Chasing Performance is a Dangerous Game

Robert Ferguson, Jason Greene, and Carl Moss
September 2009

Introduction

A common refrain from investors is that they have bad luck hiring investment managers. The future performance of the investment managers they hire tends to be disappointing relative to their expectations and the managers' historical performances. The managers tend to exhibit performance regression toward zero.

Investors' "bad luck" hiring managers may not be bad luck, but a consequence of how they hire managers. For example, there will be performance regression toward zero to the extent that good past performance relative to other managers, however measured, is a prerequisite for hiring an investment manager.

Some investors chase performance. They hire investment managers with high performance rankings and fire investment managers with low performance rankings. Chasing performance tends to have the result that hired investment managers exhibit performance regression toward zero, that good managers often will not be hired, and, if good managers are hired, that they may not be retained for long. Chasing performance is a dangerous game. Yet, most investors chase performance to some degree (directly or indirectly), hence suffer its adverse impact to some degree.

This paper illustrates just how insidious chasing performance can be. It also provides a framework for mitigating the detrimental effects of chasing performance; by putting historical performance in its proper analytical perspective, where the focus is on the investment process and historical performance plays a secondary role.

To see why chasing performance leads to disappointing future performance, consider the following parable.

The Parable of the Investment Managers

Assume a world where there are one thousand investment managers with true performance of zero. Suppose each of the managers has a ten-year performance record.

The managers' historical performance is the sum of true performance (zero) plus luck. Some lucky managers will have achieved exceptionally good performance and some unlucky managers will have achieved exceptionally poor performance.

Suppose an investor hires the ten managers with the best historical performance. These are managers with exceptional performance records. Nevertheless, the managers' true performance is zero. They were just lucky. Therefore, the hired managers' future performance over time will average about zero. The investor will attribute the managers' disappointing future performance to bad luck. However, the real reason for the disappointing future performance is the investor's hiring strategy.

This parable ignores investors' real life sophistication. Real life investors understand Modern Portfolio Theory (MPT) and know how to do sophisticated statistical performance analysis.

To see why chasing performance leads to disappointing future performance for real life investors, consider the following Modern Portfolio Theory version of the parable.

The Modern Portfolio Theory Version of the Parable of the Investment Managers

The managers are the same as before.

Now the investor knows that a manager's historical performance may be due to luck. He also knows about market related return, due to beta, and beta adjusted relative return, alpha. Perhaps most important, he knows that a manager's value added is alpha and how to test its statistical significance. Therefore, he defines performance as alpha.

The investor decides to perform a statistical test of the managers' alphas to find and hire managers who have a positive true alpha. The standard statistical approach equates this to identifying and hiring managers with "statistically significant" historical alphas.

The test procedure consists of using regression analysis to test the null hypothesis that the managers' true alphas are zero. The investor obtains each manager's quarterly returns and subtracts the quarterly interest rates from them. He also obtains the benchmark's quarterly returns and subtracts the quarterly interest rates from them. The result is two time series of quarterly excess returns, one for the manager and one for the benchmark.[1] The investor then regresses each manager's quarterly excess returns against the benchmark's excess returns.

The regression output includes each manager's historical alpha and the probability of achieving a historical alpha at least as large as that observed assuming the manager's true alpha is zero. This probability is termed the significance level of the manager's historical alpha.

Suppose that a manager's historical alpha is significant at the 1% level. Then there is only a 1% chance of achieving the manager's historical alpha or better if the manager's true alpha is zero. Either the investor has observed an unlikely event or the null hypothesis that the manager's true alpha is zero is false, the manager's true alpha is positive, and the investor has observed a reasonably likely event.

When the alpha significance level is as low as 1%, the standard statistical approach is to reject the null hypothesis that the manager's true alpha is zero in favor of the alternate hypothesis that the manager's true alpha is positive.[2]

The lower a manager's alpha significance level the better. Aside from the above statistical reasoning, this is because a low alpha significance level, such as 1%, tends to be due to the combination of a positive average quarterly alpha with a low standard deviation of quarterly alphas, a combination of attractiveness and consistency. It is the combination that matters; a high historical alpha alone does not imply a low alpha significance level.

The standard statistical approach suggests that one way of avoiding the dangers of using performance rank as a hiring criterion is to use alpha significance level rank instead. Consequently, the investor hires the ten investment managers with the lowest alpha significance levels. To his surprise, he suffers the same performance regression to the mean phenomenon. Why? Because all of the investment managers have a true alpha of zero and this is what all of the managers realise over time.

No matter how the investor ranks the investment managers' historical performance records, he is ranking on noise, not signal.

[1]Excess return is defined as return less the interest rate.
[2]A significance level of 1% does not imply that it is 99% sure that the manager has a positive true alpha.

Real Life

In real life, a manager's historical performance is the result of signal (positive, negative, or zero) and noise. Ranking on historical performance ranks on both signal and noise. The most attractive historical performances tend to have a positive contribution from noise, hence appear more attractive than they are. Performance regression toward zero is a foregone conclusion. This is true regardless of how historical performance is defined.

The difference between the parables and real life is that some managers have positive true alphas. Investors know that managers with positive true alphas will beat managers with zero true alphas over time. They know that good managers will beat bad managers over time. But how much time does it take?

Many investors believe that a ten-year performance record for a group of managers is sufficiently long to make it easy to spot the good managers. In fact, it is unlikely that the good managers will stand out.

Posit a good manager whose true average relative return is 200 basis points (bps) annually and true tracking error (standard deviation of relative return) is 800 bps annually. This manager's information ratio is 0.25. To put this in perspective, an information ratio of 0.25 typically puts a manager near or into the top quartile of managers in popular manager universes.

Posit twenty bad managers with true average relative returns of 0 bps annually, true tracking error of 1000 bps annually, hence an information ratio of 0.00.

There is a dramatic difference between the good manager and the bad managers.

Enough Bad Managers Always Beat the Good Manager

The probability that the good manager beats all twenty bad managers over a ten-year period is only about 9.6%.[3] This implies that chasing performance leaves the investor with the good manager only about 9.6% of the time and with a bad manager about 90.4% of the time. The investor's average relative return will be only 19.3 bps annually, his tracking error will be 980.7 bps annually, and his information ratio will be only 0.024. This compares with 200 bps, 800 bps, and 0.25 for the good manager.

Skeptics might argue that the number of managers at a finals presentation typically is far less than twenty. This misses the point. The adverse filtering on historical performance begins early in the manager selection process. The universe of investment managers that the finalists are drawn from far exceeds twenty. The problem may actually be worse than depicted here.

Chart 1 summarises the probability that the good manager beats all the bad managers for various numbers of bad managers, given the previous assumptions, including a ten-year performance history. It is not likely that the manager with the best historical performance is the good manager.

Skeptics might argue that the good manager will beat the bad managers over time and that all that is required to discern the good manager from the twenty bad managers is a longer historical performance record. This is correct. But how much time is enough? It is impractical to wait forever in order to identify the good manager.

[3]This probability is obtained by combining the probability density of the good manager's relative return with the probability distribution of the largest relative return among a group of twenty bad managers.

It Takes Forever for the Good Manager to Stand Out from the Bad Managers

A practical approach is to ask how long a historical performance record is necessary to be 75% sure that the good manager will beat all the bad managers, i.e., have the highest historical relative return. Assuming the same good manager as before and twenty of the same bad managers as before, a 157 year historical performance record is required to achieve a 75% probability that the good manager will beat all the bad managers.

Chart 2 summarises the probability that the good manager beats all twenty bad managers for various numbers of years in the performance history, given the previous assumptions. It is not likely that a good manager can be identified among a group of bad managers from their performance histories. Long enough performance histories do not exist.

What the Investor Should Focus On and Seek

Denote true performance by α and historical performance by $\hat{\alpha}$.
The investor should be interested in estimating the probability that a manager's true performance is positive, given the manager's historical performance, denoted by $P(\alpha > 0 | \hat{\alpha})$. This probability depends on:

- The a priori probability that the manager's true performance is positive, denoted by $P(\alpha > 0)$.
- The probability of achieving the manager's historical performance if the manager's true performance is positive, denoted by $P(\hat{\alpha} | \alpha > 0)$.
- The probability of achieving the manager's historical performance if the manager's true performance is zero, denoted by $P(\hat{\alpha} | \alpha = 0)$.

The first probability comes from an analysis of the manager's investment process without regard to the manager's historical performance. The second and third probabilities come from a statistical analysis of the manager's historical performance.

The link between these probabilities is Bayes Theorem, the probability of causes.

Bayes Theorem, the Probability of Causes

In this section, performance is defined as alpha. However, the exposition applies to any definition of performance.

The investor observes a manager's historical alpha ($\hat{\alpha}$).

There are two possible "causes" for the manager's historical alpha. It could be the result of an effective investment process, one with a positive true alpha ($\alpha > 0$) or it could be the result of an ineffective investment process, one with a zero true alpha ($\alpha = 0$).

The investor would like to know the probability that the manager's true alpha is positive in light of his historical alpha. This is termed the probability that the manager's true alpha is positive given his observed alpha. It is $P(\alpha > 0 | \hat{\alpha})$.

The investor's statistical alpha-beta analysis provides the probability of the observed alpha given that the manager's true alpha is zero. It is $P(\hat{\alpha} | \alpha = 0)$. This probability tends to be low if the manager's observed alpha is high. There is not much chance that a manager with a true alpha of zero will have a high observed alpha.

The investor can change the statistical alpha-beta analysis to provide the probability of the observed alpha given that the manager's true alpha is positive. It is $P(\hat{\alpha} | \alpha > 0)$. This probability tends to be higher than $P(\hat{\alpha} \| \alpha = 0)$ because it is more likely that a manager with a positive true alpha will achieve a high observed alpha than a manager with a zero true alpha.

The investor must also assess the a priori probability that the manager's true alpha is positive. This probability is $P(\alpha > 0)$. It is obtained from an analysis of the manager's investment process without regard to the manager's historical performance.

Bayes Theorem connects these probabilities as follows.

$$P(\alpha > 0 \mid \hat{\alpha}) = \frac{P(\alpha > 0)P(\hat{\alpha} \mid \alpha > 0)}{P(\alpha > 0)P(\hat{\alpha} \mid \alpha > 0) + (1 - P(\alpha > 0))P(\hat{\alpha} \mid \alpha = 0)} \quad (1)$$

Examples

Suppose the investor examines a manager's investment process and concludes that there is only modest reason for thinking that it should be effective. The investor decides that the a priori probability that the manager's true alpha is positive is 0.25.

$$P(\alpha > 0) = 0.25 \quad (2)$$

The investor then uses the standard alpha-beta statistical analysis to obtain the probability of the manager's observed alpha given that the manager's true alpha is zero. Suppose the observed alpha is positive. Then a reasonable expectation for this probability might be about 0.25, i.e., the observed alpha is above the manager's true alpha of zero, hence is not very likely to be achieved.

$$P(\hat{\alpha} \mid \alpha = 0) = 0.25 \quad (3)$$

Next, the investor uses the standard alpha-beta type statistical analysis to obtain the probability of the observed alpha given that the manager's true alpha is positive. Suppose the observed alpha is positive. Then a reasonable expectation for this probability might be about 0.50, because the observed alpha is then in line with the manager's assumed true alpha.

$$P(\hat{\alpha} \mid \alpha > 0) = 0.50 \quad (4)$$

The investor then evaluates Equation (1).

$$P(\alpha > 0 \mid \hat{\alpha}) = \frac{0.25x0.50}{0.25x0.50 + 0.75x0.25} = 0.40 \quad (5)$$

The investor concludes that the probability that the manager's true alpha is positive given the manager's historical performance is 40%.

Consider the same analysis applied to a manager with an investment process that is credible enough to assign a 0.75 a priori probability of a positive true alpha. This is the only probability that changes.

$$P(\alpha > 0) = 0.75 \quad (6)$$

Substitution in Equation (1) results in:

$$P(\alpha > 0 \mid \hat{\alpha}) = \frac{0.75x0.50}{0.75x0.50 + 0.25x0.25} = 0.86 \qquad (7)$$

The investor concludes that there is an 86% probability that the manager has a positive true alpha given the manager's performance history.

Note how much the answer changes based on the investor's assessment of the a priori credibility of the manager's investment process.

Conclusion

This paper presented intuition and analytical results that demonstrate that chasing performance is a dangerous game. Yet, it is a game played by almost all investors to some degree. The paper also presented an approach to hiring and firing managers that mitigates the adverse impact of performance chasing.

A careful examination of almost any investor's investment manager hiring and firing process is likely to reveal that there is a substantial component of performance chasing. Sometimes it is obvious, e.g., when there is a policy of firing a manager if he has negative performance after three years. Other times it is subtle, e.g., when the initial phase of the manager search process strongly weights attractive historical performance. No matter the form that performance chasing takes, it tends to produce future relative returns that are disappointing compared to expectations.

Historical performance alone is not an effective basis for identifying a good manager among a group of bad managers. This does not mean that historical performance is useless. Rather, it means that it must be combined efficiently with other information. The correct use of historical performance relegates it to a secondary role. The primary focus in manager choice should be an analysis of the investment process.

Chart 1

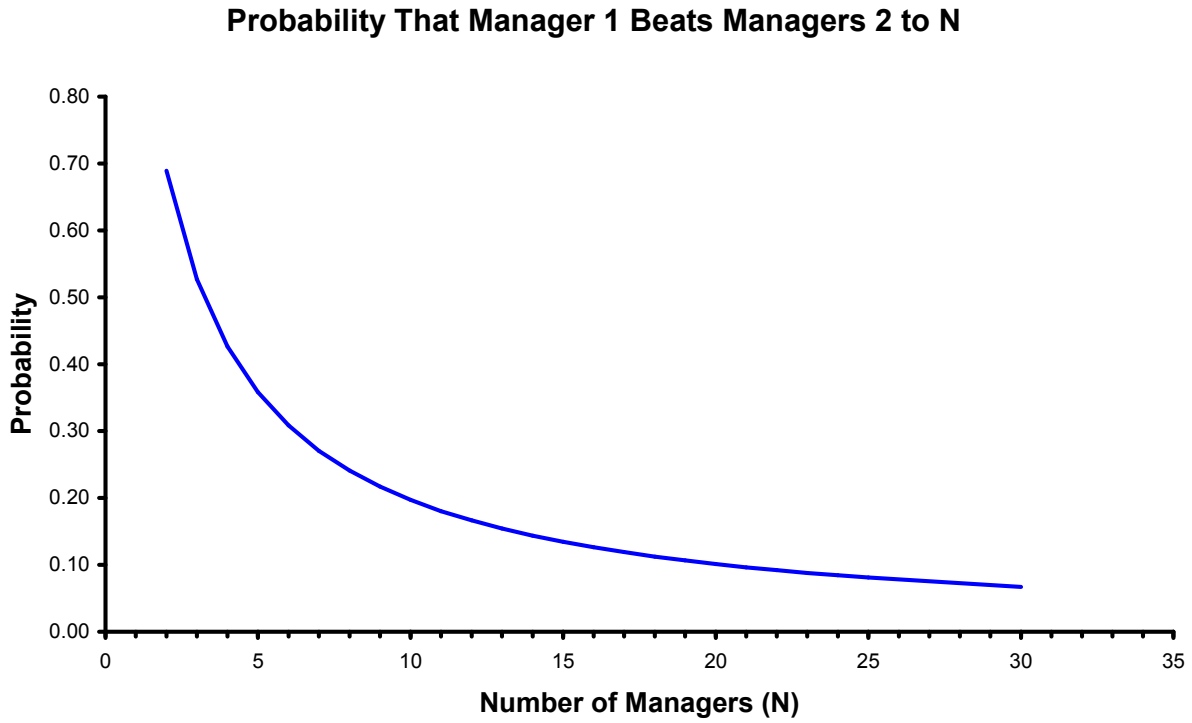

	Manager 1	Managers 2 to N
Mean relative return (annual, decimal):	0.02000	0.00000
Standard deviation of relative return (annual. decimal):	0.08000	0.10000
Information ratio:	0.25000	0.00000
Time horizon (years):	10	

Hypothetical illustration based on the assumptions shown above.

Chart 2

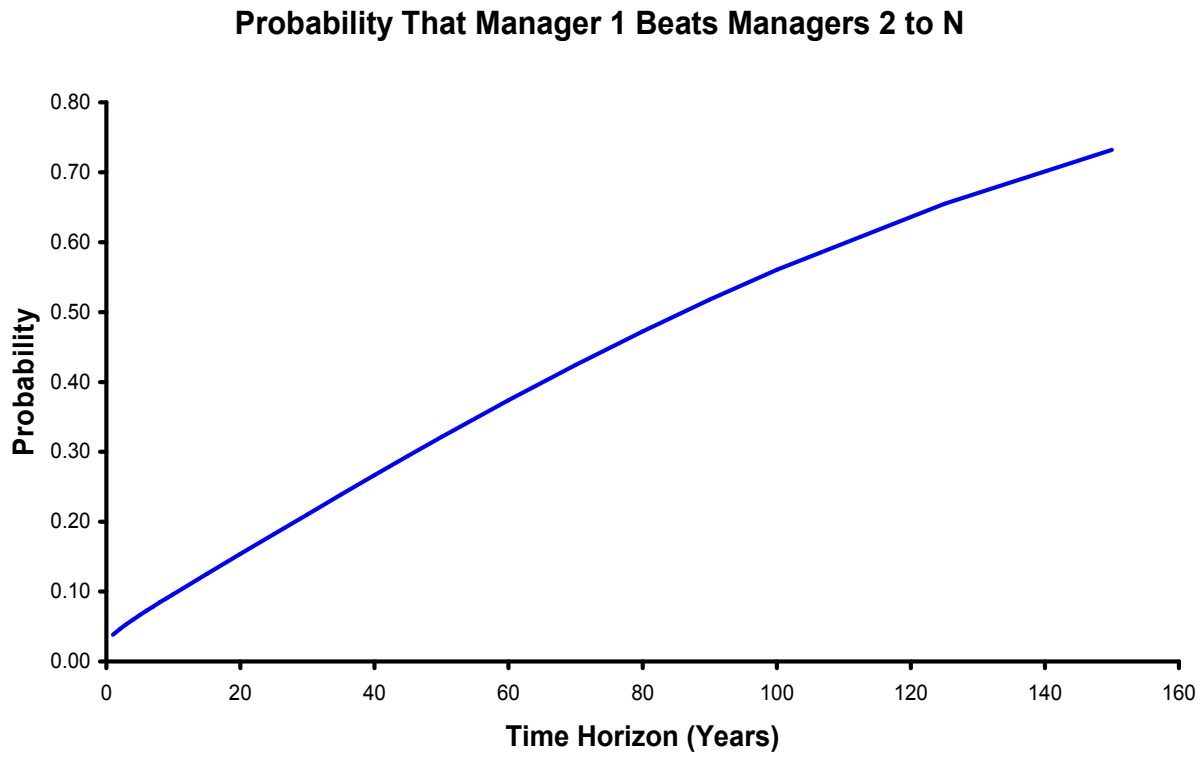

	Manager 1	Managers 2 to N
Mean relative return (annual, decimal):	0.02000	0.00000
Standard deviation of relative return (annual, decimal):	0.08000	0.10000
Information ratio:	0.25000	0.00000
Number of managers:	21	

Hypothetical illustration based on the assumptions shown above.