The Hedge Fund Game:

Incentives, Excess Returns, and Performance Mimics

Dean P. Foster and H. Peyton Young

November 2007

This version: 24 September 2008

Acknowledgements: We are indebted to Andrew Lo, Tom Norman, Andrew Patton, Joe Perkins, Tarun Ramadorai, Krishna Ramaswamy, Neil Shephard, and Robert Stine for helpful suggestions.

Abstract

We show that it is extremely difficult to devise incentive schemes that distinguish between fund managers who cannot deliver excess returns from those who can, unless investors have specific knowledge of the investment strategies being employed. Using a 'performance-mimicking' argument, we show that any fee structure that does not assess penalties for underperformance can be *gamed* by unskilled managers to generate fees that are at least as high, per dollar of expected returns, as the fees of the most skilled managers. We show further that standard proposals to reform the fee structure, such as imposing high water marks, delaying managers' bonus payments, forcing them to hold an equity stake, or assessing penalties for underperformance, are not enough to separate the skilled from the unskilled. We conclude that skilled managers will have to find ways other than their track records to distinguish themselves from the unskilled, or else the latter may drive out the former as in a classic lemons market.

1. Background

Hedge funds are largely unregulated investment vehicles that have become increasingly important in global financial markets.[1] Currently there are nearly ten thousand funds that collectively have over two trillion dollars under management. Although hedge funds pursue a great variety of investment strategies, they have two key features that we shall focus on here. One is the fee structure: the great majority of funds have a two-part structure consisting of a *management fee* plus a *performance bonus* that gives the manager a percentage of any excess returns he generates over and above some benchmark rate. Management fees are typically between 1% and 2% and the most common bonus is 20% (Ackermann, MacEnally, and Ravenscraft, 1999). A second key characteristic shared by many (though not all) hedge funds is lack of transparency: they need not, and often do not, disclose their positions or trading strategies to investors; all they are required to provide is regular audited statements of gains and losses.

These two features – fees based on excess returns and lack of information about how the returns were generated – create incentives for manipulation, as a number of authors have pointed out (Starks, 1987; Carpenter, 2000; Ackermann, MacEnally, and Ravenscraft, 1999; Lo, 2001; Hodder and Jackwerth, 2007). One problem is that the convexity of the fee structure encourages managers to employ strategies with high variance, especially as the date for meeting certain targets is approached. A second problem is that substandard returns in later periods do

[1] The term 'hedge fund' covers a wide variety of funds that are not subject to the Investment Company Act of 1940. For analyses of hedge fund performance see Fung and Hsieh (1997, 1999, 2000), Brown, Goetzmann, and Ibbotson (1999), Ackermann, MacEnally, and Ravenscraft (1999), and Ibbotson and Chen (2006).

not offset the earnings from excess returns in earlier periods unless the contract contains clawback provisions, which are fairly unusual. Furthermore it is quite easy to 'game' standard measures of performance, such as the Sharpe ratio, Jensen's alpha, and the appraisal ratio (Goetzmann, Ingersoll, Spiegel, and Welch, 2007; Guasoni, Huberman, and Wang, 2007).

However, although it is well-known that fees and performance measures can be *gamed to some extent*, there has to our knowledge been no systematic analysis of *the extent to which they can be gamed*. In this paper we propose a framework for analyzing this issue and explore its implications for the incentive structure. The basic idea is to show how a manager with no skill can manufacture apparently excess returns (*fake alpha*) for long periods of time without generating excess returns in expectation. Lo (2001) was one of the first to describe schemes of this sort, which are sometimes known as *Lo strategies*. The approach we take here is a modification that shows how a manager with no investment skill can *mimic* the performance of any skilled manager he chooses for an extended period of time. This construction allows us to make precise analytical statements about the *extent* to which fake alpha can be manufactured, and how much money it will earn under different fee structures.

The approach is a variant of a game-theoretic concept introduced by David Gale called *strategy-stealing* (Gale, 1974). This is a tool for analyzing equilibria in games that are so complex that the explicit construction of equilibrium strategies is difficult or impossible; nevertheless it is sometimes possible to compare the players' payoffs in equilibrium without knowing what the equilibrium is. The general idea runs as follows: suppose that one player in a game (say *i*) has a certain strategy s_i that results in payoff α_i. Then another player (say *j*) can copy *i*'s strategy and get a payoff at least as high as α_i. Gale originally applied this idea to a board game called Chomp, which is similar to Nim. In particular, he showed that the first mover must have a winning strategy even though he (Gale) could not construct it: for if the second mover had a winning strategy, the first mover could 'steal' it and win instead.

We shall show that a version of this argument holds in financial markets with options trading. *Namely, a trader with no skill can mimic the returns being generated by another (more skilled) trader for an extended period of time without knowing how the skilled trader is actually producing these returns*. We show how to compute the probability with which the unskilled trader can mimic the skilled one for a given period of time, and show how much he will earn in the process. A key feature of the argument is that the unskilled trader need not know anything about the actual investment strategy being employed by the skilled trader. Nevertheless, the unskilled trader has an options-trading strategy that *exactly mimics* the return sequence being generated by the skilled trader with high probability (and blows up with small probability). He therefore earns fees and attracts customers just as if he were skilled until the fund blows up. Thus, instead of stealing the other's strategy, he mimics the other's performance.[2]

We use this approach to study the question of whether the fee structure can be redesigned to screen out unskilled managers. Standard proposals for reducing the manipulability of fees include: imposing high water marks, delaying

---

[2] This is not the same as *cloning* or *replication strategies*, which attempt to reproduce the statistical properties of a given fund or class of funds. See for example Kat and Palaro (2005) and Hasanhodzic and Lo (2007).

managers' bonus payments, forcing them to hold an equity stake, and assessing penalties for poor performance.[3] It turns out that none of these proposals is sufficient to separate the skilled from the unskilled.

2. Piggybacking

A standard way to manufacture fake alpha is to write (or short) deeply out-of-the-money options at the start of the reporting period and hope that they do not end up in-the-money by the end of the period. In effect, the manager sells an insurance policy on events that have only a small chance of occurring, and the resulting premiums jack up the reported returns (while exposing investors to large potential losses if the events do occur). Lo (2001) was the first to give a specific example of this approach, to show how it can be concealed using synthetic options positions, and also to point out that such a strategy can make a great deal of money before crashing.[4]

In this section we present another way of implementing this idea that is particularly flexible and computationally tractable. Specifically, we shall show how a manager can 'piggyback' any target series of excess returns on top of the returns delivered by a benchmark asset such as bonds or stocks, and we shall

---

[3] The approach can also be used to 'game' standard performance measures but we shall not pursue that application here. See Goetzmann, Ingersoll, Spiegel, and Welch (2007) for a general discussion of how to manipulate such measures.

[4] In Lo's example the manager takes a series of short positions in S&P 500 put options that mature in 1-3 months and are about 7% out-of-the-money. He then simulates the returns from this strategy using historical data from the 1990s. Agarwal and Naik (2004) show that many equity-oriented funds have payoffs that resemble short positions in out-of-the-money puts on the S&P 500, though of course this does not prove that managers are consciously pursuing such strategies.

show how to compute the probability with which he can do this over any given period of time. This *piggyback strategy* is quite transparent and might not fool an auditor if it were discovered, though it can easily be 'dressed up' to conceal its true nature. The point, however, is not to develop the most realistic way of deceiving auditors (or investors), but to demonstrate a class of strategies that allows one to compute a lower bound on how much money can be made and the probability with which the manager can get away with it before the fund crashes.

We also want to emphasize that it is not particularly important whether a manager is deliberately trying to deceive his investors, or simply imagines that he can outsmart the market and is therefore deceiving himself. The former is a *con artist* while the latter is an *arbitrageur* with a run of good luck but no particular skill. Both types of managers want to maximize their fees, which means they are motivated to pursue high-risk strategies. From the investors' point of view the outcome is essentially the same irrespective of the managers' intentions.

Before defining a piggyback strategy precisely we need some notation. Consider first a safe asset, such as a government bond with a risk-free rate of return $r > 0$. Suppose that a fund starts with size 1 at the beginning of year 1 and runs for T years. The total return in year *t* will be denoted by $(1+r)X_t$ where $X_t = 1 + \alpha_t \geq 0$ is a multiplicative random variable generated by the fund manager's investment strategy. There are *excess returns* in year *t* if $X_t > 1$, *deficient returns* if $X_t < 1$, and *ordinary returns* if $X_t = 1$. At the end of T years the fund's *gross return* (before fees) is $R_T = (1+r)^T \prod_{1 \leq t \leq T} X_t$ and the *compound excess return* is $E_T = \prod_{1 \leq t \leq T} X_t - 1$.

This set-up can be generalized to situations where the benchmark asset delivers stochastic rather than deterministic returns. Let $\vec{Y} = (Y_1, ...Y_T)$ be a stochastic sequence of returns generated by a tradable asset, such as a weighted portfolio of stocks and bonds. These returns are given exogenously and cannot be manipulated by the manager. As before, we denote the manager's contribution to the total return in period *t* by the multiplicative random variable $X_t = 1 + \alpha_t \geq 0$. Thus α_t is the 'alpha' generated by the manager in period *t* and the total return over the period is $X_t Y_t$. After *T* years the total return is $R_T = \prod_{1 \leq t \leq T} X_t Y_t$ and the excess return is $E_T = \prod_{1 \leq t \leq T} X_t - 1$. The sequence of random variables $\vec{X} = (X_1, ...X_T)$ is generated by some investment strategy, but since we are only interested in the returns there is no real loss of generality in saying that $\vec{X}$ *is* the manager's strategy.

A specific realization of the stochastic processes $\vec{X}$ and $\vec{Y}$ over *T* periods will be denoted by $(\vec{x}, \vec{y})$, where $\vec{x} = (x_1, x_2, ..., x_T)$ and $\vec{y} = (y_1, y_2, ..., y_T)$. Any such realization generates a series of fees for the manager. We assume, as is usually the case, that the fees can be expressed as a percentage of funds under management, where the percentages depend on the sequence of realized returns up to the present. Specifically, a *fee contract over T years* is a vector-valued function ϕ that maps each realization $(\vec{x}, \vec{y})$ to a series of payments $\phi(\vec{x}, \vec{y}) = (\phi_1(\vec{x}, \vec{y}), ..., \phi_T(\vec{x}, \vec{y}))$, where $\phi_t(\vec{x}, \vec{y})$ is the *fee per dollar in the fund at the start of period t*, $0 \leq \phi_t(\vec{x}, \vec{y}) \leq x_t y_t$, and $\phi_t(\vec{x}, \vec{y})$ depends only on the realizations $(x_1, ...x_t; y_1, ..., y_t)$ up through period t. The return to investors *net of fees* in period *t* is $x_t y_t - \phi_t(\vec{x}, \vec{y})$ and the *compound net return* to the investors over the first *t* periods is

$$R_t(\vec{x},\vec{y}) = \prod_{1\le s\le t}[x_s y_s - \phi_s(\vec{x},\vec{y})]. \tag{1}$$

Consider, for example, a 'two and twenty' contract in which the two percent management fee is collected at the end of the period and the twenty percent bonus is collected on the excess returns α_t. Then the fees per dollar invested at the start of period *t* can be written

$$\phi_t(\vec{x},\vec{y}) = .02x_t y_t + .2[x_t - 1]_+ y_t \ , \tag{2}$$

where $[\cdot]_+$ denotes 'nonnegative part of'. If the bonus is collected on the excess return above some pre-specified rate $r > 0$, rather than on the excess above y_t, then the formula becomes

$$\phi_t(\vec{x},\vec{y}) = .02x_t y_t + .2[x_t y_t - 1 - r]_+. \tag{3}$$

This framework is very general and allows us to analyze many variations in the fee structure. For example we can build in high water marks as follows: let $m_t(\vec{x},\vec{y}) = \max\{x_1y_1, x_1y_1x_2y_2, ..., \prod_{1\le s\le t} x_s y_s\}$ be the maximum compound return that has ever been realized up to and including period *t*. Suppose that we impose high water marks on the version of the two and twenty fee structure defined in (3). To keep the notation simple let us assume that the benchmark rate is $r = 0$. Then the formula for the fees becomes

$$\phi_t(\vec{x},\vec{y}) = .02x_t y_t + .2[(m_t / m_{t-1}) - 1]_+. \tag{4}$$

We shall say that a manager has *no ability* or is *unskilled* if all of her feasible strategies $\vec{X} = (X_1, ... X_T)$ satisfy $E[X_t \mid x_1, ... x_{t-1}; y_1, ..., y_{t-1}] \leq 1$ for every t and all realizations $x_1, ... x_{t-1}; y_1, ..., y_{t-1}$. A manager has *ability level* $\alpha > 0$ if she has a strategy that *consistently generates excess returns* α in every period, that is, if $X_t = 1 + \alpha$ for all t and all prior realizations $(x_1, ... x_{t-1}; y_1, ..., y_{t-1})$. This is, of course, considerably stronger than saying she generates excess returns of size α in expectation, but it turns out to be an analytically useful definition as we shall see later on.

Given a fee contract ϕ, an excess returns strategy $\{X_t\}$, and a benchmark asset generating returns $\{Y_t\}$, define the manager's *take* in period t to be his expected fee divided by the expected returns he delivers during the period:[5]

$$\tau_t(\phi, \vec{X}, \vec{Y}) = \frac{E[\phi_t(X_1, ... X_t; Y_1, ..., Y_t)]}{E[X_t Y_t]}. \tag{5}$$

We shall assume that the returns and investment strategies are such that the expected return in each period is strictly positive, and thus the denominator of (5) does not vanish.

The benchmark asset $\vec{Y}$ is *safe* if it takes a constant value $Y_t = c_t$ in every period *t*; otherwise it is *stochastic*.

[5] An alternative definition would be $E[\phi_t(\vec{X}, \vec{Y}) / X_t Y_t]$, that is, the *expected ratio* rather than the ratio of the expected values. The problem with this alternative is that the manager's strategy may sometimes produce zero gross returns (the fund is cleaned out), in which case the expected ratio is undefined.

Theorem. *Let* $\vec{Y}$ *represent the returns over T periods of a benchmark asset, safe or stochastic, let* $\vec{X}$ *represent the returns generated by the manager's strategy, and let* $\phi(\vec{x},\vec{y})$ *be a non-negative contract over T periods.*

i) A mimic has a strategy that generates any desired sequence of excess returns $(x_1,...,x_T) \geq (1,...,1)$ *with probability at least* $1/\prod_{1\leq t\leq T} x_t$ *for all realizations* $\vec{y}$.

ii) For all realizations $\vec{y}$*, the mimic's take of a fund invested in the benchmark asset is at least as high as the take of any skilled manager he chooses to mimic.*

Before proving this result, note that it is a statement about *how much an unskilled manager can earn relative to a skilled one*; it does not identify how much they earn in equilibrium, nor does it specify their utility for different earnings streams. We shall have more to say about these issues in the next section. Suffice it to say here that the mimicking argument is useful precisely in situations like these where there are many players with diverse utility functions and the equilibrium is hard to pin down.

Proof. Without loss of generality assume that the initial size of the fund is $w_0 = 1$. We will first treat the case where the benchmark asset is safe and generates a risk-free rate of return $r > 0$ in every period, that is, $y_t = 1 + r$ for all *t*. Consider a manager with no ability who wants to mimic a target sequence of excess returns $(x_1,...,x_T) \geq (1,...,1)$. First he chooses a stochastic asset, such as the S&P 500, on which binary *cash-or-nothing* options are traded. Assume that the price

P_t of the benchmark asset follows a continuous-time geometric Brownian motion of form

$$dP_t = \mu P_t dt + \sigma P_t dW_t . \qquad (6)$$

We shall assume that the reporting periods are defined in discrete time by the values $t = 1, 2, 3\ldots$. Thus the total (geometric) return in the t^{th} period is $Y_t = P_t / P_{t-1}$. Let the continuous-time risk-free rate be $\tilde{r}$, so that $r = e^{\tilde{r}} - 1$, and assume that $\tilde{r} < \mu$.

Consider a given period *t*. The mimic wishes to inflate the returns during this period by the factor $x_t \geq 1$. To this end he shorts a certain quantity *q* of *cash-or-nothing* puts that expire before the end of the period. Assume that each option pays out one dollar if exercised; otherwise it pays nothing. Let Δ be the time to expiration and let *s* be the strike price divided by the current price. Without loss of generality we may assume that the current price is 1. Then the option's present value is $e^{-\tilde{r}\Delta} v$ where

$$v = \Phi[(\ln s - \tilde{r}\Delta + \sigma^2 \Delta / 2) / \sigma\sqrt{\Delta}].[6] \qquad (7)$$

The probability that the put is exercised equals

$$p = \Phi[(\ln s - \mu\Delta + \sigma^2 \Delta / 2) / \sigma\sqrt{\Delta}]. \qquad (8)$$

By selling q options he collects vq dollars now. By investing them all in the asset until expiration he can cover q options where $w + vq = q$. Thus $q = w/(1-v)$ and

[6] See for example Hull (2009, section 24.7).

the total number of shares in the fund is now $w/(1-v)$. He chooses the expiration date so that it occurs before the end of period *t*, and the strike price *s* so that v satisfies $v = 1 - 1/x_t$.[7] Thus after shorting the options and before they expire he will have $x_t w$ shares; after expiration he will have $x_t w$ shares with probability $1-p$, and the fund will be cleaned out with probability p.

It remains to show that the options expire with probability *at most* v, in which case the fund survives until the next period with probability at least $1/x_t = 1-v$. This follows immediately from (7) and (8) and the fact that $\tilde{r} < \mu$. Therefore if he had $w_{t-1} > 0$ shares at the end of period $t-1$, by the end of period t he will have $x_t y_t w_{t-1}$ with probability *at least* $1/x_t$ and zero with probability *at most* $1-1/x_t$. Therefore after T periods, he will have generated the target sequence of excess returns $(x_1, x_2, ..., x_T)$ with probability at least $1/\prod_{1 \le t \le T} x_t$, as claimed in part i) of the theorem.

Let $\phi_t(\vec{x}) \ge 0$ be the payment in period *t*, which does not depend on $\vec{y}$ because the returns are constant. Since the contract carries no penalties if the fund crashes, the mimic's *expected fee* is at least $\phi_t(\vec{x})/x_t$. Since he has no ability, his *expected total return* during the period is $y_t = 1 + r$. Hence the mimic's *take* in period *t* is at least $\phi_t(\vec{x})/x_t y_t$. But this is the same as the take of a manager who delivers x_t with certainty.

[7] Note that the strike price need not be far out of the money if the time to expiration is short, hence the Black-Scholes formula can be assumed to hold with a high degree of accuracy.

Now consider a skilled manager who generates a *distribution of excess returns* $X_t \geq 1$ in every period $1 \leq t \leq T$. We have just shown that for any realization $\vec{x}$, the mimic's take is at least that of the skilled manager in every period. Hence this also holds in expectation. This establishes claims i) and ii) of the theorem when the benchmark asset is riskless.

Next let us consider the case of a stochastic benchmark asset that generates the return sequence Y_t. In this case the mimic shorts a number of *asset-or-nothing* options that pay out one share of the asset if the strike price is exceeded. (Asset-or-nothing options are just a combination of plain vanilla options and cash-or-nothing options.) As before let Δ be the time to expiration and $\tilde{r}$ the continuous-time risk-free rate. For simplicity we shall assume there no dividend. Let s be the strike price divided by the current price. The present value of one asset-or-nothing put is (see Hull, 2009, section 24.7)

$$v = \Phi[(\ln s - \tilde{r}\Delta - \sigma^2\Delta/2)/\sigma\sqrt{\Delta}]. \tag{9}$$

The probability that it is exercised is

$$p = \Phi[(\ln s - \mu\Delta + \sigma^2\Delta/2)/\sigma\sqrt{\Delta}]. \tag{10}$$

As before we need to have $p \leq v$, which will be the case if $\mu - \tilde{r} \geq \sigma^2$. If $\mu - \tilde{r} < \sigma^2$, the mimic shorts asset-or-nothing calls instead of asset-or-nothing puts. The argument now proceeds as before. By the end of the period he has inflated the number of shares by the factor x_t with probability at least $1/x_t$. Hence he realizes a return $x_t Y_t$ by the end of the period with probability at least

$1/x_t$ irrespective of the realization of Y_t. This establishes part i) of the theorem when the benchmark asset is stochastic.

To prove part ii), let $\phi_t(\vec{x},\vec{y}) \geq 0$ be the payment in period *t*. Conditional on not having crashed before the start of period *t*, the mimic's expected fee in period *t* is *at least* $\phi_t(\vec{x},\vec{y})/x_t$, while the expected total return generated during the period is y_t. Hence the mimic's *take* is at least $\phi_t(\vec{x},\vec{y})/x_t y_t$. This is exactly the same as the take of a skilled manager who delivers the excess return x_t for sure , and it holds for all realizations y_t of the stochastic asset. The same holds if the skilled manager generates a distribution of excess returns $\{X_t\}$. This concludes the proof of theorem 1.

It should be clear from the proof that there are many different options strategies that satisfy the probability bounds given in the theorem. We also note that such a strategy need only be executed once per reporting period, and that it can be carried out within a short timeframe. For example, the mimic might choose a single day on which the options expire, and choose a strike price so that the Black-Scholes value of the option at the start of the day satisfies $v = 1 - 1/x_t$. Assuming he is not cleaned out by the end of the day, he can remain passively invested in the benchmark asset until the end of the period, and the target excess return, x_t, will have been achieved. Note, however, that the gap between the exercise probability p and the option value v increases with the expiration period Δ. Thus by using longer-dated options the mimic may be able to achieve an even lower probability of crashing and even higher earnings in expectation, that is, the bounds given in the theorem are conservative.

3. Numerical examples

To see how the theorem can be applied, let us run through an illustrative calculation. Suppose that a high-ability manager has a strategy that delivers excess returns of 10 percent every year for ten years relative to a risk-free rate of 5 percent. Thus the fund's annual return before fees is $(1.10)(1.05) = 1.155$. Under a two and twenty contract the manager's annual fee will be: 2 percent of 1.155 for management (assuming he collects on the year-end value) and 20 percent of .155 -.05 = .105 or a 2.3 percent bonus. Altogether this comes to 4.4 percent of the year-end value. The theorem says that a manager with no ability has a strategy that also delivers excess returns of 10 percent with some probability and *in expectation* earns *at least* 4.4 percent of a fund that is invested solely in the risk-free asset.

Let us see what this implies in dollar terms. Suppose that a fund starts with $100 million and is invested in the risk-free asset growing at 5 percent for ten years. The theorem shows that a no-ability manager has a strategy that, *in expectation*, earns at least 4.4 percent per annum of a fund that is compounding at 5 percent per annum before fees. After fees it is compounding at (1.05)(1 - .044) = 1.0038 per annum. This means that over ten years his *expected earnings* are somewhat in excess of $40 million. Of course, the high ability manager earns more in dollar terms; in fact, his total earnings over amount to 4.4 percent of a fund that is compounding at a rate of 15.5 percent before fees and 10.4 percent after fees. Over ten years this amounts to over $80 million. But $40 million is still a very large payoff for a manager who is not delivering positive alpha in expectation.

More generally, suppose that a skilled manager generates excess returns $\alpha > 0$ every year relative to a benchmark asset with constant returns $r > 0$. Let τ be his annual take (under a given fee contract), and assume that he takes the fraction τ out of the fund at the end of each reporting period. Let $\lambda = (1+\alpha)(1+r)(1-\tau)$. It can be shown that, over T years, his earnings per dollar initially invested in the fund will be

$$\frac{\tau\lambda(\lambda^T - 1)}{(1-\tau)(\lambda - 1)} \quad . \tag{8}$$

By contrast, a no-ability manager who generates fake α every year by the piggyback strategy will have expected earnings equal to or greater than

$$\frac{\tau\lambda(\lambda^T / (1+\alpha)^T - 1)}{(1-\tau)(\lambda - 1 - \alpha)} . \tag{9}$$

It follows that the ratio of a no-ability mimic's earnings to the earnings of an α-ability manager over T years is at least

$$\frac{(\lambda - 1)(\lambda^T / (1+\alpha)^T - 1)}{(\lambda - 1 - \alpha)(\lambda^T - 1)} . \tag{10}$$

4. Utility functions, new money, and equilibrium

The previous analysis shows that a mimic who has no special investment skill can do very well in terms of expected earnings. To achieve this, however, he takes a small risk in every period that his fund will blow up and he will go out of business. A risk-averse manager may not wish to engage in such a strategy, nor would someone who wants to build a long-term reputation. To address these issues we need to consider the utility function that a manager (of whatever skill level) is trying to optimize.

In general, utility will depend on the manager's discount rate, degree of risk aversion, and his ability to attract new funds by building a track record. All of these factors can be incorporated into an extension of the preceding framework. Of particular interest is the manager's ability to attract new money. There is a substantial literature on this phenomenon and a wide variety of models have been proposed to explain how it works; see among others Gruber (1996), Massa, Goetzman, and Rouwenhorst (1999), Chevalier and Ellison (1997), Sirri and Tufano (1998), and Berk and Green (2004). Here we shall show how to incorporate flow-performance relationships into our framework at a high level of generality without committing ourselves to a specific model and without fundamentally altering the conclusions.[8]

Let $Z_t = Z_t(x_1,...,x_{t-1}; y_1,..., y_{t-1})$ be a random variable that describes how much new money a fund will attract in the current period as a function of the returns in previous periods. We shall assume that Z_t is multiplicative, that is, its realization z_t is the *proportion* by which the fund grows (or shrinks) in period t due to the net inflow of funds. Thus the realized growth of the fund between the start and end of period t is $x_t y_t z_t$. We assume further that the manager is paid for bringing in new money only to the extent that it increases the size of the fund. Assume further that new money is accepted only at the start of each period.

[8] One issue we will not consider is the effect of performance on managers' continued employment in settings where portfolio management is delegated. For a discussion of this question see Chevalier and Ellison (1999) and Dagupta and Prat (2006).

Then the manager's earnings in period t, per dollar in the fund at the end of the previous period, can be written $\phi_t(\vec{x}, \vec{y})$ just as before. Hence *after fees* the fund grows by a factor of $(1-\phi_t(\vec{x}, \vec{y}))x_t y_t z_t$ between the end of the previous period and the end of the current period.

Let w_0 be the *initial size* of the fund. Given any realization $(\vec{x}, \vec{y}, \vec{z})$, the size of the fund at the end of period t is

$$w_t(\vec{x}, \vec{y}, \vec{z}) = w_0 \prod_{1 \le s \le t} [1-\phi_s(\vec{x}, \vec{y})] x_s y_s z_s \,. \tag{11}$$

Therefore the manager's earnings in period t can be expressed as

$$\phi_t(\vec{x}, \vec{y}) x_t y_t z_t w_{t-1}(\vec{x}, \vec{y}, \vec{z}) = \phi_t w_t / (1-\phi_t) \,. \tag{12}$$

Let $u(\cdot)$ be the manager's Bernoulli utility function for income, and let $0 < \delta < 1$ be his *discount factor*. It follows from (12) that the *utility* $U(\vec{X}, \vec{Y}, \vec{Z})$ of a strategy that delivers stochastic excess returns $\vec{X}$ over T years is

$$U(\vec{X}, \vec{Y}, \vec{Z}) = \sum_{1 \le t \le T} \delta^t E[u(\phi_t w_t / (1-\phi_t))] \,. \tag{13}$$

In principle, given $\vec{Y}$ and $\vec{Z}$ one can compute the optimal strategy $\vec{X}$ for a manager of any given ability as the solution to a Markov decision problem. (Here "ability" refers to the size of the manager's strategy set: someone with greater ability has a larger set of strategies and can generate a wider variety of returns, including higher returns, than someone with lesser ability.) An even more complex problem would be to compute the equilibrium strategies $\vec{X}^1, \vec{X}^2, ..., \vec{X}^n$ for any given set of *n* managers of different abilities, where the flow of funds to each manager *i*, say $\vec{Z}^i$, depends on the strategies pursued by the other managers. Absent more specific information we do not know how to find the equilibrium of such a game. Nor is it clear, given the complexity of the situation, that the players would be able to find it either. However the mimicking argument provides useful information about the possible outcomes whether the system is in equilibrium or not.[9]

To see this, suppose that *some* manager is following *some* strategy that delivers a stochastic return series $\vec{X}$ that earns the manager a large take of a large fund, and therefore makes a lot of money. (Empirically we know that there are such managers.) This manager's strategy may or may not be an optimum given his utility function; this is not crucial to the argument. The theorem shows that a mimic can achieve the same take. Moreover, he will attract the same amount of new money because his track record exactly mimics that of the first manager until the fund crashes. We have therefore established the following generalization of our theorem: *a no-ability manager can assure himself of a take that is at least as high as that of any skilled manager, where the former's take is from a fund that is growing at the rate of the benchmark asset and is attracting new money at the rate that the skilled manager is attracting it.*

[9] This is not to say that it is impossible to study equilibrium effects on fund managers' behavior; for example, Dasgupta and Prat (2006) study the impact of career concerns on trading behavior. In a different framework, Spiegler (2006) studies the equilibrium properties of the market for quacks that has potential application to charlatans in financial markets. Of course, these papers make much more specific assumptions than we do about how players respond to a given performance record.

5. Restructuring the incentives

As we have already said, the fact that managers can manipulate their returns in order to increase their fees is well-known. The contribution of this paper is to show how far this idea can be taken and to provide an analytical handle on how much can be earned in the process. In this section we apply the framework to the question of how to design a fee structure that screens out unskilled managers and eliminates the incentive to manipulate performance. At least four approaches have been suggested in the literature: a) impose high water marks – incentives are paid only when the fund's return is higher than anything it achieved in the past; b) pay the incentive only at the end of the fund's life on the basis of final excess returns; c) require the manager to hold an equity stake; d) levy monetary penalties if returns fall below some specified level. We shall consider each of these remedies in turn. A basic conclusion is that, while they partially reduce the incentive to manipulate returns, they cannot eliminate manipulation entirely, nor can they screen out the unskilled managers from the skilled ones.

a) Impose high water marks

In practice, a contract often stipulates that the manager can earn his bonus only when the fund's total return exceeds the highest return it ever achieved in the past. Equation (2) shows how to express this idea formally for the two and twenty fee structure. Unfortunately it is very easy for a no-ability manager to get around this requirement: all he has to do is mimic a series of excess returns (as in the proof of the theorem) while investing the funds in a safe asset such as Treasury bonds. The fund is guaranteed to grow in every period until it crashes, so the high water mark condition is satisfied and his earnings are not impaired.[10]

b) *Postpone incentive payments*

To be concrete, suppose that the fee is 'two and twenty' except that the twenty percent incentive can only be levied at the end of T years when the fund closes down and distributes its assets back to the investors. Let us ignore any new money that may come in and compute the expected earnings as a function of the realized returns $(\vec{x},\vec{y})$. Assuming that the two percent management fee is deducted at the end of each period, we find that the manager's total (undiscounted) earnings are

$$(.02/.98)\sum_{1\le t\le T}(.98)^t[\prod_{1\le s\le t}x_s y_s]+.20[\prod_{1\le t\le T}x_t-1][(.98)^{T-1}\prod_{1\le t\le T}y_t]. \qquad (14)$$

A mimic can still make a great deal of money under such a fee structure: until the fund crashes he collects the management fee on an amount that is compounding at a very high rate (namely $\prod_{1\le s\le t}x_s y_s$), and he collects the 20 percent incentive on the final size of the fund if it does not crash before then. The key point is that *the compounding from the fake excess returns exactly offsets the probability of crashing in each period*. Hence a risk-neutral and patient manager is not deterred by the

[10] Goetzmann, Ingersoll, and Ross (2007) show how to estimate the manager's expected earnings under high water mark contracts using an options pricing approach.

postponement of the fee.[11]

To see this concretely, let us do a simple calculation. Suppose that a no-ability mimic piggybacks an excess return of ten percent onto a risk-free bond yield five percent. Suppose also that the fund starts with 100 million dollars, attracts no new funds, and the manager is paid his bonus at the end of ten years. The probability that the fund does not crash before then is $(1.1)^{-10}$ or about 38.6 percent. At this point, however, the fund has grown to $[.98]^9[(1.1)(1.05)]^{10} \times 100 = 352$ million dollars. Hence the *expected* value of the final incentive payment is

$$.20[(1.1)^{10} - 1][(.98)^9(1.05)^{10}] \times 100 = 43.3 \text{ million dollars;} \qquad (15)$$

moreover this does not include any of the management fees collected along the way. If we discount the final payment at the risk-free rate of five percent, then the result is about 26 million dollars in expectation plus the management fees. This does not amount to much of a deterrent, especially when we consider that the costs of implementing a piggyback strategy are minimal.

---

[11] Risk neutrality is not a particularly restrictive assumption. The reason is that a risk-averse manager can diversify his risk by starting a family of funds, all ostensibly under separate management and all using independent piggyback strategies. In this case his realized earnings will, with high probability, be about the same as his expected earnings. For all practical purposes he can therefore be treated *as if* he were risk-neutral.

c) *Require the manager to hold an equity stake*

Let $\theta \in (0,1)$ be the size of the manager's equity stake, that is, the proportion of the fund's value that he is required to hold during the fund's lifetime T. We begin by noting that this requirement is easy to undermine, because the manager can always take positions in the derivatives market that effectively offset the gains and losses generated by his share of the fund. However, even if such offsetting positions can be prohibited, the requirement does not act as a deterrent.

To see why, let us consider the case of a safe benchmark asset that generates a fixed stream of returns $\vec{y}$ and drop them from the notation. Let us also ignore any new funds that may be attracted to the enterprise. (These simplifications do not change the conclusions in any important way.) Then we can write the fees in the form $\phi(\vec{x})$, where $\vec{x} = (x_1, x_2, ...x_T)$ is the series of excess returns generated over T years. A skilled manager who generates these returns with certainty will have final wealth (per dollar initially in the fund) equal to

$$\theta \prod_{1 \le t \le T} x_t + (1-\theta)\phi(\vec{x}). \qquad (16)$$

The theorem shows that an unskilled manager can generate this same series with probability $1/\prod_{1 \le t \le T} x_t$. His expected wealth at the end of the period is composed of two parts: the expected value of his own stake before fees, which is exactly θ

(because in expectation he cannot generate excess returns); and the expected fees from the investors, which amount to $(1-\theta)\phi(\vec{x})/\prod_{1\le t\le T} x_t$. Hence the unskilled manager's expected end-wealth, per dollar of initial fund value, is

$$\theta + (1-\theta)\phi(\vec{x})/\prod_{1\le t\le T} x_t \ . \tag{17}$$

It follows from (16) and (17) that the *ratio* of the unskilled to the skilled manager's end-wealth is $1/\prod_{1\le t\le T} x_t$, which is the same as the ratio of their earnings when they are not required to hold an equity stake.

d) Assess penalties for underperformance

Theoretically this is the most satisfactory approach, but it still does not succeed in screening out the no-ability mimics. As before, let us assume that the benchmark asset is safe and that no new money comes in, hence we can drop $\vec{y}$ and $\vec{z}$ from the notation. Consider a contract $\phi(\vec{x})$ that calls for negative payments for sequences $\vec{x}$ that deliver inferior returns. We do not need to specify which returns trigger negative payments, but we will assume that the *largest penalty* that could ever arise is δ per dollar of the fund's initial value. The payments must be enforceable, so the manager must put δ in escrow until the end of period T.

Consider some series of returns $\vec{x} = (x_1, x_2, ..., x_T)$ that can be generated with certainty by a skilled manager. Conditional on this realization of returns, the skilled manager's end-wealth, per dollar of initial fund size, is $\delta + \phi(\vec{x})$.

However, if he does not open the fund to investors and simply applies his skills to managing the amount δ, he would have $\delta \prod_{1 \le t \le T} x_t$ instead of $\delta + \phi(\vec{x})$. Therefore his participation constraint is

$$\phi(\vec{x}) > \delta(\prod_{1 \le t \le T} x_t - 1) \ . \tag{18}$$

Now consider a no-ability, risk-neutral manager who mimics the sequence $\vec{x}$. His expected end-wealth, per dollar of initial fund size, is

$$\phi(\vec{x}) / \prod_{1 \le t \le T} x_t - \delta(1 - 1 / \prod_{1 \le t \le T} x_t) \,. \tag{19}$$

Since he does not know how to generate excess returns in reality, and the mimicking strategy is essentially costless, his participation constraint is

$$\phi(\vec{x}) / \prod_{1 \le t \le T} x_t > \delta(1 - 1 / \prod_{1 \le t \le T} x_t) \,. \tag{20}$$

It follows that *any contract with penalties that keeps out the unskilled risk neutral managers keeps out all the skilled managers as well.*

6. Conclusion

In this paper we have shown how *mimicry* can be used to manipulate reward systems in financial markets. The framework is particularly useful in environments like financial markets where the game is complex and the equilibria are difficult to pin down. The approach shows how much a mimic can earn under different incentive structures, and why commonly advocated reforms of the incentive structure cannot be relied upon to screen out unskilled managers. We have also argued that it does not much matter from the investors' standpoint whether the managers are unscrupulous and trying to deceive investors or are merely deceiving themselves. It is easy for both types of managers to look like they are delivering alpha for extended periods of time, when in fact they are merely hiding risks in the tail. While we are certainly not the first to make this point, the methodology we have proposed permits one to make precise estimates of the *extent* to which the system can be gamed under a variety of modeling assumptions.

What then are the implications for the hedge fund industry? Essentially we have shown that the market is vulnerable to entry by managers who have no particular skill, but whose lack of skill is difficult to detect based solely on their track records. In other words, the hedge fund industry has a potential lemons problem (Akerlof, 1970). The combination of low entry costs and high expected rewards means that the market could be inundated by new entrants who would depress expected returns and increase the rate at which funds blow up. The upshot is that investors may eventually lose confidence and stop investing, which would put both the good and the bad managers out of business together. Thus the challenge for the truly skilled managers is to find a way to distinguish themselves from the potential mimics. This can be achieved by providing greater transparency about the strategies they are using and more precise and verifiable information about the risks that their investors are exposed to. A simple and robust conclusion that emerges from our analysis is that managers' track records are not, by themselves, sufficient to the task.

References

Ackermann, Carl, Richard MacEnally, and David Ravenscraft. 1999. "The performance of hedge funds: risk, return, and incentives." *Journal of Finance* 54: 833-874.

Agarwal, Vikas, and Narayan Y. Naik. 2004. "Risks and portfolio decisions involving hedge funds." *Review of Financial Studies* 17: 63-98.

Akerlof, George A. 1970. "The market for 'lemons': quality uncertainty and the market mechanism." *Quarterly Journal of Economics* 84: 488-500.

Berk, Jonathan B. and Richard C. Green. 2004. "Mutual fund flows and performance in rational markets." *Journal of Political Economy*, 112: 1269-1295.

Brown, Stephen, William N. Goetzmann, and Roger G. Ibbotson. 1999."Offshore hedge funds, survival and performance: 1989-1995." *Journal of Business* 72: 91-118.

Carpenter, Jennifer. 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.

Chevalier, Judith, and Glenn Ellison. 1999. "Career concerns of mutual fund managers." *Quarterly Journal of Economics*, 114: 389-432.

Dasgupta, Amil, and Andrea Prat. 2006. "Financial equilibrium with career concerns." *Theoretical Economics,* 1: 67-93.

Fung, William, and David Hsieh. 1997. "Empirical characteristics of dynamic trading strategies: the case of hedge funds." *Review of Financial Studies* 10: 275-302.

Fung, William, and David Hsieh. 1999. "A primer on hedge funds." *Journal of Empirical Finance* 6: 309-331.

Fung, William, and David Hsieh. 2000. "Performance characteristics of hedge funds and CTA funds: natural versus spurious business." *Journal of Financial and Quantitative Analysis* 35:291-307.

Gale, David. 1974. "A curious Nim-type game." *The American Mathematical Monthly* 81: 876-879.

Goetzmann, William, Jonathan E. Ingersoll, and Stephen A. Ross. 2003."High water marks and hedge fund contracts." *The Journal of Finance* 58: 1685-1717.

Goetzmann, William, Jonathan E. Ingersoll, Matthew Spiegel, and Ivo Welch. 2007. "Portfolio performance manipulation and manipulation-proof performance measures." *Review of Financial Studies*, 20: 1503-1546.

Guasoni, Paolo, Gur Huberman, and Zhenyu Wang. 2007. "Performance maximization of actively managed funds." Working Paper, Boston University and Columbia Business School.

Gruber, Martin J. 1996. "Another puzzle: the growth in actively managed mutual funds." *Journal of Finance* 51: 783-810.

Hasanhodzic, Jasmina, and Andrew W. Lo. 2007. "Can hedge-fund returns be replicated? : the linear case." *Journal of Investment Management* 5: 5-45.

Hodder, James E. and Jens Carsten Jackwerth. 2007. "Incentive contracts and hedge fund management." *Journal of Financial and Quantitative Analysis* 42: 811-826.

Hull, John C. 2009. *Options, Futures and Other Derivatives*, *7th Edition*. New York: Prentice-Hall.

Ibbotson, Roger G. and Peng Chen. 2006. "The A,B,Cs of hedge funds: alphas, betas, and costs." Working Paper 06-10, Yale International Center for Finance, New Haven CT.

Kat, Harry M., and Helder P. Palaro. 2005. "Who needs hedge funds: a copula-based approach to hedge fund return replication." Working Paper, Cass Business School, City University London.

Lo, Andrew W. 2001. "Risk management for hedge funds: introduction and overview." *Financial Analysts' Journal*, November/December Issue: 16-33.

Massa, Massimo, William N. Goetzmann, and K. Geert Rouwenhorst. 1999. "Behavioral factors in mutual fund inflows." Yale ICF Working Paper No. 00-14.

Sirri, Erik R. and Peter Tufano. 1998. "Costly search and mutual fund inflows." *Journal of Finance* 53: 1589-1622.

Spiegler, Ran. 2006. "The market for quacks," *Review of Economic Studies* 73: 1113-1131.

Starks, Laura T. 1987. "Performance inventive fees: an agency theoretic approach." *Financial and Quantitative Analysis* 22: 17-32.