

A Martingale Test for Alpha

Dean P. Foster, Robert Stine, and H. Peyton Young

December 31, 2008

We present a new method for testing whether a fund manager's track record allows us to infer that he is able to beat the market with high probability or is just plain lucky. The test is based on the martingale maximal inequality. Unlike other standard approaches the test is robust to the assumed distribution of returns while retaining substantial statistical power. The method is illustrated using stock market data from 1926-2007.

1. Estimating alpha

Fund managers who deliver returns that are consistently higher than the returns from a market portfolio are said to generate 'alpha'. This is the extra return attributable to the manager's superior investment skill. Needless to say, such managers are in great demand and many of them earn large sums of money. But how can investors tell whether a given manager actually is able to generate positive alpha over an extended period of time?

The usual method for estimating alpha empirically is to regress a given series of returns generated by a fund manager against the returns generated by a market portfolio, such as the S&P 500. Let $t=0,1,2,3,\dots$ be discrete times at which returns are reported, say the end of each month. Consider first a fund that is invested entirely in the stock market, and let r_t^m be the return generated by the market portfolio in period t . Let r_t^f be the risk-free rate of return in period t , that is, the return available on a safe asset such as US Treasury bills. The *excess*

return in the period is $r_t^m - r_t^f$. Let us compare this with an alternative managed portfolio that generates returns r_t^p and associated excess returns $r_t^p - r_t^f$, $t = 1, 2, 3, \dots$. The usual way to estimate the manager's "alpha" -- his contribution to excess returns due to superior skill -- is to estimate a linear regression equation of form

$$r_t^p - r_t^f = \alpha + \beta(r_t^m - r_t^f) + \varepsilon_t. \quad (1)$$

If the estimated value of α is positive at a high level of statistical significance (using the t -test for the intercept), we conclude that we should put at least some of our money into the alternative portfolio. This follows from the Capital Asset Pricing Model (CAPM), which we shall discuss in more detail in section 7 (see Berndt, 1996; Campbell, Lo, and MacKinlay, 1997). However, the t -test is only valid if the errors ε_t satisfy standard assumptions, in particular, they should be independent and normally distributed with mean zero and identical variance. Unfortunately, there is no reason to think that in reality the errors will be so distributed. The problem is that the manager's investment strategy (for the alternative portfolio) may be *designed* so that the errors are not i.i.d. normal. In particular, it can be shown that under standard performance fees, managers have an incentive to follow investment strategies whose returns distribution is skewed with a fat left tail (Foster and Young, 2007). In such cases the use of the t -test is problematic.

The contribution of this paper is to demonstrate an alternative test for alpha that makes *no* parametric assumptions about the distribution of the fund's returns. It is simple to compute and can even accommodate the possibility that the manager

is deliberately trying to deceive the investor by creating return series that look good, i.e., that are designed to fool the t -test. All of this comes without much sacrifice in power, as we show in section 5.

2. A martingale model of investment returns

The key idea is to show that the space of possible investment strategies – including positions in options and other derivatives – can be represented at a high level of generality by a family of martingales. We will then apply the martingale maximal inequality to test for the presence of positive alpha. The null hypothesis will be that the manager has no special investment skill and cannot deliver long-run excess returns that are higher than the excess returns generated by the market portfolio (alpha equals zero). This is a natural assumption because, in a competitive market, one would normally expect that there are relatively few arbitrage opportunities and relatively few individuals who are able to take advantage of the few opportunities that there might be.

First we need some notation. Consider a market portfolio that begins at size 1. By the end of period t it has grown by the factor $(1 + r_t^m)$. Define the ratio

$$X_t = (1 + r_t^m)/(1 + r_t^f) \geq 0. \quad (2)$$

This is the *multiplicative excess return* delivered by the market portfolio relative to the risk-free rate. Similarly the multiplicative excess return delivered by the alternative managed portfolio is

$$Y_t = (1 + r_t^p)/(1 + r_t^f) \geq 0 . \quad (3)$$

If we define $Y_t = M_t X_t$, then $M_t \geq 0$ is the factor by which the manager over-performs (or under-performs) the market portfolio in period t . We shall refer to M_t as the *multiplicative excess return of the alternative portfolio relative to the market*.

This set-up turns out to be convenient from a mathematical standpoint. It is also quite general: virtually any investment strategy yields returns that can be represented as a nonnegative martingale $M_t \geq 0$. If a manager has no particular investment skill (alpha is zero), his choice set consists only of those nonnegative martingales M_t whose conditional expectation is 1. A skilled manager has a larger choice set that includes martingales with conditional expectation greater than 1; in fact we can take the size of the choice set as a general definition of “managerial skill”.

Suppose that we observe a series of realized returns from a market portfolio and also from an alternative portfolio over the periods $t = 1, 2, \dots, T$. Since we know the risk-free rate in each period, we can compute the t^{th} - period realizations x_t, y_t, m_t respectively of the random variables X_t, Y_t, M_t . The *null hypothesis* H_0 is that the manager of the alternative portfolio is not generating positive alpha, that is,

$$H_0 : \forall t, \quad E[M_t | x_1, \dots, x_{t-1}, m_1, \dots, m_{t-1}] = 1 . \quad (4)$$

Compound Excess Return Test (CERT). For each t , $1 \leq t \leq T$, let $C_t = M_1 M_2 \cdots M_t \geq 0$ be the compound excess return of a fund through period t relative to the market. The probability that this series was generated by a manager who cannot beat the market is at most $\min_{1 \leq t \leq T} (1/C_t)$.

Corollary. A manager's performance over a given period of time allows us to infer that he can beat the market with 95% confidence provided his portfolio grows by a factor of at least twenty-fold relative to the total growth of the market portfolio over the same period.

The proof is a straightforward application of the martingale maximal inequality (Doob, 1953), which for convenience we shall derive here.

Proof sketch. Under our assumptions, the compound excess return generated through period t , C_t , is a nonnegative martingale with conditional expectation 1 in every period. Given a real number $\gamma > 0$, define the random time $T(\gamma)$ to be the first time $t \leq T$ such that $C_t \geq \gamma$ if such a time exists, otherwise let $T(\gamma) = T$. By the optional stopping theorem (Doob, 1953, Theorem 2.1), $E[C_{T(\gamma)}] = 1$. Clearly, if $\max_{1 \leq s \leq t} C_s \geq \gamma$ then $C_{T(\gamma)} \geq \gamma$, hence $P(\max_{1 \leq s \leq t} C_s \geq \gamma) \leq P(C_{T(\gamma)} \geq \gamma)$. Since $C_{T(\gamma)}$ is nonnegative, $P(C_{T(\gamma)} \geq \gamma) \leq E[C_{T(\gamma)}] / \gamma = 1 / \gamma$. Thus we have shown that

$$P(\max_{1 \leq s \leq t} C_s \geq \gamma) \leq 1 / \gamma. \quad (5)$$

It follows that any observed sequence c_1, c_2, \dots, c_T of compound excess returns was generated with probability at most $\min_{1 \leq t \leq T} (1/c_t)$. \square

Notice the interesting point that the *length* of the series is immaterial for our test: what matters is the maximum compound return that was achieved at some point during the series.

3. Downside risk

It might at first seem counterintuitive that one would need to see a fund outperform the market by a factor of twenty-fold in order to be reasonably confident that the outperformance is “for real.” The reason is that the apparent outperformance could be driven by a strategy in which the fund loses everything with positive probability. It is easy to demonstrate a strategy of this nature for which the *CERT* bound is tight. Choose a number $\gamma > 1$, and consider the following nonnegative martingale with conditional expectation 1

$$C_0 = 1, \quad P(C_t = \gamma^{1/T} C_{t-1}) = \gamma^{-1/T}, \quad P(C_t = 0) = 1 - \gamma^{-1/T} \text{ for } 1 \leq t \leq T. \quad (6)$$

In each period the fund compounds by the factor $\gamma^{1/T}$ with probability $\gamma^{-1/T}$ and crashes with probability $1 - \gamma^{-1/T}$. It is straightforward to construct a strategy based on traded options that has exactly this property (Foster and Young, 2007). For this strategy, the probability that the fund’s compound excess return C_t ever exceeds γ is precisely $1/\gamma$.

Strategies like this are not purely hypothetical; as we have previously noted, it may be *optimal* for a manager with no skill to place investment bets that clean out his fund with positive probability. However, it would be easier for investors to infer positive alpha from a given series of returns if the manager could demonstrate that the fund is not going to be cleaned out. In other words, if the manager can demonstrably limit the downside risk -- say through the purchase of an insurance policy -- then the criterion in *CERT* is less demanding.

To be concrete, suppose that the manager can guarantee that, in each period his fund will never decrease by more than some fraction $\phi \in (0,1)$ of the corresponding market portfolio (assuming both have the same value at the start of the period). In effect the manager is restricting himself to martingales M_t that have expectation 1 and satisfy the lower bound $M_t \geq 1 - \phi$.

Consider the martingale $M_t^\phi \equiv (M_t - 1 + \phi) / \phi$. This is nonnegative and has expectation 1. Hence the earlier argument shows that

$$P[\max_{1 \leq t \leq T} \phi^{-t} \prod_{1 \leq s \leq t} (M_s - 1 + \phi) \geq \gamma] \leq 1 / \gamma. \quad (7)$$

It follows that if the fractional loss in each period (relative to the market portfolio) can be limited to ϕ , then the null hypothesis can be rejected at the level

$$p = \min_{1 \leq t \leq T} \phi^t \left[\prod_{1 \leq s \leq t} (M_s - 1 + \phi) \right]^{-1}. \quad (8)$$

It is straightforward to show that this bound is tight by investing $1-\phi$ of the fund in a risk-free asset at the start of each period. Note that the smaller the downside risk ϕ , the more modest the returns M_s can be and still meet a given level of confidence.

4. Gaming the t -test

We have already pointed out that investment strategies with a small probability of large losses are not just theoretical constructs, but are actually encouraged by the fees they can deliver for fund managers. Unfortunately the most common technique for estimating alpha – the t -test – can give totally misleading results against strategies of this type.

We can illustrate this point with the previous example showing that the *CERT* bound is tight. Fix $\gamma > 1$ and a number of periods T . There exists an options trading strategy such that, with probability $1/\gamma$, the manager delivers a multiplicative excess return equal to $\gamma^{1/T}$ in every period $1 \leq t \leq T$ without crashing. Assume for simplicity that the risk-free rate is a constant $r > 0$. Then with probability $1/\gamma$ the fund's value grows by the factor $\gamma^{1/T}(1+r)$ for T successive periods. For the sake of argument, assume that there is some small variation around this growth rate (say due to a noisy execution of the strategy), so that the period-by-period returns are nearly but not quite constant. Then the regression approach to estimating alpha (see expression (1)) will yield the estimates $\hat{\beta} = 0$ and $\hat{\alpha} = (\gamma^{1/T} - 1)$ with almost no variability and therefore an extremely large value of t . (Beta is zero because, by construction, the returns are completely uncorrelated with the market.) The t -test would therefore lead us to

conclude that the manager's alpha is positive at an extremely high level of significance. But this is not the correct conclusion. The problem is that we have not seen enough data to capture the possibility that the fund could lose a large amount of money.

To be concrete, suppose that we observe a fund's quarterly returns over ten years, so that $T = 40$. Imagine that our manager is gaming the investor using the martingale construction in (6) with parameters $\gamma = 2, r = .01, T = 40$. Such a fund appears to be delivering alpha equal to $1.01(2^{1/40} - 1) \pm \varepsilon$ or about 1.77% every quarter, where ε represents a small "tremble" in the execution of the strategy. These returns would appear to be very convincing if we estimate alpha using the t -test. The analysis in section 3 shows, however, that such a series can be generated with probability one-half. Therefore in reality we can only be 50% confident that alpha is positive.

5. Power and leverage

We have argued that *CERT* protects the investor against distributions that can be manufactured by a fund manager in order to lure investors into thinking that he can generate positive alpha when in fact he cannot. It might be objected, however, that trying to protect investors against all possible distributions leads to an overly conservative test. In this section we shall argue that this is not the case.

Suppose we knew in advance that a manager's strategy generates returns that conform to the assumptions of a standard model, such as lognormality and

independence. Armed with this knowledge, we could conduct an optimal test of significance, which in this case would be a t -test on the logged returns. We claim that our test ($CERT$) involves only a modest loss of power compared to this optimal test, and has the advantage that it applies whenever the returns form a nonnegative martingale.

Assume that the fund manager is generating compound “excess” returns $C_t \geq 0$ relative to the market portfolio, where C_t can be represented in continuous time by a Brownian motion of form

$$dC_t = \mu C_t dt + \sigma C_t dW_t. \quad (9)$$

Then C_t is lognormally distributed:

$$\log C_t \sim N((\mu - \sigma^2 / 2)t, \sigma^2 t). \quad (10)$$

The null hypothesis is that $\mu = 0$ and the alternative hypothesis is that $\mu > 0$. An unknown is how much variance σ^2 the manager’s strategy involves; this will depend on how much leverage he uses, which we shall assume cannot be observed. (The investor is interested in whether $\mu = 0$ or $\mu > 0$, not in the size of the variance per se. The reason is that, by adding a sufficiently small amount of the fund to his existing portfolio, he adds little variance but increases its expected return (provided $\mu > 0$).)

Consider an investor who is using *CERT* and significance level $p \in (0,1)$ to decide whether to reject the null. He observes a series c_1, c_2, \dots, c_T at discrete time intervals and rejects if and only if the largest, say c_i , satisfies

$$\log c_i \geq \log(1/p). \quad (11)$$

By assumption, $Z_t = \frac{\log C_t + (\sigma^2/2)t}{\sigma\sqrt{t}}$ is $N(0,1)$. Hence the investor rejects the null provided that

$$Z_t \geq \frac{\log(1/p) + (\sigma^2/2)t}{\sigma\sqrt{t}}. \quad (12)$$

Of course, the manager *wants* the investor to reject, because in this case the investor accepts the alternative, which is that the manager is able to deliver excess returns. Hence the worst case for the investor arises when the manager has no skill ($\mu = 0$) and chooses σ to minimize the right-hand side of (12). This occurs when $\sigma = \frac{\sqrt{2\log(1/p)}}{\sqrt{t}}$, in which case (12) becomes

$$Z_t \geq \sqrt{2\log(1/p)}. \quad (13)$$

It follows that the *critical value* for *CERT* at significance level p is

$$c_p = \sqrt{2\log(1/p)}. \quad (14)$$

Let us now compare this with the optimal test, which in this case is the z -test. (We assume there is enough data so that the normal can be used instead of the t -distribution.) Letting Φ denote the cumulative normal distribution, the critical value at level p for the z -test is

$$z_p = \Phi^{-1}(1 - p). \quad (15)$$

It can be shown that, for every $p \in (0,1)$, $z_p < c_p$, that is, *CERT* is more conservative than the z -test. We shall now show that when p is small, the loss in power $L(p)$ is quite modest.

By definition, $L(p)$ is the maximum probability that, for some value $\mu > 0$, the z -test would reject the null when *CERT* would accept:

$$L(p) \equiv \max_{\mu} P(z_p \leq Z_t + \mu < c_p) = 2\Phi(.5(c_p - z_p)) - 1. \quad (16)$$

Claim. $\lim_{p \rightarrow 0} L(p) = 0$.

Proof. When z is large the right tail of the normal distribution has the following approximation (Feller, 1971, p.193):

$$P(Z \geq z) \approx \frac{e^{-z^2/2}}{z\sqrt{2\pi}}. \quad (17)$$

Therefore when p is reasonably small, say $p \leq .01$, we have the approximation

$$z_p \approx \sqrt{2 \log z_p + 2 \log(2\pi / p)}. \quad (18)$$

Combining (14) and (18) we find, after some manipulation, that

$$c_p - z_p \approx \frac{\log z_p + .5 \log(2\pi)}{2\sqrt{\pi} \sqrt{2 \log(1/p)}} \leq \frac{\log c_p + .5 \log(2\pi)}{2c_p \sqrt{\pi}}. \quad (19)$$

Hence $c_p - z_p \rightarrow 0^+$ as $p \rightarrow 0^+$, which establishes the claim.

Figure 1 shows the power loss function $L(p)$. Note that the loss in power is only about 30% when $p = .01$.




Figure 1. Power loss function $L(p)$.

6. A portfolio test

The test described above is for returns generated by a single fund. In practice, investors may want to identify those funds *within a given set* of funds that generated positive alpha with high probability. This is a considerably more difficult problem. To illustrate, suppose that we can observe the returns for each of n funds over the same timeframe $t = 1, 2, \dots, T$. Let C_t^i be the compound excess return for fund i through period t . Suppose that we observe a particular fund, say i^* , whose returns are sufficiently high that we would reject the null at the 5% level using *CERT*. This does not mean we have 95% confidence that fund i^* is generating positive alpha. Suppose, in fact, that *none* of the n funds is able to generate positive alpha and that the excess returns are stochastically independent. It is straightforward to construct n independent nonnegative martingales, each with expectation 1, such that on average there will be $.05n$ funds that exceed the critical threshold $c_{.05}$. For example, out of 1000 funds there will, on average, be 50 funds that pass the single-fund threshold even though none of the funds is actually generating positive alpha.

More generally, suppose that we observe n nonnegative martingales $C_t^i, 1 \leq i \leq n$, over the periods $1 \leq t \leq T$. Let the null hypothesis H_0 be that $E[C_t^i | c_1^i, \dots, c_{t-1}^i] = 1$ for all i and for all t , where the C_t^i are assumed to be independent across i for each t . By the martingale maximal inequality we know that

$$\forall c > 0, \quad P(\max_{1 \leq t \leq T} C_t^i \geq c) \leq c. \quad (20)$$

Hence

$$\forall c > 0, \quad P(\max_{1 \leq i \leq n} \max_{1 \leq t \leq T} C_t^i \geq c) \leq n/c. \quad (21)$$

It follows that, to reject H_0 at level $p > 0$, there must exist one or more funds i such that

$$\max_{1 \leq t \leq T} C_t^i > n/p. \quad (22)$$

This is known as the *Bonferroni test*. The idea is quite general and applies to any situation in which multiple hypothesis tests are being conducted (Miller, 1981).

One limitation of the Bonferroni test, however, is that it lacks power when the returns are not independent. This will be the case, for example, if fund managers exhibit herding behavior in their choice of investment strategy. Here we propose a different test that is simple to conduct and has considerable power whether or not the returns are correlated.

Portfolio excess return test (PERT). Consider a family of n funds, where each fund i generates a series of compound excess returns $c_1^i, c_2^i, \dots, c_T^i$ over T periods. Create a portfolio consisting of equal amounts invested initially in each of the funds. The portfolio's compound return series is given by $\bar{c}_t = (1/n) \sum_{1 \leq i \leq n} c_t^i$, $1 \leq t \leq T$. The probability that none of the fund managers is able to generate positive alpha is at most $\min_{1 \leq t \leq T} (1/\bar{c}_t)$.

The proof is more or less immediate. Each of the n funds generates a nonnegative martingale of excess returns C_t^i that may or may not be independent across funds. The investor's portfolio is the nonnegative martingale $\bar{C}_t = (1/n) \sum_{1 \leq i \leq n} C_t^i$. The assumption that none of the managers can generate positive alpha implies that the conditional expectation of \bar{C}_t equals 1 in every period. The conclusion follows at once from the martingale maximal inequality.

We note that this test is at least as powerful as the Bonferroni test even when the returns are independent among funds. Indeed, the Bonferroni test rejects only if some realization satisfies $c_t^i > n/p$. But in this event $\bar{c}_t > 1/p$, which implies that the portfolio test rejects also.

7. Application to stock market investing

We illustrate the approach by evaluating two different strategies for investing in the stock market. The first is a hypothetical strategy that *by assumption* delivers small but positive excess returns in each period. The second is a strategy called *TEAM* (Target Equity Allocation Management) that was first proposed by Gerth (1999). In the first example we shall assume that the returns are normally and independently distributed, so that standard methods like the t -test apply. We show that in this case *CERT* yields similar answers to the t -test: both tests concur that alpha is positive at a high level of significance (which is correct because by assumption the strategy has positive alpha). In the second example (*TEAM*), the excess returns turn out to be highly correlated and are not normally distributed. Thus the *TEAM* strategy produces returns that are outside the scope of the usual modeling assumptions for the t -test but are within the domain of *CERT*.

To set the stage for these examples, let us briefly review certain features of the Capital Asset Pricing Model (CAPM), which provides a general method for assessing the quality of investments (see for example Berndt (1996) or Campbell, Lo, and MacKinlay (1997)). Let r_t^f denote the risk-free rate in period t , r_t^m the return from the market portfolio, and r_t^p the return from a hypothetical managed portfolio. The model posits that the managed portfolio's excess return $r_t^p - r_t^f$ can be written as a multiple, *beta*, of the excess return generated by the market, $r_t^m - r_t^f$, plus a martingale difference ε_t :

$$r_t^p - r_t^f = \alpha + \beta(r_t^m - r_t^f) + \varepsilon_t, \quad E[\varepsilon_t | \varepsilon_{t-1}, \varepsilon_{t-2}, \dots, \varepsilon_1] = 0. \quad (23)$$

The multiplier β is the covariance between returns on the market and the managed portfolio, divided by the variance of the excess returns, that is,

$$\beta = \text{Cov}(r_t^p - r_t^f, r_t^m - r_t^f) / \text{Var}(r_t^m - r_t^f). \quad (24)$$

If the managed portfolio generates positive *alpha* relative to the market ($\alpha > 0$), then by taking some funds out of the market and putting them into the portfolio we could diversify our investments and earn a strictly higher return with no increase in risk. Thus it is of considerable interest to know whether α is positive or not.

The standard way to test for positive α is to estimate (23) using a series of returns and apply the t -test to the intercept. This approach is valid provided the errors are i.i.d. normal. For our first example, we simulated a hypothetical

investment strategy that by assumption has positive alpha and normally distributed errors, namely, $\alpha = 0.01, \beta = 1.0, \varepsilon_i \sim N(0, 0.10^2)$. Thus, on average, the strategy generates alpha of 1% per month with a variance of 1%.¹ For the market returns we used the actual monthly returns for the S&P 500 from January 1926 through December 2007, that is, $n = 984$ months. The risk-free rate was assumed to be the interest rate on US Treasury bills in each month.

We ran 10,000 simulations and compared the outcome of the t -test and $CERT$ at the 5% level of significance, with the null hypothesis $\alpha = 0$. For a given simulation, the t -test rejects the null at the 5% level if the estimated value, of α divided by the estimate of the standard error, exceeds 1.65. (Here the number of degrees of freedom is so large that the t is effectively the normal.) By contrast, $CERT$ rejects the null if the compound value of the simulated fund exceeds the compound value of the S&P 500 by at least 20-fold at some time during the run.

Since the null hypothesis is false, each test should reject with high probability, which is the case as Table 1 shows. More importantly, the two tests yield quite similar results: the tests render the same verdict over 92% of the time. Figure 2 illustrates the same point in a different way, namely, the empirical distributions of p -values for the two tests are quite similar. This is rather surprising given that $CERT$ holds almost independently of the assumed distribution of returns, which is certainly not the case with the t -test.

¹ By assuming that the variance equals the mean, there is no incentive to leverage the investment strategy, that is, the growth rate is optimized when $\beta = 1$.

	<i>CERT</i>	
	Accept	Reject
<i>t-test</i>	Accept	1026 180
	Reject	572 8222

Table 1. Rejection and acceptance of the null ($\alpha = 0$) by the t -test and *CERT* when the true distribution of alpha is $\alpha = 0.01, \beta = 1.0, \varepsilon_t \sim N(0, 0.10^2)$. 10,000 simulated runs, each of 984 periods.




Figure 2. Distribution of p -values under the t -test and *CERT* for 10,000 simulated runs each of 984 periods.

We now consider a second investment strategy in which the t -test turns out to be quite inappropriate but *CERT* is still applicable. This strategy, known as *TEAM*, was first proposed by Gerth (1999), and is discussed in a subsequent article by Agnew (2002). Divide discrete time periods t into successive subsequences, each

of length n . A subsequence of time periods will be called an *epoch*. At the start of each epoch i , the *TEAM* strategy targets an amount S_i to be invested in stock, with the remainder held in cash (i.e., Treasury bills). At the end of each period t within epoch i , funds are moved between the stock account and the cash account so that the amount in stock at the start of the $(t+1)^{st}$ period is $(1+r_t^f)$ times the amount at the start of the t^{th} period. Thus if $r_t^m < r_t^f$ money is moved from cash to stock, and if $r_t^m > r_t^f$ money is moved from stock to cash. At the end of the epoch, the value of the *TEAM* portfolio is $S_i \sum_t (r_t^m - r_t^f)$ plus terms that are of much smaller magnitude.

Lowering the target S_i reduces the volatility of the returns, while raising the target increases the volatility. Gerth proves that if stocks have returns that are independent and identically distributed (not necessarily normal), and if the risk-free rate of return is constant, then there exists a choice of targets such that *TEAM* yields the same level of volatility as any given mixture of stocks and cash, and has a strictly higher expected return.

There is no guarantee, of course, that these assumptions hold in practice. Our purpose, however, is not to critique the *TEAM* approach, but to use it as a test case on which to apply our theory. Suppose that some fund manager had actually used *TEAM* to manage his portfolio over the period 1926-2007. Could we conclude that he generated positive alpha given the realized returns from the market over this same period?

Following Gerth, let us evaluate the specific *TEAM* strategy in which epochs are five years long, funds are re-allocated at the end of each year, and the target at

the start of each epoch is to put 70% in stock. This results in 14 epochs over the period 1926-2007. We compute the excess returns from *TEAM* in each of these 14 epochs and regress them against the excess returns generated by the market. The OLS estimate of alpha is $\hat{\alpha} = 0.000319$ with standard error 0.000904, so the *t*-statistic is .353, which has a *p*-value of about 0.35. Thus the null ($\alpha = 0$) cannot be rejected with much confidence.²

Unfortunately the *t*-test is not appropriate, because the residuals are not independent nor are they identically distributed, as one can see from Figure 3.

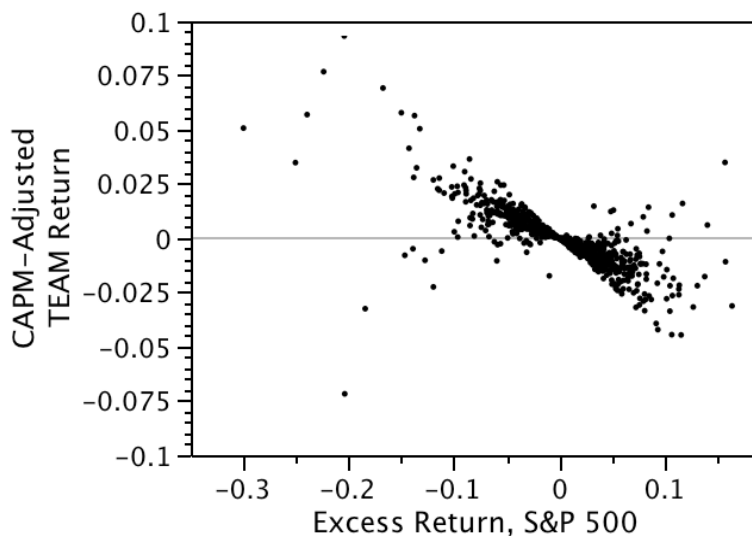


Figure 3. Scatterplot of residuals from regressing *TEAM* excess returns versus S&P 500 excess returns, January 1926-December 2007.

² Gerth (1999) made no claim that his strategy does produce positive returns relative to the market; rather, he argued that under certain conditions it will produce higher returns than a mixture of the market and cash.

In such a case *CERT* is a more appropriate test; moreover it is easy to compute. Figure 4 shows the compound excess return from *TEAM* relative to the compound returns from the market over the period January 1926-December 2007. It is evident that the series gets nowhere near the factor of 20 required to show that $\alpha \neq 0$ at the 95% confidence level. Indeed, the maximum ratio is less than 2, so we cannot be at all confident that this version of *TEAM* generates positive alpha compared to the market. Thus, in this case, the *t*-test and *CERT* yield similar *p*-values that do not permit rejection of the null hypothesis, but *CERT* is the more appropriate test because the much stronger distributional assumptions required for the *t*-test are not met.




Figure 4. Compound excess returns of *TEAM* strategy divided by compound excess returns of the S&P 500, January 1926-December 2007.

References

Agnew, Robert A., "On the TEAM approach to investing," *American Mathematical Monthly* 109 (2002), 188-192.

Berndt, Ernst R., *The Practice of Econometrics: Classic and Contemporary*. New York: Addison-Wesley, 1996.

Campbell, John C., Andrew W. Lo, and A.C. MacKinlay, *The Econometrics of Financial Markets*. Princeton NJ: Princeton University Press, 1997.

Doob, Joseph L., *Stochastic Processes*, New York: John Wiley, 1953.

Feller, William, *An Introduction to Probability Theory and Its Applications*, Princeton University Press, 1971.

Foster, Dean P., and H. Peyton Young, "The hedge fund game: incentives, excess returns, and performance mimics," Working Paper 07-042, Financial Institutions Center, Wharton School of Business, University of Pennsylvania, 2007.

Gerth, Frank, "The TEAM approach to investing," *The American Mathematical Monthly* 106 (1999) 553-558.

Miller, R. G., *Simultaneous Statistical Inference*, 2nd edition, New York: Springer-Verlag, 1981.