Debunking Some Myths of Active Management

M. Barton Waring and Laurence B. Siegel

M. Barton Waring
is a managing director and heads the Client Advisory Group at Barclays Global Investors in San Francisco, CA.
barton.waring@barclaysglobal.com

Laurence B. Siegel
is director of investment policy research at The Ford Foundation in New York, NY.
l.siegel@fordfound.org

In our travels we speak to hundreds of large investors, and a number of questions about active management and active managers come up over and over again. They come up because what we are saying is not always consistent with the conventional wisdom long accepted by traditional active managers. Looking at the persistence of some of this wisdom—often given credence only on the strength of decades of repetition—we have to wonder whether there is any merit to the saying that an untruth repeated often enough becomes the truth, and whether unchallenged myths assume the guise of fact.

We strongly believe that in the presence of skill active management can be successful. But we also believe it can be sold on its own merits without artificial arguments. So here we debunk some of the myths and stories often told in support of active management. We fear they do more harm than good, and sow confusion, misunderstanding, and ultimately distrust of healthy management disciplines.

Q: Should you use active management rather than index management in the less efficient markets like international equities, small-cap equities, or emerging markets?

Inefficiency means there is some information out there, available to some investors, that is not yet impounded into security prices. If among those who have access to this information are people with special skill at using it to forecast prices, they will generate a positive pure active return at the expense of those who do not have the information or are unable to interpret it and act on it. Having and exercising this forecasting skill is the only way to truly beat the market.

Active management in any market is a zero-sum game relative to that market's benchmark (before fees and costs). A market, after all, is just the grand sum of the investments held by all those participating in that market. Therefore it must be true that the weighted-average return of all the participants in that market is the same as the return of the market as a whole.

In essence, this is a tautology. And it has nothing to do with efficiency or inefficiency. *Even an inefficient market is a zero-sum game.*

This can best be demonstrated by a mental experiment. Consider a hypothetical emerging market country. Such a country might have fewer than a dozen stocks in its equity index. Many of these countries by any definition would be considered grossly inefficient: There is no equivalent to the U.S. Securities and Exchange Commission; there is rampant trading on inside knowledge; and no doubt there is much other evidence of inefficiency. Yet an index fund, using no selection skill whatsoever, will nonetheless capture the average return of that market—somewhere in the middle, between the winners and the losers.

Active participants in these markets will outperform persistently only if they have skill

in exploiting inefficiencies at the expense of other market participants. The unskillful active participants in such a market will lose, and there have to be some such participants, since the winners need someone with whom to trade (or there can be no winners).

As we said, active management is a zero-sum game—this is an undisputable "law of gravity" facing every investor. Inefficiency does not create an exception to this rule. Inefficiency is a necessary condition, but not a sufficient condition, for beating a benchmark other than through random luck. Special skill at forecasting is required.

Sometimes there seem to be convincing data to contradict this assertion. In the middle and late 1990s, most U.S.-based international equity managers beat their benchmarks by 5%, 7%, and more per year, leading many to argue that international investments should always be actively managed. What is going on here? Is there some flaw in the concept of the zero-sum game, a gotcha exploited by the smart guys out there doing it that the investment scientists have not yet squarely faced? The winning managers often attribute their success to the inefficiency of international equity markets.

It is now widely acknowledged that the U.S. managers gained this unusual performance by underweighting Japan relative to the MSCI EAFE when Japan was experiencing lower returns than the benchmark as a whole. Somebody had to hold the other side of this trade—one can't argue with the basic concept that the overall market is a zero-sum game. Every position taken by U.S. money managers underweighting Japan (relative to the MSCI EAFE version of market cap weights) had to be matched by an equal and opposite wager by some other investors who made losing bets, by *overweighting* Japan.

While there is no way to really know who these other players were, it seems likely that they were in large part Japanese companies, owning each other's stocks or holding their own treasury stock, and likely also individual and institutional Japanese investors. Most countries' investors hold a disproportionate percentage of their own local securities, the so-called home country bias.[1]

The interesting question is whether this underweight position by U.S. money managers evidences special investment insight or skill.

In the last two decades, Japan grew from a modest component of the non-U.S. equity market to the dominant component (about 60% at the high), and since then it has fallen back to a smaller position. It turns out that, as Japan's stock market soared during its earlier growth period, these same managers were also then underweight Japan, and their performance lagged. Over the entire period, their performance was relatively flat (*see Exhibit 1*).

So, looking at the decision to underweight Japan, over the entire history of that bet by U.S. managers, there is evidence that the recent strong performance is simply the lucky half of what in total was a mediocre market-timing bet. The managers, as a group, were unlucky in the first period while Japan's star ascended, and lucky in the second when it reversed. And the opposite is true for the investors holding the other side of this trade, with the same relatively flat result over the entire period.

So even while there are apparent anomalies, and while it might appear at first blush that active managers as a group can profit from inefficiency, the anomaly will disappear on critical examination. You should never hire active managers just because they are active, no matter how inefficient a particular market is. Skill is always necessary if you want performance that arises other than from luck. Individual managers might certainly have skill, but all managers as a group cannot make that claim.

Another caution might be in order with respect to investing in inefficient markets. Commissions, spreads, and market impact costs all go up with greater inefficiency, and may offset the ability to harvest alpha. It isn't easy to quantify the relations among inefficiency, alpha opportunity, and transaction costs, but there is some intuition that the other side of the trade will always try to protect itself against the trader who appears to have information. This will reduce whatever value a skillful player might otherwise harvest.

It has been our experience that alpha can be somewhat higher in some inefficient markets, but not by as much as one might hope, especially after transaction costs.

Q: Is it possible through special techniques to generate higher returns than the market, while taking less risk than the market?

We've heard this ability claimed by traditional active managers. This story is *sooo* appealing—we all want to believe a smart manager can add return and manage risk to a lower level. But . . . can it be true?

Let's build an easy, single-factor model for returns that separates the market-related components of return from the active components. In the spirit of simplicity we'll just look at a manager's *excess* returns (defined as returns over and above the risk-free rate):

$$r_p = \beta_p(r_m) + \alpha_p \quad (1)$$

In words, "The excess returns for a manager's portfolio can be described as the sum of the excess return for taking market risk through the benchmark (after adjustment for beta), plus an uncorrelated idiosyncratic return, a pure alpha. The beta is just a regression beta, the covariance of the manager's portfolio with the benchmark divided by the variance of the benchmark.

If that is the return model, the parallel risk model for this portfolio has to be:

$$\sigma_p^2 = \beta_p^2(\sigma_m^2) + \omega_p^2 \quad (2)$$

In words, the risk of the manager's portfolio is the sum of the beta-adjusted market variance, plus the uncorrelated variance of the alpha, known as omega (ω). There are no cross-terms, as the market and alpha terms are uncorrelated by definition.

But active managers are often over- or underexposed to the market component of the portfolio, as captured by the benchmark. One might have a single-factor beta of 0.8, for example. If a manager's single-factor beta is different from 1.0, meaning that its market-related risk differs from that of the benchmark, then the portfolio's market risk, represented in Equation (2) by the term $\beta_p^2(\sigma_m^2)$, will quite naturally differ from the natural risk of the benchmark (the benchmark of course does have a beta of 1.0).[2] Since the effect of the beta term is squared, this can be a substantial effect.

Back to the question. Is it possible to deliver higher returns than the benchmark, at lower risk?

Statistically, of course, it's possible for a particularly successful manager to deliver such a return pattern. Look again at Equation (2). There is only one way for a manager to deliver total portfolio risk below the benchmark's risk—to have a beta that is far enough below 1.0 to compensate for the pure active (omega) risk taken by the manager—and then a bit less than that. In plain words, the manager has to be persistently underexposed to the market in order to make such a claim.

And if you feed that low beta back into Equation (1), the low beta will mean that the benchmark component of returns will also be very low. So there is only one way for you as an investor to believe that this manager can deliver above-benchmark total returns—its alpha has to be extraordinarily high. But if that manager really does have the ability to deliver such a large alpha, wouldn't it be touting it directly, promising to deliver it on top of a proper benchmark exposure?[3]

Just as important, if you—the investor—really want less exposure to that benchmark (a market risk decision), a benchmark that was chosen during your strategic asset allocation (SAA) process, wouldn't you have made that decision in the course of that process? You chose to take the level of market risk represented by that benchmark, because you expect it to be rewarded over time. You don't want your manager to take you persistently out of the market by carrying a low beta. Managers are hired to deliver pure alpha, not to make unauthorized changes in your exposures to market risk from that you specified in your SAA policy.

There is a more sophisticated variation of this claim that we sometimes hear from managers: "We manage money in such a way that our active returns are negatively correlated with the benchmark, giving you lower risk than the benchmark."

But of course the part of an active return that is correlated with the benchmark isn't the pure alpha—by definition. If there is any part of the active return that is correlated with the benchmark, whether positively or negatively, this simply means the manager has a beta that is different from 1.0. So for a manager to say that it has a negatively correlated active return is just to say that it is underexposed to the market risk of the benchmark. We know this is not a good thing.

So, does active management take risk out of a portfolio? Not likely, if it's doing the right thing. Clear-eyed managers will use risk-control techniques to keep their market exposures closely aligned with those of the benchmark. And the best will use multifactor approaches, with multiple market exposure betas, or factor loadings, rather than the very simple single-factor model sketched out here, to better manage those multiple market risk exposures.

Q: Some traditional active managers claim they simply match the benchmark in up markets, but that they beat it in down markets. How can I test this claim?

This is a claim often made by traditional active managers, particularly value managers. (But it seems we hear it only during up markets.) No matter how high the manager's true mean alpha, with some predictable probability that manager will have periods of underperformance.

You can test this claim by separating the manager's periodic return data into up-market and down-market

baskets, and running regressions separately on each subset against an appropriate market benchmark (i.e., a value benchmark, or better, a multifactor set of component benchmarks). Look at the regression results to see if there is believable evidence that the alphas are higher in down markets than in up. Be careful to keep track of the statistical significance figures for these values.

In practice, it isn't likely you'll find believable support for the claim very often. But it isn't impossible.

Unfortunately, what is more likely happening to support such a manager's claim is very simple, and not good, as it has nothing to do with alphas or even with successful market timing. The manager may just be holding a low-beta portfolio, which will indeed outperform in down markets. But it will underperform by an equal amount in up markets. The manager is hoping you don't notice, since below-market returns in up markets are still up. The manager is taking market risk in addition to selection risk, and that isn't what that manager was hired to do.

Q: How are alphas distributed? Is it possible to hire a manager who only outperforms?

Position-by-position, day-by-day, the alpha distribution for any manager will be very much a function of that manager's specific style and trading behavior.

But we don't need to know the day-by-day distribution to understand the distribution of alphas over longer periods of time, which is the distribution that is of most interest to investors. Remember the central limit theorem? It can be summarized (very briefly) as saying that the average of a large number of independent random variables will approach being normally distributed—no matter what their underlying distributions are.[4]

So while we don't know what the trade-by-trade distribution of alpha is, we can say something useful about the distribution of alpha over many trades, many days, or more usefully, many months or years. So long as fresh new decisions are always being made on fresh new information, the average alpha of a manager's returns will approach a normal distribution over many periods (or lognormal, to be precise). So since managers are hired for long periods, their prospective alpha over any likely holding period should best be viewed as normally distributed even if their daily process is best characterized by a distribution that is Poisson, uniform, exponential, or something even more exotic.[5]

Why is this useful? It allows us to use standard tools of statistical analysis more confidently, many of which assume normality. If we didn't know the distribution of alphas over time, we'd be less confident of the usefulness of such statistics.

And it also tells us that we can't expect to ever find an active manager who never underperforms. No matter how high the manager's true mean alpha, with some predictable probability that manager will have periods of underperformance. Likewise, bad managers, with no skill and a true negative mean alpha, will predictably enjoy some portion of their time in positive alpha territory. Hence the difficulty in separating luck from skill.

Q: Since active managers are professionals, can't we expect them to generate positive alpha by taking advantage of the non-professionals in the market, even if the overall market is a zero-sum game?

Managers sometimes argue that, while across all market participants it may be true that active portfolio management is a zero-sum game, the professional participants like them take advantage of the non-professional participants, and predictably add value above the benchmark. As a result, it makes sense to hire professional managers simply because they are pros, who can be expected to predictably beat their benchmark over time.

We can find a little evidence that directionally supports this claim, but the extent of the evidence isn't sufficient to make the case that professional active managers should automatically be preferred over index funds.

What is that evidence? Brinson, Hood, and Beebower [1986] and Brinson, Singer, and Beebower [1991] have conducted empirical studies of the returns achieved by professional investors (that is, pension plans and other institutions and their managers, and mutual funds).[6] The average alphas (after all fees and costs) are slightly different in these studies, but it isn't unfair to summarize the results as indicating a negative alpha, about –0.5% per year. So if there is any "professional effect," it is clearly not sufficient to reliably beat the benchmarks.

Hmmm. The average professional return may be negative, but how does it stack up against a fair expectation based on theory? We mentioned that there might in fact be some professional effect, at least directionally. Here is why. The notion that markets are a negative sum game after fees and costs would justify a prediction that—in the absence of a "professional effect"—the average loss would be equal to fees and costs. We estimate these fees and costs to be well over 1% per year, the sum of manager fees at, say, 40 to 90 basis points, and explicit and implicit trading costs (slippage, opportunity costs, and market impact,

which are not directly observable) of 50 to 150 basis points. So, since the average underperformance isn't nearly as bad as theory predicts, the implication is that the professionals did in fact beat the non-professional market participants over the time periods studied.

But sadly not by enough to cover fees and costs, or to give the professionals a return higher than that of their passive benchmark. The non-professionals have done even worse (by implication—we don't have the data to prove it), but that's little consolation when you're getting beaten by your benchmark.

Q: Does it really take 80 years to show an active manager's alpha is statistically significant?

We can do some very quick-and-dirty statistical work to answer this question. We'll use the rule of thumb for calculating the statistical significance, a t-statistic, noting that it is just the pure information ratio (IR) multiplied by the square root of time.[7] If it requires 80 years to get a significant two-tailed t-statistic of 2, and calling the square root of 80 an even 9, this suggests a manager believes its information ratio to be in the proportions of 2/9, or only a bit more than 0.2. This isn't a trivially low information ratio; neither is it especially good. Our experience generally is that top-quartile managers will have an information ratio of 0.5 or above (see Grinold and Kahn [2000, p. 130]).

But what if this were a stronger manager, with a higher information ratio? How long would it take? Using the 0.5 IR that we believe is a marker for top-quartile managers, only 16 years are required to find a significant t-statistic. How about an outstanding manager, with an IR of 1.0? Only four years of such performance are required for statistical significance.

So, no, it doesn't take 80 years to show statistically significant alpha if the manager is really any good. Remember the basic reasons for thinking about statistical significance: One should reject positive alpha histories that are not statistically significant, as you can't support a conclusion from the data that the alpha is meaningful (see Waring and Siegel [2003]).

So while we as realists have to acknowledge that we aren't likely to persuade every investor to ignore completely all one-, three-, and five-year track records that are not statistically significant, we will at least caution you to ignore our advice at your own risk. And even when historic alphas are significant, use them carefully—you can't absolutely prove skill with a finding of statistical significance, but it's fair to consider it carefully as evidence of such skill.[8]

Q: Does holding too many active managers make a fund a closet index fund?

As we argue in Waring and Siegel [2003], active portfolios of anything, whether of securities or of managers, if not informed by skill (positive information coefficient, or IC) are necessarily just poorly constructed index funds. The expected return of such a portfolio, before costs, is the same as that of the benchmark.

Of course such a portfolio won't be a very good index fund. It will experience high tracking error—but that's all right because the costs will be high as well. And if it is not carefully constructed, it may have misfit risk as a result of style and size (capitalization) exposures, as well as potentially other exposures such as sector, that are at odds with those of the benchmark, when the components are summed across all the managers.

So the term *closet index fund* properly refers to any portfolio that claims to be active but that isn't in fact informed by skill, and so it will outperform (or underperform) only randomly. Its performance will be all a matter of luck. It doesn't matter whether it is one manager or 30 managers; a portfolio assembled without at least a fair assertion of skill is a closet index fund.

In the presence of skill, the opposite is true. One can hold a lot of managers yet not be a closet index fund if the managers are chosen skillfully.

There are other reasons not to hold a lot of managers, but by itself a large number or a diverse set of styles and approaches across your managers does not necessarily mean the investor holds a closet index fund.

Q: Should I hold active managers instead of index funds—so that someone can move me in and out of the market to protect the fund's value in volatile markets?

The ability to move a fund to a lower-beta position, either by going to cash or bonds or otherwise being more conservative, is an advantage frequently asserted by traditional active managers when they compare their approach to many more modern portfolio construction approaches—index funds, enhanced index funds, or any fund that has a fully invested requirement.

Adjusting one's market risk—beta risk—over time is an appropriate active management discipline. But it is subject to the usual active management caveats, plus one more. The usual caveats point out that, across all investors making market-timing bets, the return of the portfolios will be equal to the return of the market, less fees and

costs of course. Only those with special skill and insight into making market-timing bets will win with any degree of persistence; any other winners will be just lucky, and are as likely to be unlucky as lucky in the next period. Of course, we have to assume also that there is a bit of inefficiency related to the timing decision, inefficiency that the manager's skill can insightfully exploit.

The other usual caveat with regard to active decisions is that they add risk. It turns out that market timing adds a great deal of risk. Kritzman [2000, Chapter 5] does an excellent job of explaining this.

Here is the extra caveat that we don't need to bother with in discussing security selection approaches, and it's a subtle one: Market timing is what we call a *low-breadth* active management discipline. This means there are very few investment decisions made.

To be successful, even a skillful investor needs a lot of bets across which to diversify good and bad forecasts. Because of the low breadth, the bottom line is that special levels of skill are required to justify engaging in market-timing activity. Some investors no doubt have such skill, but it must of necessity be relatively rare. This is why modern market-timing approaches tend to attack a lot of asset classes and countries and to process a lot of independent valuation signals. They are building breadth, so that managers can achieve a higher information ratio with whatever skill level they do have.

So, to answer the question, if you believe that the manager in question has demonstrated skill in the specific art of making market-timing decisions of low breadth, then, yes, this is a good attribute and might justify hiring the manager. Market timing, if skillfully executed, is a legitimate active management discipline. Like any other active process, it doesn't have to be right every time, but over time it should be more right than wrong if it is skillfully done.

If you hire managers simply because they are capable of moving your portfolio out of equities, however, you may be just adding expected active risk rather than expected alpha—unless skill at market timing has been adequately demonstrated.

Q: I have a manager who disdains all talk of benchmarks, and denies that he should be measured against one. How does this connect to your benchmark-relative framework?

It really isn't *our* benchmark-relative framework. The term *benchmark* is in one aspect misleading. it is usually used simply to denote a yardstick for measuring performance, but the truth is that in investing, a benchmark signifies something far more important. It represents the return available on the asset class, as well as the return (before costs) that is available on the aggregation of all active managers in the asset class.

Another way to say this is that the benchmark represents the market-related risk component of a manager's returns, the systematic, undiversifiable, beta component. Unlike the residual, diversifiable, unsystematic, active component, the benchmark component is unconditionally rewarded by the market with an expected return related to its level of market risk, and a realized return that is completely determined by the movement of the market without any connection to manager skill. It is only the active component of returns that represents the manager's contribution to returns, and this contribution to returns is positive in expectation only if the manager has skill.

The long-only manager who resists being measured against benchmarks is surely willing to acknowledge that he or she is in fact exposed to the market. This statement is equivalent to saying that some component of their returns is systematic and some other component is pure active, and this in turn is equivalent to acknowledging that the manager is a relative-return manager appropriately evaluated by reference to a benchmark. There really is no escape; there are always two components of returns.

It can be hard to identify a manager's true benchmark—especially when the manager makes a lot of market-timing bets, switching between benchmark components. The historic returns will reveal unstable betas in a regression against the multiple components chosen. Yet there is, hidden in there somewhere, some normal portfolio representing the manager's typical level of market risk (or risks), of necessity.

In extreme cases, it is best for the analyst to simply default and use the closest general benchmark—the S&P 500, for example—as a good-enough proxy. But a better way is usually simply to understand the manager's articulation of his or her process. A value versus growth style rotator, for example, might acknowledge favoring value; the investor might estimate the normal portfolio as 60% large value and 40% large growth. (These style weights are really the same thing as multifactor betas.)

What if the manager argues that his or her style rotation isn't explicit or intentional, but simply *happens*, an accident of the manager's stock selections? We would respond that such a manager is taking many uncompensated style rotation bets, and would generate more consistent returns if he or she were to use risk controls to dampen them.

Of course, any time a manager shows such heavy

levels of style-timing activity, the logical question to ask is whether you believe the manager has skill at that activity. Market timing and style rotation are legitimate active management disciplines, but because of their low breadth they require greater skill than security selection approaches if they are to deliver similar information ratios.

It often seems that managers who most strongly resist the concept of benchmark-relative investing are managers who simply want to avoid accountability for producing pure alpha and thereby hope to maintain the investor's mandate for long periods without rigorous review.

Q: What makes a manager a good manager?

A manager's task is to generate pure active return without taking too much pure active risk. Active managers are too expensive to hire just for exposing the portfolio to market, or beta, risk. The investor hires them to add value over and above that provided by being invested in some predetermined mix of market exposures. These predetermined exposures come right out of the investor's strategic asset allocation policy; the sum of the normal portfolios of the managers in an investor's portfolio should look very much like the benchmark for the asset class (unless the investor—not the manager—wants to make an explicit bet otherwise).

So, that being said, here are some things to look for in identifying good managers.

Good managers "get it." Good managers know the difference between true alpha and market returns. And they know that active management is a zero-sum game; they know they have to be on the right side of the bell curve (by more than enough to cover fees and costs) if they are going to be successful at it. Igor Sikorsky, reflecting on his early years developing helicopters, is reported by his son to have said, "We were ignorant, and we were ignorant of the fact that we were ignorant. That is ignorance squared, and it can lead to disaster" (Sikorsky [2003, p. 138]).

You don't want managers that represent ignorance squared.

Most traditional managers (probably not too strong a characterization) still don't understand that they aren't hired to deliver persistent misfit risk relative to their benchmarks. Examples are found in the so-called core-plus strategies in fixed-income (typically involving high-yield debt persistently and intentionally held in a portfolio benchmarked to the Lehman Aggregate Bond Index, which does not include such debt); in the persistent underweight to Japan by active international managers, as discussed earlier; and in the observation that many style-biased managers (value managers, for example), still happily report their performance to clients with reference to the S&P 500 index if they beat it, instead of to the proper style index. (Or they report their performance relative to whichever index they did beat—stylized or unstylized—evidencing a willingness to prove competence through having beaten the lower of two benchmarks.)

But in smaller and more subtle ways, many traditional active managers simply have no appreciation of the notion of a normal portfolio to represent their natural market-risk exposure, nor do they have a full appreciation of how their style fits into the investor's benchmark. They don't even know what they don't know.

Use managers who get it.

Good managers have low or moderate active risk levels. Good managers take little active risk (for long-only funds) and carefully budget and manage it. The long-only constraint is a powerful force limiting the ability of even the smartest managers to harvest their skill (in the form of alpha) at high levels of active risk. Risk budgeting disciplines and other techniques that aid the portfolio manager, both in placing the right sizes on the intended bets that are being made and in minimizing uncompensated bets, are hallmarks of the modern portfolio management process and serve to keep active risk at low or moderate levels.

All the fundamental characteristics of managers that have traditionally informed the good manager selection process are in fact important. But they need to be understood by the investor not simply as screening or exclusionary devices—as is the norm. Rather, they need to be seen as potentially illuminating signals that may help the investor make the key evaluation: whether or not the manager is capable of making skillful forecasts and of building efficient portfolios from them. Stability of personnel, compensation of personnel, education of personnel, tools available, disciplines used, richness of signals, longevity of record, and so on are all useful in forming a skillful manager selection decision.

So a good manager's process will be focused on maximizing pure alpha, and on controlling risk—that is, on building a high information ratio. High information ratios mean consistency of performance over time. As Grinold and Kahn [2000, p. 148] demonstrate, the information ratio is proportional to the square root of breadth (the number of independent decisions to which the manager's skill is applied). This observation suggests—although it does not literally require—that ideally many people contribute to the process.

Thus, the good manager today is less likely to be one single smart person, or even a small group of smart people managed by a single smart person. Instead, we are finding that an increasingly effective model of good management is to finally move from the preindustrial period of building one handmade musket at a time, using individual, non-standard crafts, to Eli Whitney's industrial model, where smart people organize an entire process and build a lot of identical muskets (much more cheaply and efficiently) using interchangeable, mass-manufactured parts.

The analogy in investing is to create processes that concentrate and refine the essential inputs that generate the desired end-product—in this case, alpha. Today, the best processes are likely to be industrial-strength information management and evaluation systems, built and directed by teams of smart people looking for insights that are generalizable across large numbers of securities, rather than individuals or small teams of people trying to pick one security at a time.

Q: How do I forecast a candidate manager's expected alpha?

We suggest in Waring and Siegel [2003] that one could get started by using Grinold and Kahn's [2000] forecasting relation:

$$\text{Alpha} = \text{IC} \times \text{Volatility} \times \text{Score} \tag{3}$$

where IC (information coefficient) is a measure of your manager selection skill (the expected correlation between one's forecasts and the subsequent realizations of those forecasts); volatility is the standard deviation of the return being forecast (omega risk, volatility of the pure alpha); and score is the strength of the manager being evaluated, expressed in standard deviations above or below zero (a score of +2 or –2 would be considered very strong, and a score near zero would be weak).

But this is just a start, a quantitative take on what ultimately must be a qualitative decision. This question deserves a full development on its own. Suffice it for us to repeat here that an investor is a portfolio manager, managing a portfolio of managers who in turn manage portfolios of stocks. So an investor must have skill at selecting managers who have skill at picking stocks (or other securities). There are two levels of skill at work in this game.

There are no recipes for success. If there were, the process would no longer be a zero-sum game—but, mathematically, it has to be. While skill needs to be informed by as much fundamentally informative data as possible, at the end of the day skill can be acted on only through visceral judgments, judgments that over time show up as being somewhat more right than wrong.

This is the same requirement for skilled judgment in forecasting that we all require of active managers; after all, an investor is just an active manager of managers. Of course, the essential problem—generating alpha by forecasting returns and by constructing appropriate portfolios that reflect those forecasts—is the same at the investor level, looking at managers, as it is at the manager level, looking at stocks.

So anticipate more from us on this topic—but we won't offer a recipe or a formula for choosing the best manager. We will continue to emphasize that your thoughtful and educated inputs of judgment are needed for you to be a successful active investor.

ENDNOTES

[1]Part of the explanation lies in the benchmark, which until recently was not float-adjusted, so that illiquid shares (such as the cross-holdings common in Japan) were counted as part of the market when they should not have been. The float-adjusted version of the benchmark, had it existed then, would have had a smaller relative weight than did the non-float-adjusted version actually then in use. The relative size of the U.S. managers' underweight positions would have been smaller by reference to a float-adjusted benchmark.

[2]Grinold and Kahn [2000] identify the portion of the beta that is different from 1.0 as a positive or negative "active beta." This is a very useful way to think about it, and has other benefits in facilitating calculations. A variable active beta is used for efforts to generate returns from tactical asset allocation processes. We're presuming a static value of this active beta when we discuss typical portfolios, rather than the variable value associated with market timing strategies.

Remember that we are dealing with the persistent beta of a manager, not a temporary variation for market-timing purposes. Such temporary moves create another type of alpha, and a parallel risk, but these are almost identical to our alpha and omega terms. See Grinold and Kahn [2000], Appendix to Chapter 4.

[3]Alpha can come from either stock selection or beta timing. Some managers like to have the flexibility of making market-timing decisions by occasionally going into cash, or by rotating among styles, and some sponsors are agreeable to that.

[4]As long as the underlying distributions have a finite variance.

[5]There can be some very surprising occurrences buried in the period-by-period underlying distributions of returns. See Taleb [2001]. Our advice is to minimize the potential for

damage from these events by building very "high-breadth" strategies, with many fine-grained, low-correlation signals across many securities.

[6]For generally similar results and an updated and much clearer interpretation, see Ibbotson and Kaplan [2000] and Surz, Stevens, and Wimer [1999].

[7]Any standard regression package will readily calculate the precise t-statistic for the analyst. More precisely, it is the arithmetic average residual return times the square root of the number of periods (usually we would use monthly data), all divided by the standard deviation of those returns—the standard error—stated in the same periodicity. Our quick method in the text works with an annual periodicity, as that is how information ratios are usually stated.

[8]We're only scratching the surface of how to determine the statistical significance of manager histories. We've used a two-tailed test framework; some would move to a one-tailed test. There are issues about using data from different regimes of management (and other issues) that need to be thoughtfully addressed. Our point is that we need more rigorous work here than currently available.

REFERENCES

Brinson, Gary P., L. Randolph Hood, Jr., and Gilbert L. Beebower. "Determinants of Portfolio Performance." *Financial Analysts Journal*, July/August 1986, pp. 39-44.

Brinson, Gary P., Brian D. Singer, and Gilbert L. Beebower. "Determinants of Portfolio Performance II: An Update." *Financial Analysts Journal*, May/June 1991.

Grinold, Richard C., and Ronald N. Kahn. *Active Portfolio Management*, 2nd ed. New York: McGraw-Hill, 2000.

Ibbotson, Roger G., and Paul D. Kaplan. "Does Asset Allocation Policy Explain 40%, 90% or 100% of Performance?" *Financial Analysts Journal*, January/February 2000.

Kritzman, Mark P. "Half Stocks All the Time or All Stocks Half the Time." In *Puzzles of Finance*. New York: Wiley, 2000.

Sikorsky, Sergei. "Igor Sikorsky." *AOPA Pilot,* February 2003, p. 138.

Surz, Ronald J., Dale Stevens, and Mark Wimer. "The Importance of Investment Policy." *The Journal of Investing,* Winter 1999.

Taleb, Nicholas Nassim. *Fooled by Randomness*. New York: Texere, 2001.

Waring, M. Barton. "The Myth of the Absolute Return Investor." Forthcoming, *Investment Insights*, Barclays Global Investors, San Francisco, 2004.

Waring, M. Barton, and Laurence B. Siegel. "The Dimensions of Active Management." *The Journal of Portfolio Management,* Spring 2003a.

——. "Understanding Active Management." *Investment Insights*, Barclays Global Investors, San Francisco, 2003b.

To order reprints of this article, please contact Ajani Malik at amalik@iijournals.com or 212-224-3205.