Illustration by Elliott Golden

The Not-So-Real McCOY

By Dean P. Foster and H. Peyton Young

The hedge fund industry has consistently attracted some of the world's best and brightest investors with the promise of earning enormous wealth by delivering outsize returns. But as longtime academics Dean P. Foster and H. Peyton Young show, it is surprisingly easy for far less talented managers to create "fake" alpha, mimicking the returns of their more skilled and scrupulous peers. The only way to protect the industry, the authors argue, is through greater transparency.

The current shakeout in the hedge fund industry is capturing headlines, as managers who suffered large double-digit losses in 2008 are likely to have little choice but to shutter their funds. But even if the past year's turbulent financial markets had been more forgiving, the business would be facing a serious challenge from another source. Hedge funds' fee structure, combined with an almost total lack of transparency, makes the industry vulnerable to invasion by low-quality entrants who could undermine returns and trigger a collapse of confidence. The trouble is that when hedge fund investors are not allowed to look "under the hood," it can take a very long time for them to tell whether a manager is consistently generating true alpha or is simply having a run of good luck. Even worse, this information gap provides an opportunity for outright charlatans to enter the market — looking just like the real McCoys — without getting caught.

In this article we show how easy it is to generate "fake alpha" when investors cannot see what you are doing, and how much money you can make in the process. Although this potential problem has been discussed in the academic literature and is understood by some sophisticated players in the industry, it is less widely appreciated how serious it is and how difficult it will be to fix. The problem cannot be rectified simply by tinkering with the hedge fund fee structure. It is not enough to defer performance fees for managers or require them to hold an equity stake, nor will it work to levy monetary penalties on managers who underperform. Returns can, in effect, be manipulated using a certain options trading strategy. We argue that the only solution is far greater transparency and that the industry itself has a strong interest in providing it.

To understand this idea, consider Figure 1, which shows the total value (before fees) of an imaginary hedge fund compared with the value of a fund invested solely in the Standard & Poor's 500 index from January 1998 through December 2007 (the good old days). This is a pretty impressive performance.

It turns out that you can reproduce exactly this sequence of returns with high probability without actually generating alpha in the long run. In what follows, we shall show how to fake these returns, and more generally how to mimic the performance record of any skilled manager you wish to target without having the slightest idea about how that manager is actually generating returns. In fact, the target sequence of returns need not correspond to the track record

of an actual manager; it suffices that the record look as if it could have been generated by a manager with real talent.

The technique for achieving this, which we call "piggybacking," is simple and almost costless to implement. It is not the same as replication or cloning, which are strategies designed to reproduce the statistical properties — that is, the sources of returns and the corresponding risks — of a given hedge fund or class of funds. The point of piggybacking is to reproduce a specific sequence of returns with a high degree of probability. In this way the fund manager earns the performance or incentive fees associated with these returns and attracts new money into his fund just as if he were a star manager.

Let's say you want to deliver a performance record like that in Figure 1. Managers with similar records are lionized and attract large amounts of money. In fact, it is surprisingly easy to duplicate this performance using the piggyback strategy as long as investors cannot see your positions, and as long as you are willing to accept a small annual probability that your fund could go bust.

FIGURE 1: CUMULATIVE TOTAL RETURN FOR HYPOTHETICAL HEDGE FUND VS. S&P 500 INDEX

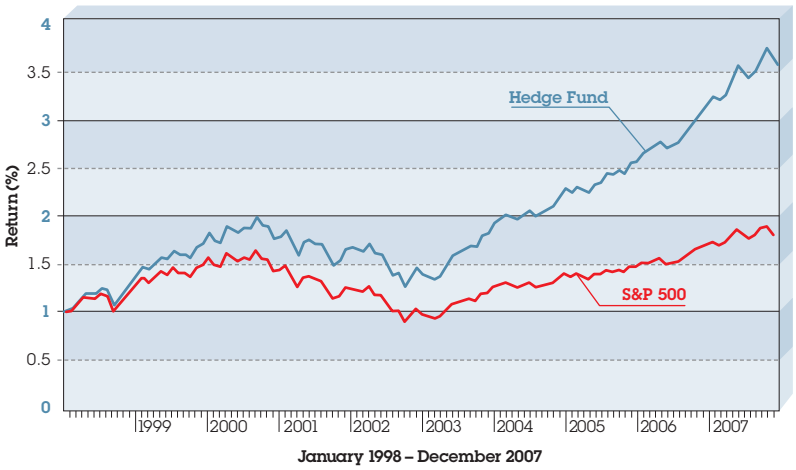

FIGURE 2: CUMULATIVE TOTAL RETURN FOR PIGGYBACKING VS. RISK-FREE BOND YIELDING 4 PERCENT*

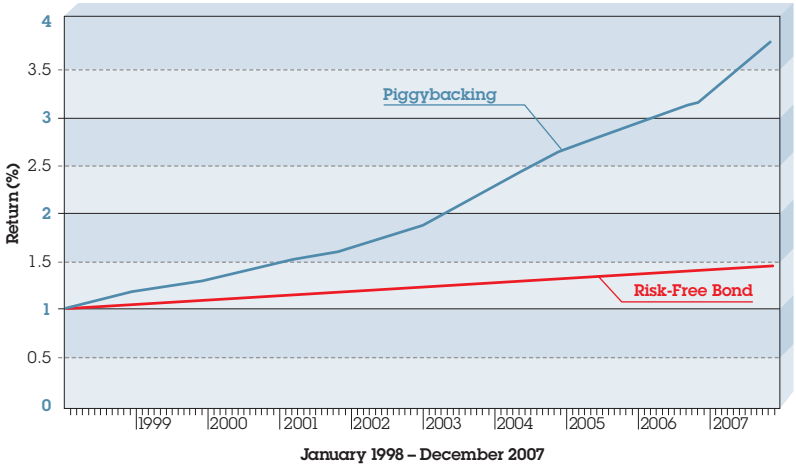

*The size of the piggybacked return is varied from year to year, the average being 10 percent.

Source: Foster and Young.

Here's how to do it: At the start of the year, invest all your funds in the S&P 500. Once a year take a short position in a bundle of asset-or-nothing puts on the S&P 500 that have a nearby expiration date. (An asset-or-nothing put pays out one share of the index if and only if the closing price is less than the strike price on the expiration date. You can create such a derivative by combining a plain-vanilla European put with a cash-or-nothing put, both of which are routinely traded on exchanges.) Let α be the target amount by which you plan to inflate your total return over the next twelve months — the amount of fake alpha. For example, to achieve the result in Figure 1, you would take $\alpha = 0.07$, which will generate an annual return that is 7 percent higher than the return on the S&P 500.

Once you have established the target value of α, choose the strike price so that the options are exercised with probability $\alpha/(1 + \alpha)$. In our present example, the strike price would be chosen so that the probability of exercise is approximately 0.07/1.07 = 0.065, or 6.5 percent. Now go short the maximum number of puts you can cover. The idea is to go for broke: If the puts are exercised, the fund will be cleaned out; but if they are not exercised, you will increase the number of shares in the fund by the factor $(1 + \alpha)$.[1] In the latter case you just sit back and wait until the end of the year (or any 12-month period), at which time you report that the fund grew by a factor of $(1 + \alpha)$ times the total return on the S&P 500 in that year. (During the entire time you have been fully invested in the S&P 500, which you used as collateral on your options position.) To the investors it looks as if you generated excess returns, and you collect a substantial performance fee. In reality you took a gamble and got lucky.

There is always the chance your fund will go bust, but this is not much of a deterrent. The reason is that piggybacking generates sizable values of alpha without much risk. As we have just seen, for example, you can generate alpha of 7 percent a year, and the risk of crashing is only 6.5 percent a year.

Of course, someone's suspicions might be aroused by the fact that the fund is returning exactly 7 percentage points more than the S&P 500 year in and year out. However, this can be avoided by varying the target amount of fake alpha from one year to the next. In some years you might inflate the index's total return by 5 percent, in other years by 9 percent, and so forth. The same basic method can be used to recreate any target sequence of returns that you choose.

Using the piggyback strategy, you can generate M times the return on the S&P 500 (or any other market portfolio) with a probability of at least 1/M over any range of time you choose. For example, you can deliver twice the total cumulative return on the S&P 500 over five years, or ten years, or any other target period, and the chances are one in two that your fund will not crash before then. You can deliver three times the return on the S&P 500 over any target period and the chances are one in three the fund will not crash, and so forth.

Naturally, the more you pump up returns by these methods, the higher the probability that the fund will crash. As

previously shown, however, the probability of crashing is not all that high on an annual basis for excess returns that look very impressive. Furthermore, if you are risk-averse, there is a simple way to spread the risk: Just start several funds under different names and run them in parallel, using independent piggybacking strategies. The probability is high that at least one of them will survive, yielding performance fees that make up for poor results at some or all of the others.

There is another reason why the probability of crashing does not act as much of a deterrent: As long as your fund does not crash, the size of the fund — and hence your postponed earnings — keeps growing at an exponential rate. Consider the previous example with fake alpha equal to 7 percent. The probability that the fund crashes within ten years is about 50 percent. But if it hasn't crashed in that time, it has grown by 7 percent annually over and above the growth rate of the underlying asset (in this case, the S&P 500). This is a very large number. From 1998 to 2007 your fund will have grown by more than 3.5-fold, and your postponed earnings would have grown at the same rate. So even if there is a 50 percent chance your fund won't survive for ten years, it could be a bet worth taking.

The issue is not merely one of keeping charlatans out of the market. A manager with no skill who thinks he can beat the market could also follow the piggybacking strategy, or something similar. Like the charlatan, the unskilled manager takes on large risks because that is how he maximizes his earnings, but he believes (mistakenly) that he is seizing on arbitrage opportunities that will produce above-average returns for his investors. One type of manager is deceiving himself; the other is deceiving his investors. Operationally, there is no way to distinguish between these situations.

Many commentators seem to think that fees are to blame. We disagree. Indeed, we shall show that none of the commonly proposed reforms to the hedge fund fee structure are going to keep out the charlatans or, more generally, the managers with no skill who mimic the performance of others. One feature that is already written into many contracts is the high-water mark, which stipulates that no further performance fees can be earned until a fund's return exceeds the level at which the previous such fees were paid. In the situation shown in Figure 1, for example, the manager would not be paid any performance fee from 2001 through 2003, because the fund's value decreased during that time (even though it stayed well ahead of the S&P 500).

It is easy to change the game to circumvent this constraint. All you have to do is piggyback fake alpha on top of the returns from a completely safe asset, such as short-term government bonds, which do not fluctuate much in value. For example, the piggyback strategy can generate fake alpha of 10 percent a year on top of a risk-free rate of 4 percent a year, and the probability that the fund will not crash in any given year is 1/1.10 or about 91 percent.[2] In this case the fund's annual growth is more than 14 percent[3] until it crashes. Of course, this steady growth might look a bit suspicious, but you could dress it up by targeting a value of alpha that varies from one year to the next (see Figure 2). In either case the high-water mark keeps getting higher, and you keep collecting performance fees until the fund crashes.

Another common proposal is to postpone all performance fees for a predetermined number of years. Doing so would appear to fix the problem, because a piggybacker's fund will probably crash before the fees are paid. The difficulty is that the postponement may have to be very long to have much effect. Suppose that no performance fee is paid until five years have elapsed. A mimic can set up shop and pad the returns on the S&P 500 by an 7 percentage points a year. After five years he will collect his performance fee with a probability of more than 70 percent[4]. If the fee is postponed for ten years, he will collect it with probability of more than 50 percent.

> "The industry is vulnerable to entry by managers who have no particular skill but whose lack of skill is difficult to detect."

These are pretty high numbers, but there is a further kicker: The asset base on which the postponed fee is computed grows at the rate of the S&P 500 compounded by an additional 7 percent per year. Over long periods of time, the S&P 500 has grown at about 9 percent a year (though it is easy to forget this given its recent performance). Therefore, as long as the fund does not crash, its annual growth rate is 16.6 percent[5]. In five years such a fund will more than double; in ten years it will more than quadruple. Hence even if the incentive fee is postponed for ten years, the increased size of the fee at the end of the waiting period offsets the risk that the fund might crash in the meantime.

The only real deterrent in this case is the manager's impatience to get the money, but waiting ten years for the prospect of a payment in the tens or hundreds of millions of dollars does not seem like much of a hardship. (And remember that the manager is earning management fees all along the way.) We conclude that only a lengthy postponement — say, 20 years or more — would have much bite. But postponing fees for this long would dissuade many skilled managers from participating in the market too.

Finally, we consider the notion (rarely seen in practice)

ON THE WEB
Visit **www.alphamagazine.com** to read the full working paper, "The Hedge Fund Game: Incentives, Excess Returns, and Performance Mimics," by Dean P. Foster and H. Peyton Young, on which this article is based.

that managers should be penalized for bad performance just as they are rewarded for good performance. In our view this is not likely to happen, but in any event it won't solve the problem. The essential difficulty is this: Really bad performance means that the manager loses a lot of money. How can the investors be sure the manager is going to pay the penalty? The answer is to hold the potential penalty in escrow in case the bad outcome occurs. (Obviously, this money cannot be invested in the fund or it could be blown away too; hence, it must be kept in something safe, like bonds yielding a risk-free rate.)

If the amount held in escrow is sufficiently large, mimics would be deterred from entering the market. But there is no way to differentiate ex ante between the mimics and the real McCoys. Thus the latter would also be required to post a bond. A little calculation shows that they won't agree to this because they would do better by investing an equivalent amount in their private hedge fund and not taking any money from outsiders. (This rather surprising result is proved in the companion paper on our Web site.)[6]

What are the implications of this analysis for the hedge fund industry? Essentially, we have shown that the industry 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 short, the hedge fund industry has a potential "lemons" problem. This term was coined by economist George Akerlof to describe the used-car market, where sellers tend to have much more information about the reliability of their cars than do potential buyers. This leads buyers to insist on lower prices to compensate for their risk. But then the owners of cars that actually are of high quality will withdraw them from the market, which means that the remaining cars will be, on average, even riskier. The result is a downward spiral in prices and a situation where no one can sell a car at a reasonable price.

The hedge fund industry could be facing a similar situation. The root of the problem is lack of transparency: If investors have only track records to go on, they cannot be sure whether they are dealing with a skilled manager or a low-quality manager who is merely mimicking a skilled one. The problem is potentially much worse than in the used car market because it is so easy for unskilled (or unscrupulous) entrants to set up shop. It is as if anyone could manufacture a car in their garage that looks turbocharged for a while and eventually blows up.

What is the prognosis for the industry as a whole and what are the possible remedies? We predict that, as low-quality imitators come in, average performance will deteriorate and the number of fund closures will rise. This process may already be under way, though it is being masked for the moment by closures resulting from market turbulence.

The problem cannot be fixed simply by reforming the fee structure. The reason is that it costs relatively little to enter the business and it is easy to mimic the track records of highly successful managers. Hence any reform in the fees that is potent enough to drive out the mimics is likely to drive out skilled managers too. Rather, the root of the problem is lack of transparency. Skilled managers need to find a way to distinguish themselves from the low-quality entrants. Here the analogy with the used-car market provides some clues to the solution. Just as a buyer can hire a mechanic to look under the hood of a potential purchase, so hedge fund managers may have to allow professional intermediaries (acting on behalf of potential investors) to have extensive access to their books and trading strategies, not just once but on an ongoing basis. Alternatively, individual fund managers may find it advantageous to operate under the umbrella of a large organization that can guarantee the product, just as the owner of a used car may prefer to sell it through a dealer rather than try to market it himself.

The bottom line is that the present arrangement, in which investors have no guarantees against downside risk and must rely solely on funds' track records to determine fund quality, is likely to unravel sooner or later. The industry leaders have a strong interest in making sure that this doesn't happen.

Dean P. Foster *is the a professor of statistics at the Wharton School of the University of Pennsylvania. He has been a council member of the Game Theory Society since 2005. His primary areas of research are machine learning, game theory and finance. He has also done work on measuring risk in financial markets, making good predictions in the context of data mining and learning the strategic choices made by an adversary.*

H. Peyton Young *is James Meade Professor of Economics at Oxford University and a senior fellow in Economic Studies at the Brookings Institution in Washington. Previously, he was Scott and Barbara Black Professor of Economics at Johns Hopkins University. He is a past president of the Game Theory Society, a fellow of the British Academy and a fellow of the Econometric Society. His current research interests are strategic learning, game theory and applications of game theory to finance. He has also written on the diffusion of innovations and the evolution of social norms.*

ENDNOTES

1. Here is the calculation: Suppose that you have s shares in the fund and you sell n puts. Each put is worth $\alpha/(1+\alpha)$ shares — its expected value — because by assumption the time to expiration is short. (For longer-dated options the computations involve Black-Scholes pricing, but the idea is similar.) By selling n puts you can therefore buy an additional n $[\alpha/(1+\alpha)]$ shares, so you now have $s + n\ [\alpha/(1+\alpha)]$ shares altogether. Solving the equation $n = s + n\ [\alpha/(1+\alpha)]$, it follows that you can sell or go short $n = (1+\alpha)s$ puts, and this is the number of shares you will have if they are not exercised.

2. At all times you keep your fund fully invested in the desired benchmark asset, in this case short-term government bonds that mature in a year. Once a year you take a short position in cash-or-nothing puts on a stochastic asset (like the S&P 500) that have a nearby expiration date. The strike price is chosen so that the puts are not exercised with probability 1/1.10. More generally, if α is the target amount by which you want to inflate your returns, choose the strike price so that the probability is $1/(1+\alpha)$ that the puts are not exercised.

3. The fund grows by a factor of 1.144, or 1.04 x 1.10, a year.

4. The calculation is: 1.07^{-5}, or 0.713.

5. The annualized growth is 1.166, or 1.07 x 1.09.

6. See "The Hedge Fund Game: Incentives, Excess Returns, and Performance Mimics," Working Paper No. 07-42, Wharton Financial Institutions Center, University of Pennsylvania, November 2007, at http://fic.wharton.upenn.edu/fic/papers/07/p0742.htm.

By Dean Foster and H. Peyton Young