, Goetzmann, and Ibbotson (1999). There are a total of 1200 log steps between the lower and upper boundaries with the initial fund value X_0 centered in that space. The risk aversion coefficient of the manager's power utility is $\gamma = 4$.

Table 1
Standard Parameters

Time to maturity	T	0.25	Interest rate	r	0.05
Log value steps below/above X_0		600/600	Initial fund value	X_0	1.00
Risk aversion coefficient	γ	4	Mean	μ	0.07
Number of time steps	n	60	Volatility	σ	0.05
Initial high-water mark	H_0	1.00	Incentive fee rate	c	0.20
Exit boundary at t=0	Φ_0	0.50	Basic fee rate	b	0.02
Manager's share ownership	a	0.10			
Future nodes for the Normal approx.		$1+2 \times 60 = 121$			
Log X step		$(\log(1/0.5))/600 \approx 0.001155$			

A. The Effect of the Liquidation Barrier


Before displaying results for our reference case, it is useful to build some intuition by examining a sequence of simpler situations. In Merton (1969), an individual (analogous to our manager) dynamically chooses the optimal allocation of available funds between shares and the riskless asset. In the case where there is no intermediate consumption (between 0 and T), she chooses that investment strategy to maximize her expected utility of terminal wealth (W_T). Merton's analysis is in continuous time (as opposed to our discrete-time framework); however, that description otherwise matches the situation of our manager if she had no incentive option and there was no liquidation boundary. In Merton's framework, the optimal proportion allocated to the risky investment would be constant and using our standard parameters implies:

$$\kappa = \frac{(\mu-r)}{\gamma\sigma^2} = 2. \quad (5)$$

Our model also generates a flat optimal kappa surface at $\kappa = 2$ when there is no liquidation boundary or incentive option. Thus, our discrete-time analog of Merton's analysis generates the same solution. That is not surprising since optimally allocating a constant proportion to the risky investment does not exploit the rebalancing capability (in either discrete or continuous time). This changes dramatically when we add the liquidation boundary.

Figure 1. Optimal Risky Investment Proportion (κ) with No Incentive Option and No Managerial Share Ownership

In this figure, the manager receives as compensation only a management fee ($b = 2\%$) and has neither an incentive option ($c = 0$) nor an equity stake ($a = 0$). Other parameter values are as specified in Table 1.



The liquidation boundary effectively turns the manager's compensation function into a knockout call with a rebate equal to the severance compensation of equation (2). The manager's

optimal kappa levels are depicted in Figure 1 for our standard parameters except for setting $a = c = 0$.¹¹ That is, the manager has neither ownership in the fund nor an incentive option.

The manager exhibits essentially four different areas of economic behavior. Two of these areas are intuitively rather straightforward. In “Merton Flats” to the right in Figure 1, the manager’s optimal kappa recedes to the continuous-time Merton solution, which in the current case is 2. This represents an area where fund value is far enough from the liquidation boundary (given the time left to T) that it plays essentially no role in her decision making.

“Gambler’s Ridge” in the far left corner of Figure 1 is also not surprising. Here the manager is in a situation just prior to T that could be described as “heads: I win, tails: I don’t lose very much.” She is thus willing to gamble with a very large kappa. In practice, there would be limits on a fund’s ability to undertake leverage. In our case, we excluded kappa values where we did not get a good approximation for the normal distribution -- the maximum available kappa here is only 10. Nevertheless, her gambling behavior is pronounced.

More interesting and perhaps more surprising are the “Valley of Prudence” toward the left boundary and the “Hill of Anticipation” toward the center of Figure 1. The Valley of Prudence can be interpreted as a region where the manager chooses a very low kappa (zero or only slightly higher) in order to dramatically reduce the chance of hitting the liquidation boundary at an early date.¹² Hitting that boundary early incurs an implicit cost since the manager is now unable to improve on her severance compensation by managing the portfolio a little longer. Approaching the terminal date, the remaining potential for her gaining from continuing to manage the portfolio

¹¹ For better readability of the graphs, we depict the kappa surfaces in Figures 1 - 4 on a grid structure which is much coarser (by a factor of 5 to 25) than the underlying data.

¹² Since we approximate the normal distributions very accurately, there is still some exceedingly small probability of crossing the boundary as long as kappa is not exactly zero. The manager does not entertain negative kappa strategies as these are risky and can thus hit the boundary. Moreover, their expected return is less than the riskfree rate.

becomes progressively smaller. Eventually, the possible upside from a high-kappa bet comes to dominate the alternative of carefully managing the portfolio, as she encounters Gambler's Ridge.

The Hill of Anticipation is a novel area of managerial behavior. It occurs a few percent above the lower boundary and starts some two months before the end. Here, the manager increases the risk of the controlled process substantially but not in the indiscriminate manner of the Gambler's Ridge. She has more to lose and more time left to manage the fund than on the Gambler's Ridge area, and this moderates her behavior regarding kappa. Nevertheless, she finds it attractive to increase kappa above the Merton optimum since the potential loss is limited and the time to maturity is relatively short. If she is fortunate and her higher-kappa bet pays off with a large increase in fund value, she heads toward Merton Flats. There the higher kappa level is too risky and gets revised downward. Hence, the Hill of Anticipation tails off to the right approaching Merton Flats. If she is unfortunate, then there is still consolation (and utility) in the knowledge that she can bet on Gambler's Ridge one last time. The Hill of Anticipation also tails off to the left, dropping into the Valley of Prudence where she prefers to wait until very close to T before undertaking the high-kappa bets associated with Gambler's Ridge.

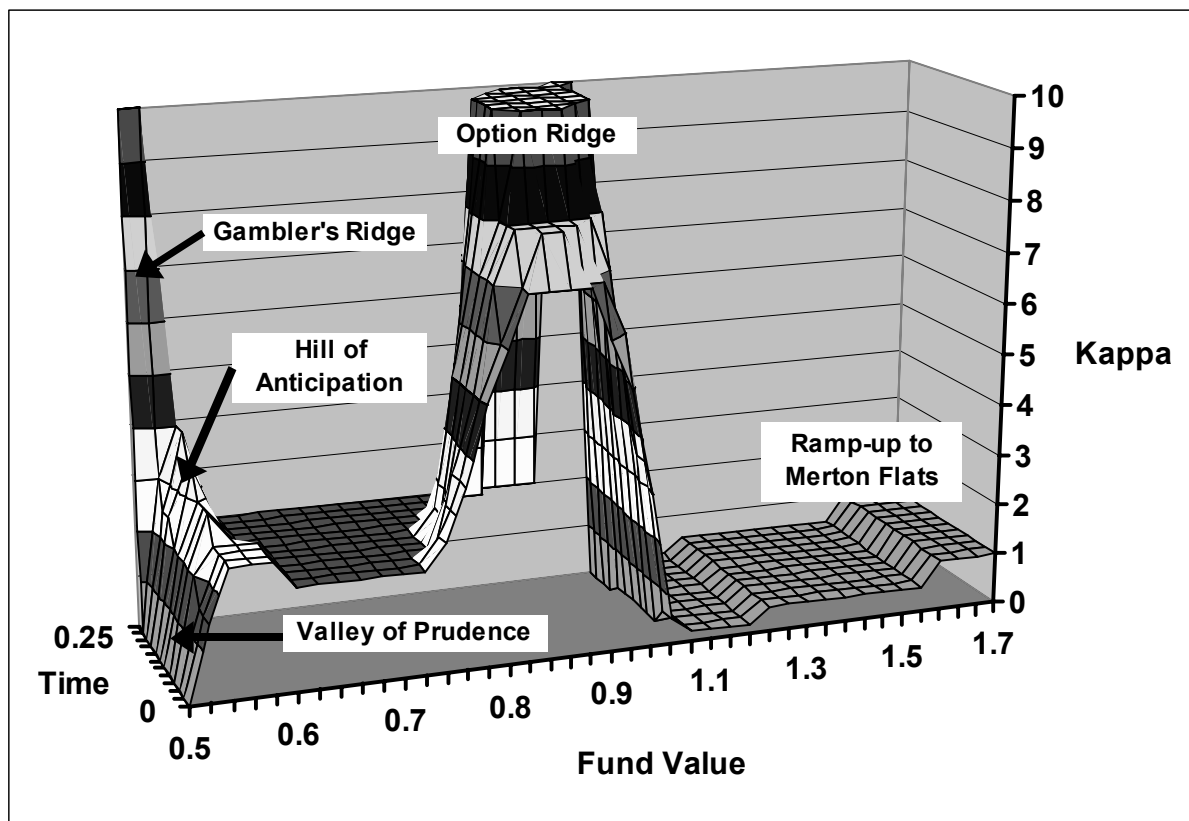
Thus, introducing a liquidation boundary causes the manager to follow an optimal strategy that is much richer than the constant kappa solution. A key factor in these results is the absence of dead-weight liquidation costs or some penalty which reduces the manager's severance compensation. Even a relatively small penalty that reduces her severance compensation by as little as 3% can eliminate her gambling behavior both at the boundary and on the Hill of Anticipation. In that case, we only see the Valley of Prudence along the lower boundary; and that valley extends substantially further before blending into Merton Flats at higher fund values.

B. The Effect of an Incentive Option

We now consider the effect of adding an incentive option (struck at the high-water mark) to the manager's compensation structure and also return to not penalizing her severance compensation. There is still no share ownership by the manager ($a = 0$), but otherwise, the parameters are as in Table 1.

Figure 2. Optimal Risky Investment Proportion (κ) with an Incentive Option but No Managerial Share Ownership

In this figure, the manager receives as compensation a management fee ($b = 2\%$) and an incentive option ($c = 20\%$), but she still does not have an equity stake ($a = 0$). Other parameter values are as specified in Table 1.



In Figure 2, we see the same features as in Figure 1 plus a new region of high kappa values, which we term “Option Ridge”. This region is centered just below the terminal high-water mark of $H_0e^{rT} = 1.0125$. Again, the manager dramatically increases the fund’s riskiness as she approaches the terminal date. Now the motivation is to increase the chance of finishing with her option substantially in-the-money. She thus increases the kappa considerably if the fund value is either somewhat below or slightly above the strike price.


Somewhat above the strike price, Option Ridge drops into a valley where kappa decreases dramatically and can go all the way to zero near maturity as the manager locks in her bonus. If the fund value at maturity were just at the strike price, the incentive option would have a zero payoff. Even a couple of grid points into the money, the option payoff is quite small. Consequently, near maturity and at or slightly above the strike level, the manager has an incentive to choose extremely large kappa values. This incentive tails off rapidly as the fund value increases since the manager starts having more to lose if her option finishes out-of-the-money. This leads to a lock-in style behavior, particularly near maturity and slightly above the money. From the outside investor perspective, such lock-in behavior is another perverse effect of the incentive option. Depending on the level of fund value, that option can induce both dramatically more and dramatically less risk-taking compared with the $\kappa = 2$ preferred by an outside investor with the same utility function as the manager.

There is a Merton Flats region between the Hill of Anticipation and Option Ridge. This is because the liquidation boundary is relatively far below the high-water mark. If the liquidation boundary is sufficiently close to the high-water mark, the incentive option starts to affect the Hill of Anticipation causing it to spread into Option Ridge and eliminating the Merton Flats region in between. There is also another Merton Flats region that is far to the right. To reach that upper Merton Flats, the manager’s incentive option has to be sufficiently deep in the money that it acts

like a fractional share position. Gambler's Ridge and the Valley of Prudence are driven almost exclusively by the lower boundary and therefore do not change noticeably when an incentive option is added to the manager's compensation.

Figure 3. Optimal Risky Investment Proportion (κ) with both an Incentive Option and Managerial Share Ownership

In this figure, the manager receives the complete compensation package: a management fee ($b = 2\%$), an incentive option ($c = 20\%$), and also an equity stake ($a = 10\%$). Other parameter values are as specified in Table 1.



C. The Reference Case

We now reintroduce the manager's share ownership ($a = 10\%$) and examine the effect on her optimal kappa choice in Figure 3. The most dramatic differences between Figure 2 and Figure 3 are that Gambler's Ridge almost disappears and that the Hill of Anticipation vanishes.¹³

In previous figures, Gambler's Ridge and the Hill of Anticipation were induced by partial protection of the basic management fee ($b = 2\%$ annually) when fund value hits the liquidation boundary. However, over a three-month interval, that management fee represents only 0.5% of fund value; and its effects near the lower boundary are largely overwhelmed by the manager's 10% ownership stake (the incentive option being almost worthless near that boundary). Consequently, this part of the picture is consistent with Fung and Hsieh's (1999, p. 316) comment about managerial share ownership inhibiting excessive risk taking. Note that this qualitative result depends importantly on the degree of managerial ownership. Moreover, Option Ridge remains an area of very high kappa values, although somewhat narrower than previously. Above Option Ridge, the manager's optimal kappa does not drop as low as in Figure 2 and also ramps up faster towards an upper Merton Flats region which again exists at high fund values.

Since we now have the manager's optimal kappa at each grid point, we can readily calculate the probability of reaching any terminal grid point given a starting location. This provides another approach for assessing the implications of the manager's risk-taking behavior. For example, starting at the beginning of the grid with the initial fund value X_0 , the manager optimally takes risks which cause the fund return to exhibit a bimodal distribution. Her desire to finish in-the-money with her incentive option, leads her to gamble so much on Option Ridge that she either ends up with large profits (and a sizeable incentive) or much poorer. This is again a striking contrast to what would be preferred by an outside investor with the same utility function.

¹³ They appear compressed in Figure 2 due to the change of horizontal scale relative to Figure 1.

In particular, that investor would prefer a constant kappa strategy, which would generate a lognormal return distribution. Implicitly, that investor is accepting the manager's behavior in order to gain access to the fund's investment technology. However, there would appear to be considerable room for altering the manager's compensation structure to better align her interests with the investor's.

III. Endogenous Shutdown and Effort Choices

In this section, we explore two major extensions of our model. We first consider the possibility that the manager can choose to shut down the fund voluntarily at asset levels above the liquidation boundary. We subsequently examine a variation of our model where the manager can enhance the return characteristics of the fund's (proprietary) technology by expending extra effort.

A. The Managerial Decision to Shutdown the Fund

Instead of simply using a prespecified liquidation boundary, the model can be readily adapted to include a managerial shutdown option. This is an American-style option where the manager can choose to liquidate the fund at asset values above the prespecified lower boundary. Whether she will choose to do so depends on her other opportunities relative to continuing to manage the fund. What largely motivates the manager to keep the fund alive are the possibility of earning the incentive fee by exceeding the high-water mark plus the ability to manage her own

invested capital ($a > 0$) using the fund's superior return technology.¹⁴ If the value of her outside opportunities is large enough to offset those effects, she will choose to shut down the fund.

We model her outside opportunities in a simple manner, using L to represent an annual compensation rate which is independent of the fund value.¹⁵ If the manager chooses to shutdown the fund at time τ at some fund value X_τ above $\Phi_\tau = 0.5 H_0 e^{r\tau}$, she receives at maturity:

$$W_\tau = aXe^{r(T-\tau)} + b(1-a)\tau Xe^{r(T-\tau)} + L(T-\tau) \quad \text{for } 0 \leq \tau \leq T \quad (6)$$

The first two terms of (6) indicate that the manager recovers her share of the fund (aX) plus a prorated fraction of the management fee (with no incentive payment). These two amounts are invested for the time remaining until T at the riskless rate r , since the manager no longer has access to the fund's investment technology after shutdown. She also earns L prorated over the time remaining until T . As we work backward in time through our grid, we compare the indirect utility of receiving (6) with that from choosing the optimal $\kappa(X,t)$ and continuing to manage the fund. When the indirect utility of (6) dominates, it indicates that the manager would voluntarily choose to shut down the fund at that grid point.

In our experience, this endogenous shutdown has only occurred at fund values below the lower edge of Option Ridge, where the probability of reaching the high-water mark becomes very small and essentially disappears as an influence on the manager's decisions. However, depending on the value of L , shutdown can potentially occur well above the prespecified lower boundary.

¹⁴ This is consistent with Brown, Goetzmann, and Ibbotson (1999) who indicate a belief that funds are terminated because it appears unlikely that performance will reach the high-water mark (presumably within a "reasonable" time frame).

¹⁵ More complicated specifications of the manager's outside opportunities are possible; however, the intuition remains the same.

On the other hand, when the value of her outside opportunities is relatively low, the manager will not voluntarily choose to shutdown and must be forced to liquidate the fund at the lower boundary.

Note that if a shutdown occurs, outside investors incur a resetting of their high-water mark when switching to another fund. In effect, they are forced to forgo the possibility of gains in the current fund without triggering incentive fees. Moreover, outside investors can experience a pattern of heavy gambling along Option Ridge with fund closure at perhaps only slightly lower asset values. This could be described as “heads: the manager wins a performance incentive, tails: outside investors have their high-water mark reset.” That description sounds rather unappealing from the perspective of an outside investor but serves to illustrate the importance of being able to address the manager’s optimal actions in an American option framework.

Fung and Hsieh (1997, p. 297) point out the possibility that relatively poor performance may trigger fund outflow which is sufficiently large that “assets shrink so much that it is no longer economical to cover the fund’s fixed overhead and the manager closes it down.”¹⁶ This suggests that the fund’s cost structure as well as the manager’s external opportunities play important roles in her decision whether or not to shut down the fund. We have not explicitly included operating costs, but this can be readily done – at least in simplified form. Variable costs can be modeled via adjusting μ and r to a net-of-cost basis. Fixed costs can be represented as a drag on expected returns that is greater at lower fund values. Both types of costs reduce expected future fund values and the manager’s expected compensation. Hence, they lead to an endogenous shutdown decision at higher fund values than when such costs are not considered.

¹⁶ They also mention the possibility that a young fund with good performance may go unnoticed, the managers get impatient, close down the fund, and return to trading for a financial institution.

B. Managerial Effort

Presumably, outsiders invest in a hedge fund because they believe the manager has an expertise that they cannot replicate for themselves (or that replication is too costly). Previously, we modeled the manager as working with equal effort and skill at all grid points where the fund was in operation. We now consider the possibility that the manager has some control over the effort (and skill) she uses in managing the fund. We model this by assuming that she can enhance μ (the expected return of the risky investment technology) via expending more effort.¹⁷ However, expending effort reduces her utility.

We use ψ to denote the level of effort expended. We use $\psi = 0$ to denote the normal effort level and increase ψ in steps of 0.01 to a maximum of 0.02 (maximum effort level). The enhanced drift for the risky investment technology becomes $\mu + \psi$, and the manager's indirect utility function takes on the modified form of:

$$G_{X,\psi,t} = E[G_{X,\psi,t+\Delta t}] - 0.5g\psi^2 \quad (7)$$

where g is a parameter that scales the manager's aversion to effort.

At each grid point, the manager jointly chooses κ and ψ to maximize her indirect utility (G). We employ the same basic procedure as previously and select the highest indirect utility.¹⁸ We denote that value as the optimal $G_{X,\psi,t}$ as we loop backward through time. We also record

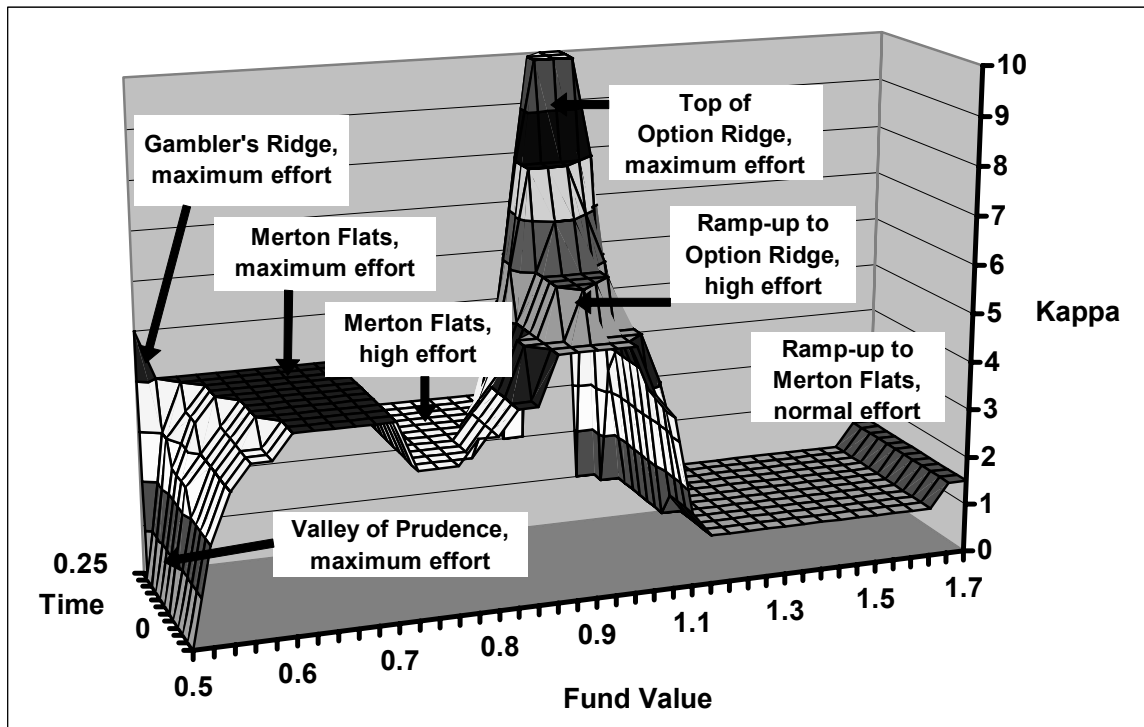
¹⁷ Alternatively, we can model her effort as reducing the volatility (σ) of the risky technology. Altering σ affects both the drift and volatility of the fund value, whereas altering μ affects just the drift. However, the qualitative effects are similar.

¹⁸ Previously, we had a discrete set of kappa values that allowed us to calculate a matrix of probabilities (with one probability vector for each potential kappa value). Now, we change that matrix to have a probability vector using the appropriate drift and volatility for each combination of κ and ψ . Our augmented probability matrix is again the same throughout the grid.

the optimal kappa and psi values for each grid point. For modest numbers of effort levels (we use three – normal, high, and maximum effort), this augmented procedure is not onerous.¹⁹

Figure 4. Optimal Risky Investment Proportion (κ) with both an Incentive Option and Managerial Share Ownership plus a Choice of Three Effort Levels.

In this figure, the manager receives the complete compensation package: a management fee ($b = 2\%$), an incentive option ($c = 20\%$), and also an equity stake ($a = 10\%$). Other parameter values are as specified in Table 1. The manager can also increase the drift by 0, 1, and 2% per year through exerting normal, high, or maximum effort.



We use for our results an effort aversion coefficient of $g = 2500$. Figure 4 displays typical results for the situation where the fund is liquidated at the lower boundary (one-half the high-water mark) as in Section II. We observe that the manager expends only normal effort at relatively high fund values. These are scenarios where she expects a relatively high terminal

¹⁹ We have experimented with up to ten effort levels. This provides more refinement, but the overall qualitative results are much the same as with three effort levels.

payoff and incremental income is less valuable in terms of her utility than at low fund values. Hence, she is less willing to expend additional effort at high fund values. On the other hand, she expends greater effort along the lower boundary, along Option Ridge, and approaching Gambler's Ridge. These also tend to be locations where she chooses high kappa values. As a somewhat loose generalization, she tends to exert maximum effort to get her incentive option into the money and to avoid liquidation.

Compared with Figure 3, the optimal kappa levels are higher below Option Ridge except for the lower portions of the Valley of Prudence -- where the manager is trying to avoid hitting the liquidation boundary by choosing very low kappa values. Kappa values are the same above Option Ridge in both figures. This is consistent with the manager expending only normal effort ($\psi = 0$) in Figure 4, while we have $\psi = 0$ in Figure 3 by construction. Option Ridge is now wider, indicating higher kappa values on the shoulders of that ridge. Intuitively, positive psi values increase the Sharpe Ratio for the risky technology and make greater investment (larger kappa) more attractive. This motivation is very clear in the Merton Flats region below Option Ridge. In Figure 3, the optimal kappa for that region is 2. In Figure 4, the optimal kappa increases to 3 with high effort ($\psi = 0.01$) and to 4 with maximum effort ($\psi = 0.02$). Using equation (5) with μ replaced by $\mu + \psi$, one can readily see that these are the appropriate optimal kappa values conditional on those levels of effort.

We now add the possibility of a voluntary shutdown using the same modeling structure as in the previous subsection. Above the endogenous shutdown level, the optimal risk-taking and exertion of effort is virtually identical to Figure 4. As in the previous subsection, the manager's ability to voluntarily shut down damages outside investors by forcing them to reset their high-water marks at other funds. However, the increased kappa value of 4 in the maximum-effort

portion of Merton Flats causes the manager to choose a slightly lower endogenous shutdown level as compared with that in the previous subsection.

Including effort as a managerial choice variable yields some interesting results; but ones that are intuitively reasonable after some reflection. We see increased effort only on and below Option Ridge. The manager becomes something of a “slacker” when things are going well. Admittedly, the model is simplified; however, this result suggests that the typical hedge-fund incentive structure may not elicit intensive managerial effort at high fund values. It is also interesting that increased effort goes together with higher kappa values rather than resulting in a tradeoff between the two. Thus, one needs to exercise some caution before inferring whether a relatively high kappa is the result of just gambling or in response to enhanced upside probabilities resulting from extra effort by the manager.

IV. Managerial Control and Risk Taking

Recently there has emerged a growing literature examining the nature and effects of incentive compensation mechanisms for money managers. Although using different valuation technologies and somewhat different incentive structures, some of these papers have generated results that can be related to portions of our Figure 3. It is instructive to make those comparisons. It not only promotes a better understanding of how these papers fit together but also strengthens our knowledge of how shares, options, knockout barriers, and horizon times interact in influencing managerial behavior.


Carpenter (2000) utilizes an equivalent martingale technology to determine the optimal trading strategy for a risk averse money manager whose compensation includes an option component. The manager seeks to maximize expected utility of terminal wealth, which is

composed of a constant amount (external wealth and a fixed wage) plus a fractional call option on the assets under management with a strike price equal to a specified benchmark. There are substantial similarities to the incentive option in our model, with Carpenter's benchmark corresponding to our high-water mark at time T . There are also important differences. Carpenter's manager does not have a personal investment in the fund ($a = 0$) and also does not earn a percentage management fee ($b = 0$). These two differences remove the manager's fractional share ownership – see equation (1). Also, Carpenter does not have a knockout barrier where the fund is liquidated or the manager is fired for poor performance.

In Figure 5, we superimpose a graph similar to Carpenter's figure 3 on a stylized time slice from our Figure 3. Carpenter finds results that qualitatively correspond to our manager's behavior when the fund value is above the high-water mark. Starting from the high-water mark, there is the upper slope of Option Ridge followed by a pronounced dip in κ before a gradual ramping up to an upper Merton Flats at high fund values. However, her manager behaves very differently from ours as fund value drops below the strike of the incentive option. Her manager continues to increase volatility as the fund value declines and there is, counterfactually, no limit to this behavior since it is costless to the manager. On the other hand, our manager moderates volatility and gradually reduces the risky investment proportion to the level prevailing in the lower Merton Flats. This difference in behavior is induced by our manager owning a fractional share in the fund which makes it very expensive for a risk averse manager to increase risk without limit. Parenthetically, even if our the manager did not explicitly own a fractional share ($a = 0$), having a percentage fee based on the (terminal) value of funds under management ($b > 0$) generates similar results as in Figure 5.

Figure 5. Comparison of Risk Choices in Different Models I: Hodder & Jackwerth, Merton, and Carpenter

We depict a stylized time slice of the surface of risky investment proportions (κ) from our Figure 3 where the manager receives the standard compensation (management fee $b = 2\%$, incentive option $c = 20\%$, and equity ownership $a = 10\%$). We also graph Merton's optimal solution which is constant at $\kappa = 2$. Finally, we overlay the result from Carpenter (2000) where we assume that her incentive option is aligned with our standard assumptions.



The liquidation boundary and the extent of severance compensation also play important roles in our model whereas Carpenter does not have such a lower boundary. This aspect of the analysis is partially examined in Goetzmann, Ingersoll, and Ross (2003) (GIR). That paper has a fee structure that is similar to ours (except for no explicit managerial ownership) as well as a liquidation boundary. In most of their paper, the hedge fund's investment policy is fixed. However, in section IV they briefly explore a simple extension with the state space (measuring fund value) split into multiple regions, where different volatilities can be chosen by the manager. GIR use an equilibrium pricing approach with a martingale pricing operator based on the attitudes

of a representative investor in the hedge fund. Hence, they cannot directly address choices based on managerial utility. However, they are able to examine volatility choices which maximize the capitalized value of fees (performance plus annual) earned by the fund.

In that context, they examine two alternative cases (GIR, p. 1708). With no lower liquidation boundary, they find that “the volatility in each region should be set as high as possible if the goal is to maximize the present value of future fees.” When they have a liquidation boundary, GIR find that “volatility should be reduced as the asset value drops near the liquidation level to ensure that liquidation does not occur.” They also point out that “this conclusion is inconsistent with that of Carpenter (2000) in which volatility goes to infinity as asset value goes to zero.”

Clearly the liquidation boundary plays a vital role. Carpenter does not have such a boundary (or managerial share ownership). Hence, at low asset values her manager is motivated only by the probability of getting back into the money prior to the evaluation date. The further out-of-the-money and the shorter the time to maturity for her incentive option, the more the manager is willing to gamble. In contrast, GIR have a boundary at which fees go to zero. If the objective is to maximize fees, such a boundary is to be avoided, and this drives their result that volatility should be decreased as asset values approach the boundary. In effect, this is our earlier result where a penalty imposed at the lower boundary causes the manager to reduce κ (and volatility) as the fund value declines near the boundary.

An important but perhaps subtle issue in the GIR model is the timing of performance fees. In GIR, such fees are earned continuously whenever the fund value reaches the high-water mark. In our model as well as Carpenter’s, such fees are earned only on an evaluation date. This difference means that GIR’s manager can never be deep-in-the-money. Similarly, their manager cannot lose an accrued incentive fee by falling out-of-the-money prior to an evaluation date.


Hence, the GIR manager would always want to increase volatility as the fund value moves further away from the liquidation boundary. This serves to emphasize the role of timing in performance measurement. If performance evaluations are quarterly or annual, then the sort of complicated risk-taking behavior seen in Figure 2 and Figure 3 is more realistic than GIR's continuously increasing volatility.

Another related paper is Basak, Pavlova, and Shapiro (2002) (BPS). That paper examines the use of benchmarking to control the risk-taking behavior of a money manager. The manager maximizes expected utility with respect to a terminal payoff function and exercises continuous control of the investment process. One version of their model examines optimal behavior with a single risky plus a riskless asset and generates results which can be fairly readily compared with ours.

Figure 6 qualitatively illustrates the GIR and BPS results compared with ours and with Merton's. As discussed above, GIR's liquidation boundary and incentive structure with continuous earning of performance fees results in volatility being optimally zero at the liquidation boundary and then increasing as the fund value rises. Their paper does not examine this situation graphically, but we illustrate the qualitative result at the left-hand side of Figure 6.

Figure 6. Comparison of Risk Choices in Different Models II: Hodder & Jackwerth, Merton, GIR, and BPS

We depict a stylized time slice through the surface of risky investment proportions (κ) from our Figure 3. There the manager receives the standard compensation (management fee $b = 2\%$, incentive option $c = 20\%$, and equity ownership $a = 10\%$). We graph Merton's optimal solution, which is constant at $\kappa = 2$. Next, we overlay the result from Goetzmann, Ingersoll, and Ross (2003) (GIR) with their lower boundary behavior aligned with our Valley of Prudence. This is a hypothetical graph since GIR do not graph that result in their paper. Finally, we overlay the results from Basak, Pavlova, and Shapiro (2003) (BPS) where we assume their fund flow (digital option) is aligned with our incentive option. Again, we assume that their risk choices for fund values slightly below (0.8 - 0.9) the option strike price align with our own results.



Our illustration of BPS results in Figure 6 is based on their figure 1a.²⁰ We have also aligned their benchmark with our high-water mark, and we are plotting fund value discounted at the riskless rate on the horizontal axis (in that framework, $H_T = 1$). BPS does not have a

²⁰ In their model, the benchmark is risky. An example would be the S&P 500. Consequently, it's possible for their manager to follow a strategy which is either more or less risky than the benchmark. In the current version of our model, the high-water mark is known and it's not possible to follow a less risky strategy than setting kappa to zero (investing completely in the riskless asset). Hence, BPS figure 1b is not relevant in our situation.

liquidation boundary. Consequently, they do not get the types of boundary induced behavior (depending on the severance compensation structure) that occur in our model or in GIR. Instead, the BPS manager optimally pursues a Merton Flats strategy toward the left of Figure 6. This is because their manager’s compensation in that region is effectively a fractional share. As fund value increases toward 1 (our high-water mark), the portfolio weight in the risky security rises²¹ then dives dramatically to zero before rising gradually back to a Merton Flats strategy for high fund values. This behavior around the high-water mark is due to the way BPS model funds flow, which provides an implicit performance incentive for their manager.

In their model, the manager’s compensation is simply proportional to terminal fund value (assets under management in their terminology) and the same as our management fee b . Although they can use other approaches, fund flow is modeled in that paper by adjusting the terminal fund value using a multiplier which takes on just two values $f_L < 1$ for poor performance and $f_H > 1$ when performance is good. Using z to denote the proportionality coefficient between terminal fund value X_T and the manager’s payoff W_T , the BPS compensation structure is equivalent to:

$$W_T = zf_L X_T + z(f_H - f_L) X_T 1_{\{X_T \geq H_T\}} \quad (8)$$

The indicator variable takes on the value one in good performance states, where X_T equals or exceeds what corresponds to our high-water mark. The BPS manager’s compensation as portrayed in equation (8) is effectively a partial share of fund value plus a binary “asset or nothing” call option struck at the high-water mark. This modeling choice implies a rather

²¹ However, the exact shape of this Option Ridge (our terminology) in BPS will depend on the parameter choices and can differ from our model.

extreme response for fund flow compared with empirical estimates by Chevalier and Ellison (1997) which portray fund flow as a much more gradual function of past performance.

Still, there are clear similarities between the BPS compensation structure and our manager's payoff in equation (1) when she does not hit the liquidation boundary prior to date T . In both cases, the manager has a partial share plus an incentive option. However, the binary option in equation (8) has an at-the-money value of $z (f_H - f_L) H_T$. In other words, the incentive structure of equation (8) implies a jump in the manager's compensation (as well as an increased slope) when performance just reaches the benchmark. That jump is what causes the BPS manager's optimal κ in Figure 6 to dive to zero when fund value touches the strike price of 1. In effect, that jump is sufficiently valuable to the manager that she chooses to lock-in the at-the-money position and hold it until date T . At fund values further above the strike price, the BPS manager's risk taking heads back toward a Merton Flats strategy. Effectively, BPS has a payoff structure which is equivalent to Merton (1969) plus a binary option. At asset values which are far enough from that option's strike price, their manager exhibits the same behavior as in Merton.

Comparison of these models highlights the importance of seemingly minor changes in the manager's compensation structure. For example, whether or not the manager has a share position as well as an incentive option can substantially mitigate risk-taking behavior – compare our results and those of BPS with the more extreme risk-taking in Carpenter. The nature of the incentive option (e.g. plain vanilla call versus binary asset-or-nothing) also makes a difference, with the binary option inducing more dramatic shifts in risk-taking because of the jump in value at the strike price. On the other hand, both types of options can cause active managers to lock-in on a high-water mark (or benchmark) months before an evaluation date. Such behavior is presumably undesirable from the perspective of outside investors. We also get the message that liquidation barriers as well as the frequency of evaluation can have dramatic effects. In summary,

there is a lot to be seen in this relatively simple comparison. Our Figure 3 may not depict the “whole elephant,” but it does illustrate how managerial behavior can vary dramatically in different parts of the state space.

V. Concluding Comments

Exploring the effects of a typical hedge-fund compensation contract as well as the implications of differing shutdown alternatives, we find a range of rich and interesting managerial behavior. If fund value is near the lower liquidation boundary and there is only a little time left until the manager’s evaluation date, she may be inclined to take extreme gambles. This behavior is prompted by an asymmetry in payoff structure caused by a liquidation boundary which truncates her down-side compensation risk. She gambles more in this situation the lower her shareholding in the fund and vice versa. Such gambling can also be reduced or eliminated by explicitly penalizing her compensation for hitting the liquidation boundary.

Having a performance incentive for exceeding a high-water mark also induces extensive risk-taking as she tries to push that incentive option into the money. Once that is achieved, she dramatically lowers her risk-taking behavior and pursues a lock-in style strategy. For an outside investor with the same utility function, this behavior is far different from what he would prefer. Indeed, that behavior results in a bimodal distribution for the managed hedge fund returns. It is not clear why this highly nonlinear compensation contract would be used. Seemingly, a linear contract would provide a much better alignment of the manager’s risk-taking and the preferences of external investors. Bebchuk and Fried (2003) suggest an interesting approach by viewing the compensation contract demanded by a powerful manager as part of the agency problem rather than its solution. We intend to explore this question in future research.

Such a bimodal return distribution, particularly when coupled with an endogenous shutdown option, suggests a potentially serious problem with survivorship bias in reported hedge fund returns. Liang (2003) documents such an empirical survivorship bias. The bimodal distribution will also cause potentially large errors in derivative prices based on the erroneous assumption that asset values follow an uncontrolled geometric Brownian motion. For example, a European call struck at the high-water mark is more than twice as valuable with managerial control (using our standard parameters) than its Black-Scholes price with the constant volatility preferred by the outside investor. This option corresponds to the manager's incentive fee structure. However, there are also instances where options on hedge fund values have been part of external transactions. A well-cited example is the seven-year call on \$800 million of Long Term Capital Management (LTCM) shares sold by Union Bank of Switzerland (UBS) to LTCM. Again, these are issues we intend to explore in the future.

We frequently find that seemingly slight adjustments in the compensation structure have dramatic effects on managerial risk-taking. In addition to our comparisons in Section II, this was again illustrated in Section IV (see Figure 5 and Figure 6), where we examine results from recent papers by Carpenter (2000), Goetzmann, Ingersoll, and Ross (2003), and Basak, Pavlova, and Shapiro (2003). Although we can explain results from those papers using our model and put them into a more general context, the dramatic divergence of results across those papers illustrates that one needs to be cautious with generalities about managerial behavior. Even minor additions to the model can have major implications.

Allowing the manager to voluntarily shut down the fund adds an American-style option to the analysis. Our methodology can readily handle this situation, and it adds an interesting aspect of managerial discretion. Two key drivers in the shutdown decision appear to be the manager's outside opportunities and the likelihood that her performance incentive option will finish out-of-

the-money. Moreover, it is possible that the manager chooses to shutdown at a fund value well above what outside investors would prefer.

Allowing the manager to enhance the fund's Sharpe Ratio via increased effort adds another interesting dimension to the analysis. One striking result is that maximum effort tends to go together with relatively high risk-taking. This makes sense because the effort enhances the Sharpe Ratio, which makes greater risk-taking more attractive. A second potentially important result is that standard hedge fund compensation contracts do not appear to provide much incentive for maximum effort levels when the fund is doing well. Our modeling of effort choice is simplified, but such issues clearly warrant further attention.

Managerial control of the hedge fund's investments implies controlling the stochastic process for fund value. An underlying theme of the paper is developing a methodology for valuing payoffs (derivatives) based on such a controlled process. The basic approach we developed here can be applied to other situations where a portfolio return process is controlled by a utility maximizing individual. With some added constraints, a mutual fund manager clearly fits this description, as does a currency trader at a bank.

In a more approximate manner, we can think of a firm being controlled by an individual manager (the CEO). A useful comparison is Merton (1974), where risky debt is valued based on an exogenous underlying process for the firm's asset value. An alternative perspective is to model this asset value process as being controlled via investment and hedging decisions, in a manner analogous to an investment portfolio. From that perspective, not only risky debt but any derivative based on firm value (such as stock or options) is implicitly based on a controlled process. Hence, the basic valuation technology developed in this paper has numerous potential applications.

Appendix: Numerical Procedure

The basic structure of our model uses a grid of fund values X and time t , with $\Delta(\log X)$ constant as well as time steps Δt of equal length. The initial fund value X_0 is on the grid, and it is convenient to have the fund values increase over each time step at the riskfree rate $e^{r\Delta t}$. This choice implies that in the limiting case where $\kappa = 0$ (the manager chooses to only invest in the riskless asset) the value process will still reach a regular grid point. Thus, the grid structure will not prevent the manager from switching to the riskless strategy. Maintaining this structure for the lower boundary implies having $\Phi_t = \Phi_0 e^{rt}$ where t is a multiple of Δt and $0 \leq t \leq T$.

To calculate expected utilities, we will need the probabilities of moving from one fund value at time t to all possible fund values that can be reached at $t+\Delta t$. The possible $\log X$ moves are $r\Delta t + i\Delta(\log X)$ where the $r\Delta t$ term is due to the riskless drift in the X grid. We use i to index the grid points to which we can move. In the current implementation, the range for i is from $-60, \dots, 0, \dots, 60$. The probabilities for those possible moves depend on the choice of kappa which determines the process for X over the next time step. For a given kappa, the log change in X is normally distributed with mean $\mu_{\kappa, \Delta t} = [\kappa\mu + (1-\kappa)r - \frac{1}{2}\kappa^2\sigma^2]\Delta t$ and volatility $\sigma_{\kappa, \Delta t} = \kappa\sigma\sqrt{\Delta t}$. Note that this mean and variance do not depend on the level of X . They do depend on Δt but not on t itself. Since the normal distribution is characterized by its mean and variance, the probabilities we need are solely functions of κ and not the grid point.

We now use the discrete normal distribution. For a given kappa, we calculate the probabilities based on the normal density times a normalization constant so that the computed probabilities sum to one:

$$P_{i,\kappa,\Delta} = \frac{\frac{1}{\sqrt{2\pi\sigma}} \text{EXP} \left[-\frac{1}{2} \left(\frac{r\Delta t + i\Delta(\log X) - \mu_{\kappa,\Delta t}}{\sigma_{\kappa,\Delta t}} \right)^2 \right]}{\sum_{j=-60}^{60} \frac{1}{\sqrt{2\pi\sigma}} \text{EXP} \left[-\frac{1}{2} \left(\frac{r\Delta t + j\Delta(\log X) - \mu_{\kappa,\Delta t}}{\sigma_{\kappa,\Delta t}} \right)^2 \right]} \quad (\text{A1})$$

We keep a lookup table of the probabilities for different choices of kappa which we vary from 0, 0.1, 0.2, 0.5, 1.0, 1.5, 2.0, 2.5, 3, 3.5, 4, 5, 6, 7, 8, 10, to 20. However, the ends of this range are problematic and can result in poor approximations to the normal distribution. For low kappa values, the approximation suffers from not having fine enough value steps. For high kappa values, the difficulty arises from potentially not having enough offset range to accommodate the extreme tails of the distribution.

To insure reasonable accuracy, we compare the standardized moments of our approximated normal distribution $\hat{\mu}_j$ with the theoretical moments of the standard normal, $\mu_j = 1 \cdot 3 \cdot \dots \cdot (j-1)$ for j even and $\mu_j = 0$ for j odd. In particular, we calculate a test statistic based on the differences of the first 10 approximated and theoretical moments scaled by the asymptotic variance of the moment estimation – see Stuart and Ord (1987, p. 322):

$$\frac{1}{10} \sum_{j=1}^{10} \left(\frac{\hat{\mu}_j - \mu_j}{\frac{1}{n}(\mu_{2j} - \mu_j^2 + j^2 \mu_2 \mu_{j-1}^2 - 2j \mu_{j-1} \mu_{j+1})} \right)^2, \text{ where we set } n = 1 \quad (\text{A2})$$

After some experimentation, we discard distributions with a test statistic of more than 0.01. For our standard model, this results in eliminating the distributions associated with the kappa level of 0.1 and the kappa levels greater than 10. We finally have a matrix of probabilities with a probability vector for each kappa value in our remaining choice set.

We now calculate the expected indirect utilities and initialize the indirect utilities at the terminal date J_T to the utility of wealth of our manager $U_T(W_T)$ where her wealth is solely determined by her compensation scheme. Our next task is to calculate the indirect utility function at earlier time steps as an expectation of future indirect utility levels. We commence stepping backwards in time from the terminal date T in steps of Δt . At each fund value within a time step t , we calculate the expected indirect utilities for all kappa levels using the stored probabilities and record the highest value as our optimal indirect utility, $J_{X,t}$. We continue, looping backward in time through all time steps.

In our situation, using a lookup table for the probabilities associated with the kappas has two advantages compared with using an optimization routine to find the optimal kappa. For one, lookups are faster although coarser than optimizations. Second, a sufficiently fine lookup table is a global optimization method that will find the true maximum even for non-concave indirect utility functions. In such situations, a local optimization routine can get stuck at a local maximum and gradient-based methods might face difficulties due to discontinuous derivatives.

When implementing our backward sweep through the grid, we have to deal with behavior at the boundaries. The terminal step is trivial in that we calculate the terminal utility from the terminal wealth. The lower boundary is also quite straightforward. We stop the process upon reaching or crossing the boundary and calculate the utility associated with hitting the boundary at that time. For our basic model, the manager's severance pay is reinvested at the riskfree rate until time T . Consequently, she receives a terminal wealth of $W_T = aX_\tau e^{r(T-\tau)} + 0.5(1-a)b\tau H_0 e^{rT}$ for sure. Because that terminal payoff is certain, its expected utility is simply the utility of terminal wealth for that payoff. We use these values in calculating the expected indirect utility at earlier time steps.

For the numerical implementation, we also need an upper boundary to approximate indirect utilities associated with high fund values. We use a boundary 600 steps above the initial X_0 level. For fund values near that boundary, our calculation of the expected indirect utility will try to use indirect utilities associated with fund values above the boundary. We deal with this by keeping a buffer of fund values above the boundary so that the expected indirect utility can be calculated by looking up values from such points. We set the terminal buffer values simply to the utility for the wealth level associated with those fund values. We then step back in time and use as our indirect utility the utility of the following date times a multiplier which is based on the optimal Merton (1969) solution without consumption: $\exp[\Delta t(\mu - r)^2(1 - \gamma)/(2\gamma\sigma^2)]$. We do not assume that these values are correct (they are based on a continuous time model while we work in a discrete time setting) but they work very well. This approach is potentially suboptimal, which biases the results low. However, the distortion ripples only some 20-50 steps below the upper boundary, affecting mainly the early time steps.

References

- Basak, Suleyman, Anna Pavlova, and Alex Shapiro (2003), "Offsetting the Incentives: Risk Shifting and Benefits of Benchmarking in Money Management," working paper, MIT.
- Basak, Suleyman, Alex Shapiro, and Lucie Teplá (2003), "Risk Management with Benchmarking," working paper, London Business School, December.
- Bebchuk, Lucian A. and Jesse M. Fried (2003), "Executive Compensation as an Agency Problem," *Journal of Economic Perspectives* 17, 71-92.
- Brown, Stephen J., William N. Goetzmann, and Roger G. Ibbotson (1999), "Offshore Hedge Funds: Survival and Performance, 1989-95," *Journal of Business* 72, 99-117.
- Carpenter, Jennifer N. (2000), "Does Option Compensation Increase Managerial Risk Appetite?" *Journal of Finance* 55, 2311-2331.
- Chevalier, Judith and Glenn Ellison (1997), "Risk Taking by Mutual Funds as a Response to Incentives," *Journal of Political Economy* 105, 1167-1200.
- Fung, William and David A. Hsieh (1997), "Empirical Characteristics of Dynamic Trading Strategies: The Case of Hedge Funds," *Review of Financial Studies* 10, 275-302.
- Fung, William and David A. Hsieh (1999), "A Primer on Hedge Funds," *Journal of Empirical Finance* 6, 309-331.
- Goetzmann, William N., Jonathan E. Ingersoll, Jr., and Stephen A. Ross (2003), "High-Water Marks and Hedge Fund Management Contracts," *Journal of Finance* 58, 1685-1717.
- Liang, Bing (2003), "On the Performance of Alternative Investments: CTAs, Hedge Funds, and Funds-of-Funds," working paper, Case Western Reserve University, April.
- Markowitz, Harry (1959), *Portfolio Selection: Efficient Diversification of Investments*, Cowles Foundation Monograph #16 (Wiley 1959); (reprinted by Blackwell 1991).
- Merton, Robert (1969), "Lifetime Portfolio Selection under Uncertainty: The Continuous Time Case," *Review of Economics and Statistics* 51, 247-257.
- Merton, Robert (1974), "On the Pricing of Corporate Debt: The Risk Structure of Interest Rates," *Journal of Finance* 11, 449-470.
- Mossin, Jan (1968), "Optimal Multiperiod Portfolio Policies," *Journal of Business* 41, 215-229.
- Ross, Stephen A. (2004), "Compensation, Incentives, and the Duality of Risk Aversion and Riskiness," *Journal of Finance* 59, 207-225.
- Stuart, A., and S. Ord (1987), *Kendall's Advanced Theory of Statistics*, Vol. 1, 5th ed. Oxford University Press, New York.