InvestmentInsights

The Investment Research Journal from BLACKROCK

4.03

APRIL 2003
VOLUME 6 ISSUE 1

UNDERSTANDING ACTIVE MANAGEMENT

PART ONE: *The Dimensions of Active Management*

PART TWO: *Debunking Some Myths of Active Management*

by BARTON WARING *and* LAURENCE SIEGEL

Investment Insights

Barclays Global Investors is now BlackRock.

BlackRock is pleased to provide these research insights to you. They were originally produced by Barclays Global Investors prior to our merger in 2009.

BLACKROCK

EXECUTIVE EDITORS

Ronald N. Kahn
415 670 2266 phone
415 618 1574 facsimile
ron.kahn@blackrock.com

Russ Koesterich
415 670 2576 phone
415 618 1875 facsimile
russ.koesterich@blackrock.com

EDITOR

Marcia Roitberg
850 893 8586 phone
415 618 1455 facsimile
marcia.roitberg@blackrock.com

CONTRIBUTING WRITERS

Barton Waring
Head of Client Advisory Group

Barton Waring joined BGI in 1995 and has run our Client Advisory Group since 1996. He has an extensive background in classical investment strategy and policy, asset allocation, and quantitative investment management, gained from his prior work with large defined benefit and defined contribution plans in both consulting and advisory roles. Barton received his BS degree in economics from the University of Oregon, his JD degree from Lewis and Clark, with honors, and his MBA degree in finance from Yale University.

Laurence Siegel
Director, Investment Policy Research, Ford Foundation

Guest co-author Laurence Siegel is director of investment policy research at the Ford Foundation in New York, where he has worked since 1994. Previously, he was a managing director of Ibbotson Associates, an investment consulting firm that he helped to establish in 1979. Larry is chairman of the investment committee of the Trust for Civil Society in Central and Eastern Europe and serves on the investment committee of the NAACP Legal Defense Fund. He also assists the investment committees of numerous other nonprofit organizations. He serves on the editorial boards of *The Journal of Portfolio Management* and *The Journal of Investing*, and was the founding editor of *Investment Policy* magazine. Larry received his BA in urban studies and his MBA in finance from the University of Chicago.

The authors gratefully acknowledge the input and assistance provided by Peter Bernstein, Bruce Goddard, Blake Grossman, Ron Kahn, Ken Kroner, Gerry Rocchi, Clinton Stevenson, and Linda Strumpf. Many other very helpful comments were received; we apologize in advance to those not listed. "The Dimensions of Active Management" is scheduled to be published in *The Journal of Portfolio Management*.

UNDERSTANDING ACTIVE MANAGEMENT

Table of Contents

IN THIS ISSUE

EXECUTIVE SUMMARY

This paper takes a look at active management. The first part explores the role that active managers play in a portfolio and the "dimensions" in which they vary. In the second part we address some of the myths about the nature of active management. Active management is a good thing, used wisely. These myths get in the way of such wise use.

Plan sponsors hire active managers to generate "pure" active return, or "pure" alpha (the portion of the return that is not explained by market or beta risk exposures of the portfolio), knowing that doing so will generate pure active risk. Building a portfolio of managers involves optimizing the trade-off between pure active return and active risk.

Current practice typically draws only a fuzzy distinction between a manager's pure alpha and the market exposures (including style and other systematic risk factor exposures) delivered by that manager. By making this distinction clear, and by quantifying it, one can greatly improve the payoff to the decision to hire active managers.

Existing practice for building portfolios of managers, such as historical performance comparisons, style boxes, and heavy reliance on traditional active management are not clearly focused on adding value to the portfolio. The process of building portfolios of managers can be better attuned to maximizing active return while controlling active risk. We suggest the following rules of thumb:

1. Be disciplined in forming expected alphas, and in giving the greater weights to managers with higher ratios of expected alpha over active risk squared. Be rigorous when examining historic alphas.

2. The portion of the portfolio that moderates its overall risk will consist of some combination of "good" risk-controlled active funds (enhanced index funds) and traditional index funds. Together, these should represent one-third to two-thirds of the total fund). Risk-controlled active funds will be more heavily weighted than equally skillful traditional active funds at most risk levels chosen by institutional investors. Among traditional active managers, prefer skillful lower active risk managers to higher risk, concentrated managers. Bias toward diversified portfolios. "Good" market-neutral long-short funds will receive a substantial weight for investors not limited to long-only managers.

3. Set the overall active risk at a comfortable level, your "risk budget." For US equities, a typical overall active risk level is 1.5 to at most 2%, with the very largest investors preferring even less active risk (between 0.75 and 1.25%).

4. Keep a careful eye on misfit risk, trying to minimize it while still maximizing expected alpha. If you don't use an optimizer designed especially for this purpose, you'll have to use a style map or "effective asset mix" table as a supporting tool.

In short, for the plan sponsor confident of his or her skill at selecting managers, the question is not whether to use active management, but how. The plan sponsor's task is made harder by the fact that too often active managers "spin" their craft, creating a number of myths about active management in the process. In the second part of this issue we address some of those misconceptions.

PART ONE:
The Dimensions of Active Management

Every reader of this paper either employs active managers or is one. But what exactly are active managers? What is their role in the investor's portfolio? What *should* their role in the portfolio be? What are the real "dimensions" in which active managers vary? Are current practices for building portfolios of managers the best that can be designed, or is there a better way?

Most investors know *something* about the answers to the questions we posed, as common sense goes a long way here. Active managers select securities. They invest the portfolio, and it is hoped that they will add "alpha," an incremental return over their benchmark. Investors, for their part, work hard to hire the best managers they can identify. Our ambition here is to support that effort.

To that end, we hope to convey two fundamental insights. The first is that one hires active managers to generate what we'll call "pure" active return, or "pure" alpha, knowing that in their efforts to do so they will generate pure active risk. (We use the unconventional modifier "pure" for reasons discussed below). We also note that building a portfolio of managers is like building a portfolio of anything: It's an optimization problem. And if pure active return and risk are the key dimensions describing active managers, then building a portfolio of managers involves optimizing the tradeoff between these two dimensions.[1]

Second, pure active return, subject to a penalty for pure active risk, is what managers should be hired to deliver. And it is what investors should pay active fees for. This is in contrast to current practice, which typically draws only a fuzzy distinction between a manager's pure alpha and the market exposures (including style and other systematic risk factor exposures) delivered by that manager. By making this distinction clear, and by quantifying it, one can greatly improve the payoff to the decision to hire active managers.

Holding active managers, then, is like any other proposition in finance, the balancing of the hope for gain against the concern about risk. Existing practice does not do a good enough job of managing that tension, and doesn't even explicitly focus on it, but there are new technologies and ways of thinking that do.

I. What is active management?

You already know that active managers try to beat an asset-class or style benchmark, using securities held in other than benchmark weights. Can this inherently be successful on average? What can such managers do for your portfolio? How should you choose them? How should you weight them in your portfolio? Does more active risk mean more active return? If so, under what circumstances?

Active managers are *forecasters.* They use the information available to them, and whatever their native talents are, to make stock-by-stock (or factor-by-factor, or market-by-market) forecasts of pure active return based on information that they believe is not yet impounded in the price, and then they translate these forecasts into portfolios.[2]

Well, not exactly. The traditional manager typically doesn't make explicit forecasts of stock-specific returns, nor does he optimize those forecasts to construct their portfolios. (These practices *are* followed by the best of the risk-controlled active managers.) Instead, traditional managers may establish "price targets" or express forecasts in some other forecasting space that does not translate directly to expected alphas. At the end of the day they do hold a portfolio of securities, often more or less equal-weighted, that they hope will beat their benchmark. This portfolio is informed more by traditional research, intuition and experience than by optimization and risk control considerations.

However, whether the manager knows it or not, stock-specific alpha forecasts are always there, implied by the bets in the portfolio. They can be backed out of the portfolio holdings, using "reverse optimization."[3]

If you can't forecast better than the average market participant, you shouldn't be an active manager. And if you can make superior forecasts, but don't agree with the forecasts implied by your holdings, then you should have different holdings, consistent with forecasts with which you do agree.

Likewise, the *investor* makes alpha forecasts for the *managers* it holds. These forecasts are rarely explicit in current practice, but they are also always implicitly there. They can also be backed out through reverse optimization—and they can be embarrassingly large.[4] Investors should only hire active managers if they have skillfully formed positive alpha forecasts for the managers. And if they have these forecasts they should use them. We will get to this in more detail later.

Superior forecasting, then, is what generates the positive expected alpha that we hope to earn from hiring active managers. Forecasting is important at two levels, by the managers looking at stocks, and by an investor looking at managers.

MARKET RISK AND ACTIVE RISK, REWARDED AND UNREWARDED RISK

To begin to identify the dimensions of active management, we must first break up total risk into its "gross" component parts. Following William F. Sharpe's most famous work, the total risk of any investment can be broken into: (1) *market risk,* also called policy, systematic, undiversifiable, or beta risk; and (2) *pure active risk,* also called specific, unsystematic, idiosyncratic, diversifiable, or alpha risk (risk goes by many names).[5] Policy risk is usually managed by investors through their adoption of a strategic asset allocation (SAA) policy, with its attendant asset-class benchmarks, all of which are fully diversified and which therefore carry only market risk by definition. Active risk comes into the investor portfolio through the use of active managers who, while holding some components of the policy benchmarks, try to beat them.

Finance academics have focused much of their research in the last few decades on market risk, observing that if markets are efficient then "you can't beat the market." In practice, however, investors don't seem to completely believe them. Investors still routinely hire active managers, whose returns can differ widely from benchmark returns. In fact, the all-passive institutional portfolio is a rarity, the academic insight being honored solely in the breach. So active risk is in fact taken, in the search for pure active return, and this tradeoff needs to be managed as well as the tradeoff between market return and risk that is at the heart of asset allocation policy.

WHERE'S THE BEEF?

Does this mean that investors should only index, shunning active managers entirely? Not necessarily, although it may for some. The market efficiency story is incomplete: Under a couple of fairly easily satisfied conditions, you *can* beat the market.

Note carefully: As long as a market is not completely efficient (and we believe that none are) and as long as there are native differences in human intelligence and skill levels (of course there are), *some managers will outperform through real skill, not just by virtue of random variation.*

It is important to note that while market inefficiency is a necessary *condition for "good" active managers to exist, it is not a* sufficient *condition—skill is also required.*

Capital markets must function so that the expected return on the overall equity market must be higher than that on fixed-income investments.[6] As a result, fully diversified market risk (such as one takes by buying an index fund) must be rewarded. Pure active risk, in contrast, cannot be rewarded *on average,* since active managers are betting against each other in a zero-sum game.[7] Particular active managers, of course, will beat the market by a large margin—but they are doing so at the expense of the others, either through luck or through special skill. But active managers (and other market participants) in aggregate *are* the market, so they must again in the aggregate earn the return on the market—minus fees, transaction costs and other expenses, which can be substantial.[8]

Under these conditions, the notion that there can be an "expected alpha" or pure active return makes perfect sense. It is important to note that while market inefficiency is a *necessary* condition for "good" active managers to exist, it is not a *sufficient* condition—skill is also required.

But no manager will ever have so much skill as to be able to manage the portfolio without adding some active risk. We measure skill with an "information coefficient" (IC), a correlation coefficient of forecasts with realizations. Forecasting skill will seldom carry a very high information coefficient, and so there will always be risk left over, and a chance of underperformance even in the hands of the "quite good" manager. This risk represents variability in active returns not explained by skill but rather by luck. The luck component

Definitions

Alpha, "pure alpha," "pure active return," "selection return," "specific return": The portion of the return that is not explained by the market risk exposures of the portfolio, the real return generated over and above a manager's custom benchmark or normal portfolio.

"Pure active risk," "omega": The standard deviation or other measure of volatility of the alpha.

Active return, "simple active return": The sum of the alpha plus any additional return resulting from the manager's holding less (or more) than a "beta equals one" exposure to the benchmark.

Active risk, "simple active risk": The standard deviation or other measure of volatility of the active return.

Utility: A form of risk-adjusted return. Expected returns are decreased by a penalty proportional to the risk (variance) of the return stream.

Traditional active management: The typical form of active management most commonly practiced, characterized often by heavy use of security analysis, often by equal-weighted portfolios. Active risk levels run from around 4% up to as high as 25%, with the norm being around 5 to 6%.

Risk-controlled active managers: Heavy use of technology to gather data and evaluate insights about individual securities, and heavy use of advanced processes to form these insights into optimal portfolios having minimal uncompensated risks. Tend to have low active risk levels, 1 to 3% or so.

of risk, period-by-period, is nearly as large for the skillful as for the unskillful manager. Over long time spans, of course, positive alpha is much more likely to accumulate for the skillful.[9]

So the bad news is that active management overall is in fact a zero-sum game. But the good news is that for the skillful investor, there is an opportunity to add value to the portfolio over and above market returns. This really is good news, especially if you contrast it with the returns to be expected from your benchmark and SAA decisions. You can't influence or control the return of your SAA policy: The market is going to do what the market is going to do. Other than making a risk level decision—to be more or less aggressive in your SAA policy (either in an asset-only decision or preferably with respect to your liability)—you're just a passenger. But if you have skill at security selection (or market timing or sector rotation, any active process), you have some *control* over returns, and this will add value, pure alpha, over and above the return of the SAA policy. The search for such alpha is, arguably, the investor's highest calling.

Remember, we're not saying it is easy. It's not. But it can be done.

II. Defining "pure" active return and risk: What should managers be hired to deliver?

We agree on the separation of market or policy risk and pure active risk, in principle. Let's go one step further and apply it in practice, separating these "gross" components of market risk (and return), and pure active risk (and return) in a real world context.

The returns that usually pass for alphas—the simple differences between the benchmark return and the manager's return—are properly known just as "active returns" (without the "pure"). But these "simple" active returns might have more or less exposure to market risk than the amount implicit in the investor's benchmark.

Market risk can be measured by a CAPM single-factor beta, or, more usefully, in a multifactor manner. The most intuitive of these multifactor approaches measures market risk in terms of *style factors*—the familiar large-capitalization, small-capitalization, value, and growth categories.[10] Other, often more complex, factor models have been identified that more completely explain market risks, but some of these are hard to describe in plain English so we'll stick with style factors in this discussion. Style factors also have the benefit of convenient investability through low-cost style index funds.

The manager's returns, then, can be explained in terms of the exposures, or betas, to an intuitive series of style factors that express the manager's return from market risks; plus a pure alpha, or the return generated over and above market returns. Simple active return and pure active return are only the same if the manager's factor exposures are weighted the same as they are in the benchmark. Since the benchmarks flow from the SAA decision, managers can and should be chosen with factor weights such that, at least in sum across the managers, they are consistent with the benchmark's factor weights.

So once a set of market factor weights—a custom benchmark such as "large-cap growth" or "80% large-cap value and 20% large-cap growth"—is set for a given manager, the proper objective of the manager is simply to beat that customized benchmark. That's why you hire active managers. You hire them to give you the levels of market risk exposure that you expect from them and which you assigned to them through their customized benchmark, and to beat that benchmark. If a manager does anything else, with or without the knowledge of the investor, it is stealthily changing an important aspect of the investor's SAA policy. The common term for this is "misfit" risk, but it's really the risk that the particular mix of benchmarks representing one's SAA policy is not being delivered.

In other words, the valuable and important return added by a manager isn't the total return that he or she delivers, but only that part of the return that is *beyond* what could be delivered through a set of index funds reflecting the manager's persistent style biases, its market risk exposures. This unique contribution of the manager to the return is what we're calling *pure* active return, or *pure* alpha. We know we're being redundant to say "pure" alpha, but we're trying to call attention to this precise definition, one that doesn't contain a market (or style) risk component, for the reasons stated in the prior section. It is also what Sharpe meant by his term "selection return" in his work on style analysis, and what those familiar with the very detailed factor models of BARRA know as "specific return."

Realized pure alpha is easily separated from market risk factors and measured by regression analysis.[11] The regression determines the effective

style weights of the manager, or mix of style benchmarks that has the "best fit" to the manager's actual returns. The pure alpha is then the residual, the manager's actual return in excess of the return on this amount of market risk. On a forward-looking basis, we *assign* to a manager his or her customized benchmark or "normal portfolio," capturing the style and other market risk exposures that will best describe that manager's neutral, or "home," position. It's easy for index funds or risk-controlled active managers. For others, this customized benchmark might be informed by the historic regression and by any other information that is useful to characterizing the normal style biases of the manager. Even a TAA manager or a style rotator has such a "home" position.

The view of market risk that we've been describing, by the way, is continuous and scalar—that is, a manager can have any amount of exposure to a single benchmark or to multiple style (or any other factor) benchmarks. The market exposure or style weight, at its essence, is just a *beta,* after all. And betas are a good way to determine or describe the level of exposure to any market risk.[12]

One additional idea: The investor may or may not be able to collect a portfolio of managers whose normal portfolios, in the aggregate, look like the benchmark. The misfit risk of a single manager goes away if it is cancelled by misfit risk of an opposing character from another manager (a growth manager is offset by a value manager, a large-cap manager with some small-cap exposure is offset by a small-cap manager with some large-cap exposure, etc.). But such perfect offsets of style and other factors are not often the case. So when we're optimizing a manager structure, as we do below, we'll be optimizing pure active return against *total* active risk—and on the risk side we'll control not only the pure active risk added up from each manager, but also the *net* misfit risk taken across all managers, calculated properly using scalar values for the managers' exposures to all the market risk factors we are tracking. This net misfit risk is a part of the active risk investors actually face.

Too often, performance evaluation practices confuse the benchmark return and the pure alpha, apportioning credit and blame incorrectly.

THE INFORMATION RATIO SUMMARIZES THE KEY DIMENSIONS OF ACTIVE MANAGEMENT

The standard deviation of the period-by-period pure alphas may be thought of as the *pure active risk,* representing the tracking error to the manager's customized benchmark. These two parameters,

- Pure active return, or pure alpha, α
- Pure active risk, ω (known as the Greek letter "omega")

can be combined to arrive at a single measure of manager achievement (either historical or expected), the:

- Pure information ratio, $IR = \frac{\alpha}{\omega}$

representing the amount of pure active return delivered (or expected) per unit of pure active risk taken (or expected) by an individual manager, relative to its customized benchmark.[13]

Across the portfolio of managers held by the investor, the denominator would be the aggregation of the ω terms plus any net misfit risk remaining across the group of managers. We indicate this "simple" active risk as σ_A. One wishes for the misfit component to be zero across all managers, but in practice it is difficult to make every last bit of misfit risk go away.

PURE ACTIVE RETURN AND RISK AS A LEVEL PLAYING FIELD FOR EVALUATING MANAGERS

Now that we've defined pure active return and risk, we can use these measures (and particularly the ratio of return to risk, the pure information ratio) as well as misfit risk, to compare any manager with any other—across asset classes, styles, and risk levels—creating a level playing field for all managers.

Even more importantly, we can use these measures to properly separate investment results that are the investor's responsibility from those that are created by the manager. The returns delivered by the capital markets on the particular mix of styles that constitute the manager's custom benchmark are the responsibility of the investor who selected the manager, if only because the investor is the only party in a position to control the market risk exposures across his or her whole portfolio of managers.

Too often, performance evaluation practices confuse the benchmark return and the pure alpha, apportioning credit and blame incorrectly. Even the smartest and most well intentioned investors are sorely tempted to blame the active manager, rather than themselves, when the manager's asset class delivers a poor policy return (no matter what pure alpha the manager achieved). With the pure active return and risk clearly defined and calculated, these errors need no longer occur. As a common example, think of the value manager that boasts of beating the S&P even when he fails to beat the value benchmark. Which one should he really be held against? If a manager persistently chooses to exercise his expertise in one domain of market risk such as "deep value," isn't that the domain against which his value-added should be measured by clear-eyed investors?

III. What are active managers for? Maximizing active return minus a penalty for active risk (maximizing expected active utility*)*

So why hire active managers? They provide the possibility of adding pure active return, of course, but they also add active risk.

The risk added by active management is, in and of itself, undesirable, so a manager has to do more than just have an expected alpha that is positive. He or she must add enough to more than compensate for the added risk. Active managers are there to add *utility,* not just expected return.

The methodology for determining the utility of active decisions is parallel to that for SAA decisions. In general, the expected active utility (or usefulness, or desirability) of a portfolio of active managers is equal to its expected alpha, minus a risk penalty for active risk:

$$E(U_\alpha) = E(\alpha_p) - \lambda_\alpha \, E(\sigma_A^2) \qquad (1)$$

where

$E(U_\alpha)$ is the expected utility of active management in the portfolio;

$E(\alpha_p)$ is the expected alpha on portfolio p;

λ_α is the active risk aversion parameter for the investor (or the rate at which risk is translated into a negative return, or disutility);

and $E(\sigma_A^2)$ is the simple active risk (expected variance) of portfolio *p* (including both omega risk from each manager and the net misfit risk across the managers).

How do we figure out whether a portfolio that includes active managers provides incremental utility over the benchmark; that is, whether the combination of managers selected adds enough expected alpha to justify the extra risk? How, taking this one step further, does one *maximize* expected utility? Through optimization. Although Harry Markowitz created this tool a half century ago, and although it's the dominant practice for building efficient portfolios of asset classes and, increasingly over the last 15 years, of securities within an asset class,[14] optimization has just begun to be used for building efficient portfolios of *managers* within the investor's portfolio over the last six or seven years.

But why? Building a portfolio of managers is like building a portfolio of anything: It's all about balancing risk and return, trying to find the best tradeoff. Optimization is the technology that explicitly calculates these tradeoffs in search of the highest utility portfolio (of anything) for a given investor.

So increasingly, and as a result of thinking about how managers interrelate in the total portfolio's utility function through optimization, investors are beginning to see their task as one of building *portfolios* of managers, in place of the conventional approach of looking at each manager in isolation. To accomplish that task, investors hiring managers are increasingly using optimization as a portfolio-building tool. Specifically, one must optimize on the managers' pure active return and risk—the real dimensions of active management—while dealing with the other (policy and misfit) return and risk components in a sensible way such as we describe.

IV. Building portfolios of managers: The current common sense

Before going more deeply into our prescriptions for better approaches based on optimization, let's hold up to the mirror a couple of the chief features of current practice for building portfolios of managers. These features represent a good portion of the "common sense" that guides us in manager structure today. Albert Einstein said, "Common sense is the collection of prejudices acquired by age eighteen."[15] But is this good enough?

PERFORMANCE "HORSE RACES": MINING THE HISTORICAL DATA

The most widespread current practice for building portfolios of managers is to assume, tacitly if not explicitly, that managers will continue to earn whatever alpha they've been earning in the past. Great effort is put into identifying and sorting candidate managers based on their historic active returns, despite massive evidence that past performance is at the very best a weak predictor of future results, and despite the muddling of these returns with market components.[16]

That's probably because analysis of past performance is something that investors can do with a sense of objectivity and confidence. It's hard to argue with actual historical returns. They provide a feeling that one is dealing with something real and concrete.

Except when making forecasts.

While past performance should not be ignored—it is one of many factors that should be considered when evaluating a manager—the investor should remember that performance track records do not, by themselves, distinguish between luck and skill. Two managers, one lucky (but producing random variation around the properly style-adjusted

benchmark) and the other truly skillful, can have the same track record. This can even be true over fairly long time horizons.

We'd like to give credence only to "statistically significant" performance data, right? Statisticians use the t-statistic to test whether data are statistically significant—more precisely, to test whether one can conclude with reasonable confidence that any given alpha, positive or negative, was achieved through skill rather than luck. There is general agreement (with only minor quibbles) that if a manager's alpha over the period studied is more than two standard deviations away from that of the benchmark (that is, if its t-statistic is greater than two), we can say the manager has a "statistically significant" alpha. This simply means that there is a very high probability (roughly 95% if alphas are normally distributed) that the manager's alpha really is different from zero. It is evidence of skill rather than luck.[17]

Here is why we digress on statistics: If a manager's historic alpha is not statistically significant in this way (high t-statistic), it makes no sense to even consider whether the manager's historical alpha will repeat going forward. Since a lack of statistical significance says that we don't even understand whether the underlying non-random component of the alpha was different from zero or not, the data are meaningless noise and should not be used. That's what "not statistically significant" means. The investor will have to make his or her alpha forecast for that manager based on a more qualitative or fundamental type of analysis, which of course can be done—it is just hard.

On the other hand, if a performance history does display statistical significance, then it is fair to include it among one's other inputs when evaluating a manager; but still one doesn't just extrapolate it into the future without thought. A high t-stat doesn't by itself prove skill, but a low t-stat should be interpreted as meaning that the performance record, at least, shows no evidence of skill.

If a performance history does display statistical significance, then it is fair to include it among one's other inputs when evaluating a manager; but still one doesn't just extrapolate it into the future without thought.

If we cannot usually rely on past performance to select active managers, then how can we select them at all? We don't have a recipe, and we know that there aren't any recipes—if there were one, everyone would be following it and of course then it wouldn't work. Each investor will have to develop his or her own methodology for forecasting manager alphas, but the key ingredient is the tough one—one has to have great insight and ability. It's no different from how excellent active managers pick stocks: they use a tremendous amount of research into fundamentals, and at the end of the day they make a judgment call informed by their trained instincts. In "Those pesky expected alphas," below, we touch on other issues related to forecasting alphas for managers.

Unfortunately, significant t-statistics on managers are quite rare—although such rarity is exactly what theory predicts.[18] We almost never see t-statistics used, because using them would require the investor to throw out and not use the historic alphas in most all cases. Ignoring historical returns seems at first blush like a peculiar practice, but isn't it right to reject historical data when they contain no useful information?

STYLE BOXES AND STYLE MAPS

A practice that has been dominant over the past couple of decades or so, and that is less than ideal, is to divide the equity universe into "style boxes"—where the boxes are large value, large growth, small value, small growth, and perhaps mid-cap and/or core categories—and then to "staff" each style box with managers. This practice seems useful for breaking up the arduous manager-selection task into manageable pieces. The idea is to assign managers to each of these buckets, and that once the buckets are all staffed, the plan is well structured.

Since a cap-weighted combination of all the style boxes gives you back the market portfolio, it might appear that there's nothing really wrong with style boxes *per se.* Unfortunately, value managers, growth managers, large-cap managers, and small-cap managers don't all come prepackaged cleanly in these boxes: Some are stronger ("deep value") or weaker (growth at a reasonable price) than others, the amount of "valueness" or "growthiness" not always coming in units of one. Many large capitalization active managers are equal-weighted, and as a result end up having some amount of "small-cap" exposure and less than full "large-cap" exposure (and vice versa for most small-cap managers, who seem to hedge a bit toward large stocks). Anyone who has conducted style analysis on managers knows that it is common for, say, a growth manager to be characterized as 70% in large growth, 15% in large value, 10% in small growth, and 5% in small value.[19] Particularly for traditional active managers, style exposures usually come from a continuous spectrum, not from an "all or nothing" bucket. (The percentage weights, by the way, are just a convenient way to express the betas of the manager relative to the style factors.)

Rather than simply using boxes to represent styles, some investors plot their managers—and their total portfolio of managers—on a style map. Managers who don't fit neatly into a style box can be hired in such an improved framework, as it is usually calculated in a manner that accommodates continuous, scalar values. The investor's total portfolio style map displays the net misfit of the portfolio. One can also use it to simulate changes in managers, or in manager weights, to see the effect on misfit risk and other portfolio characteristics.

Style maps are much better than style boxes, since they recognize the continuous nature of styles and the importance of managing *net* misfit risk, across all managers. If an investor isn't going to formally optimize as we suggest in this paper, style maps and tables of the "effective asset mix" are the next best thing.

The fundamental problem with using either style boxes or style maps to organize the manager structure effort is that neither one requires *the investor to deal "face to face" with the managers' expected alphas.*

But the fundamental problem with using either style boxes or style maps to organize the manager structure effort is that neither one *requires* the investor to deal "face to face" with the managers' expected alphas. Used in the conventional way, style boxes and maps can subtly invite the investor to just "fill them up," without critically examining his or her skill at assessing expected manager

alpha or the relative expected alphas of the candidate managers. For investors that aren't strongly disciplined, style boxes and maps invite taking on active risk without a real expectation of pure active return.

HEAVY RELIANCE ON TRADITIONAL ACTIVE MANAGEMENT: "I WANT MORE JUICE!"

As we've already noted, many of the investors who rely heavily on recent past performance and on style boxes as aids in selecting managers are also strongly biased toward traditional active management, with its high degree of active risk. They do so because they equate greater active risk with greater active return. Such investors often express disdain for risk-controlled active strategies, and a preference for concentrated, high-risk active strategies. The refrain, when faced with a low-risk active strategy, is "But I want more juice!"

Such investors make two mistakes. First, having learned that "risk is related to return," they fail to distinguish between policy risk (which is in fact associated with a higher expected return) and active risk, which is only rewarded conditionally on skill. More juice, by itself, isn't going to accomplish anything.

Second, such investors don't realize that their *optimal* portfolio of managers, at whatever active risk level with which they are comfortable, will be constructed from good low-risk active managers in preference to higher-risk managers, for two independent reasons. The first of these is that the mathematics of optimization dictate that an investor's manager allocation should be dominated by managers that have a combination of the highest forecast information ratios and the *lowest levels of active risk.* Next, due to the no-short-selling constraint, high-active-risk portfolios tend to have *lower information ratios* than low-active-risk portfolios at the same level of manager skill. Feed this lower information ratio back into the first reason, and you see a loop that requires unimaginably high skill levels to justify giving large allocation to concentrated managers.

We conclude that sophisticated investors don't really want "more juice"—what they really want, or should want, is higher alphas, and less risk. We cover these concepts more thoroughly in the next section.

CURRENT PRACTICES: CONCLUSION

Although in practice they are informed by the best efforts of investors, the conventional dimensions within which active manager decisions are framed—historical alphas, style boxes, the search for "more juice"—are not clearly focused on the investor's goal of adding value, or utility, to the portfolio. Practices for building portfolios of managers could be better attuned to the real dimensions of active management, maximizing expected alpha while controlling active risk. Can we suggest practices that point in that direction?

V. Building portfolios of managers: Tomorrow's best practices

FROM SECURITIES TO MANAGERS: TREAT EACH MANAGER LIKE A STOCK

Active return/active risk optimization may not be what investors first think of when they set out to build a portfolio of managers, but it has been growing in use for building active portfolios of *stocks* for well over 15 years. It is becoming common practice in that context, having many advantages over traditional methods of constructing active stock portfolios. What we are suggesting now, is simply to treat each manager as one would

Exhibit 1

EFFICIENT FRONTIER OF ACTIVE MANAGERS

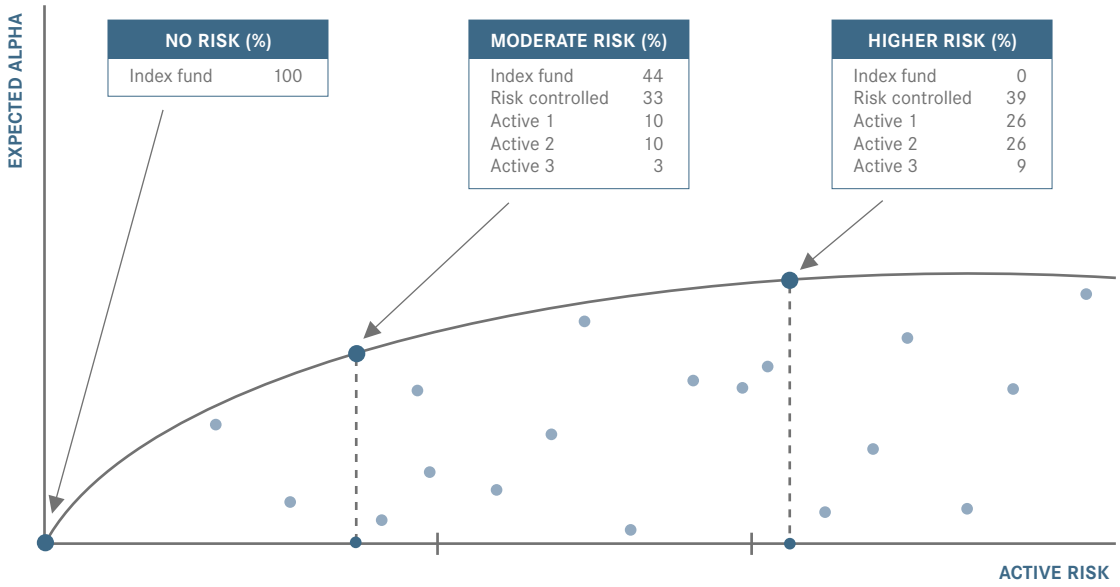

a stock, giving the investor the same powerful tools for managing its portfolio of managers that the best managers have for managing their portfolios of stocks. Specifically, one must optimize the managers' pure active return and risk—the real dimensions of active management—while dealing with the other (policy and misfit) return and risk components in a sensible way.

We set up the utility function for this optimization earlier. Now, let's implement it using managers as the "securities" or "stocks" across which we will optimize. From this, we can construct an active efficient frontier across managers, an example of which is shown in Exhibit 1.

Thus, following Waring et al. (2000), who call their method "manager structure optimization" (MSO) to echo Markowitz's MVO, or mean-variance optimization, we can implement this utility function on a practical basis.[20] We need only to form the investor's estimate of

1. The expected pure alpha and pure active risk of every manager.
2. The expected correlation of every manager's alpha with that of every other manager.
3. The market factor betas that fairly characterize each manager's normal portfolio.
4. The expected return and risk of each market factor used, and the expected correlation of each factor with every other factor.

Whew! This looks more daunting than the usual list of MVO inputs. However, it's easier than it looks. The correlations of the pure alphas are often simply estimated as zeroes. This is for two reasons. First, we have taken the market-related risk factors out by calculating a regression alpha, the pure active return that is uncorrelated with the market factors. While this doesn't necessarily mean that all common factors have been removed from the alphas such that they are necessarily uncorrelated with each other, they will by virtue of the process tend to have low correlations. Second, we observe in alpha histories that the correlations are in fact low, generally running between –0.2 and +0.2. Given the sample error

in those observations, one fair solution is to regard all the alpha correlations as zero on a forward-looking basis unless there is a specific reason to expect some pair to be different.[21]

The third and fourth items are the descriptions of the managers' customized benchmarks, expressed in terms of the market risk factors that are relevant to the manager structure question at hand, and the capital market assumptions related to those market risk factors. Analysts are used to providing capital market assumptions for SAA work; the assumptions used for manager structure optimization should be consistent with the SAA assumptions (but may need to be more detailed).[22]

The difficult item is the first, and in particular the return component of it, expected alpha. We'll come back to this in a minute.

Notice in Exhibit 1 that at zero active risk only an index fund is held, but that as risk tolerance increases—going to the right on the efficient frontier—the proportions held in risk-controlled active and traditional active increase. Another, more detailed, example is in the Appendix.

THOSE PESKY EXPECTED ALPHAS

"It's tough to make predictions, especially about the future."
—usually attributed to Yogi Berra[23]

In our experience, practically everyone sees our point that the key dimensions of active management are expected alpha and active risk, that these need to be balanced, and that building portfolios of managers should be an optimization problem. In theory. They follow along and nod their heads in agreement at each step in the progression of the discussion.

But also in our experience, as soon as it sinks in that they will need to form specific numerical estimates of the expected alpha of their candidate managers, the mood often changes. For some, the efficient-market and zero-sum-game alarm lights go on in their heads. They can't see themselves estimating alphas so that the optimization problem for managers can be solved in practice, and they mentally start moving back to their comfort zone, filling out style boxes.

But these same investors—who quail at the task of forming expected alphas on efficient-market and zero-sum-game grounds—almost always do in fact hire active managers. The resulting portfolios contain *implied* alpha forecasts, which, as we pointed out above, can be calculated through reverse optimization. Thus, investors are forecasting expected alphas whether they resist the notion or not. Worse yet, they don't know what alphas they have implicitly forecast, and if they did know they would quite likely reject many of them as unreasonable.

How can one reconcile these conflicting impulses? If you don't think you can forecast alphas, expressing a quantified degree of confidence in a given manager, then maybe you shouldn't hire active managers. You should index instead. If you are going to hire active managers, you might as well make alpha forecasts explicitly, rather than implicitly, since that is more honest and productive than simply selecting some active managers and hoping that the alpha forecasts implied by the holdings weights turn out to be accurate.

Having done that, run an optimizer to select the portfolio.[24] However, if the prospect of specific numerical alpha forecasts for managers is too daunting, then at the very least *think about* manager selection as an optimization problem. Your goal, as in any problem in portfolio construction, is to control risk while maximizing expected return.

Don't worry—you don't have to be a prophet, just a good forecaster. In other words, you don't have to be right about every forecast, just a little bit more right than wrong, across all forecasts, to add value over time (see footnote for suggestions for forecasting alpha).[25]

WHY LOW-ACTIVE-RISK MANAGERS ARE PREFERRED: REASON NUMBER ONE

All of these prescriptions embed the idea that active risk is, in and of itself, bad—that of two equally skillful managers (that is, having the same information ratio), the lower-active-risk manager is to be preferred. But why? Grinold (1990) and Kahn (2000) show that the optimal solution to the utility function for active management must, as a matter of mathematical necessity, be consistent with the following relationship:

$$h_{mgr} \sim E\left(IR_{mgr} \frac{1}{\omega_{mgr}} \right) \qquad (2)$$

where:

h_{mgr} is the holdings weight of portfolio or manager (its percentage allocation),

IR_{mgr} is the expected information ratio of the manager, and

$\frac{1}{\omega_{mgr}}$ is the expected volatility of the manager's pure alpha around a properly established benchmark.

Here's the translation: The size of a manager's allocation is *directly* proportional to the expected information ratio and *inversely* proportional to the level of active risk. If this seems to reward expected alpha once and to punish for risk twice (because active risk is also in the denominator of the information ratio), it does. Risk squared, or variance, is the real operator. So another way to state this result is that a manager's allocation will be higher in direct proportion to a higher expected *alpha* and in *inverse* proportion to active *variance.* Stated this way, it makes intuitive sense.

Remember the "I want more juice" discussion? No, you don't want more juice. You want managers with more alpha, and a lot less risk.

WHY LOW-ACTIVE-RISK MANAGERS ARE PREFERRED: REASON NUMBER TWO

So far we've seen that for a group of managers with equal forecast information ratios, those with greater active risk will have smaller optimal holdings in the portfolio. But there is another effect to consider as well, one that also tends to diminish the size of allocations to managers that have greater active risk.

Most managers are prohibited by their mandates from taking short positions, and hold long-only portfolios. This limits a manager's ability to make bets on its insights, limitations that get exacerbated as active risk is ratcheted up and requires larger and larger active positions. As a result, *even with skill levels held constant,* the information ratio that can be achieved goes down as a manager constructs a portfolio with greater concentration and active risk.[26] Exhibit 2 illustrates this concept.

What are these limitations? A skillful alpha forecasting process is as much at home generating sell signals (negative expected alphas) as buy signals (positive expected alphas). But few of the negative alphas can be acted on in a constrained, long-only portfolio. A given security can only be sold down from the benchmark weight to a zero weight, and not below, no matter how strong the negative signal. The manager's insights are wasted, and the effect is greater as greater risk is taken on (bigger negative positions are indicated but can't be achieved without violating the constraint). The amount of alpha per unit of active risk thus goes down as the active risk level goes up, generating a declining information ratio for a constant level of skill. This effect is remarkably powerful: At a given skill level, enhanced-index funds and market-neutral long-short funds are shown in Kahn's (2000) study to have roughly *twice* as high an expected information ratio as their long-only, traditional active counterparts that have moderate risk.

Exhibit 2

THE IMPACT OF THE LONG-ONLY CONSTRAINT ON PORTFOLIO EFFICIENCY

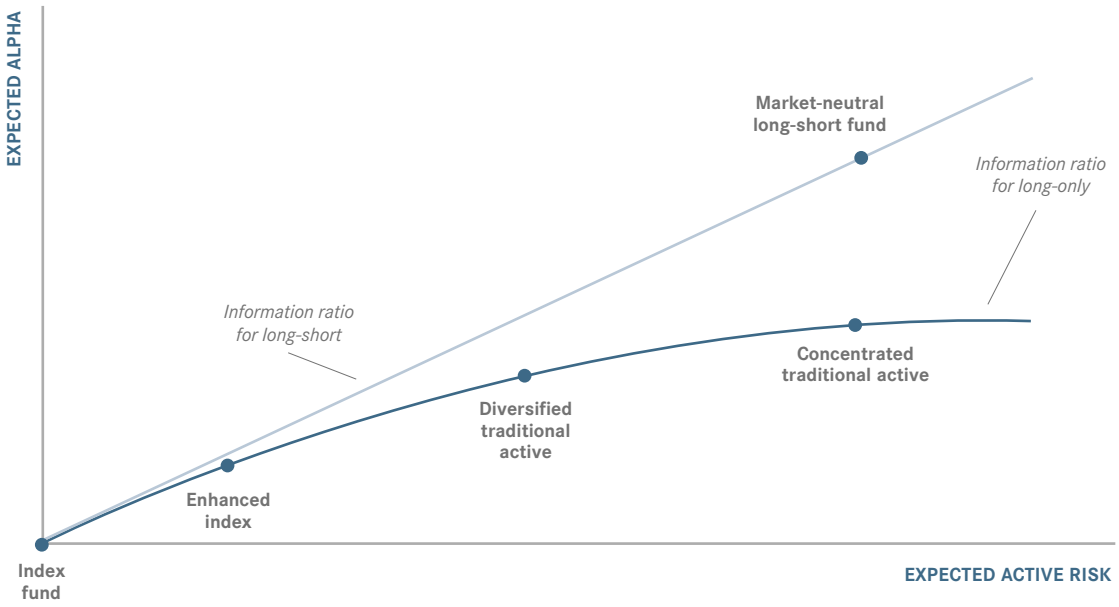

Source: Barclays Global Investors.

So for an additional reason, higher-risk long-only active management is suboptimal. Again, "more juice?" Not. How about a high-protein shake?

HOW MUCH ACTIVE RISK IS ENOUGH?

Since pure active risk (omega) is uncorrelated with policy risk, the relationship is Pythagorean, and total risk is less than the ordinary sum of policy and active risk and typically only slightly larger than policy risk alone. For example, if policy risk is 9% and the active risk is 3%, the total risk will be $\sqrt{9\%^2 + 3\%^2} = \sqrt{.009} = 9.5\%$. So total risk only goes up by 0.5% as a result of adding 3% risk from active management. Does this relatively small increment to total risk suggest that investors should take more active risk, choosing a more aggressive position on the active efficient frontier?

Good question. Theory doesn't give us much help in the practical domain of putting a value on the risk aversion term in any optimization. This term, called λ (lambda), determines how much risk an investor will take in search for the available expected return. We have to look at sources other than theory to get a proper sense for the appropriate ranges of active lambdas for investors.

One source to look at is our own behavior. Recalling the Brinson et al. (1986, 1991) studies and the other related articles cited earlier, about 90% of the variance—not the returns, the variance—of a typical portfolio's returns is attributable to SAA decisions (market risk), and only about 10% is attributable to active decisions (security selection and tactical asset allocation). What does this mean? Let's translate numbers that we know—from the familiar turf of standard deviation—into variance, to get a hint. An investor who takes policy risk

with a 9% standard deviation (about mid-range) has a policy variance of 9% squared or 0.0081; if this same investor's active risk (standard deviation) is 3% then he or she has a 0.0009 active variance, smaller than the policy variance in the ratio of 90 to 10. Thus Brinson's observed variance ratios (90% to 10%) are consistent with more or less ordinary standard deviation numbers (9% and 3%) for policy and active risk respectively. In other words, investors, voting with their feet, reveal a preference for taking far less active risk than market risk.

Why might this be? Some have asked the question, "if 'risk is risk,' why don't investors choose to take on the same amount of variance with respect to active bets as they do with market risk bets?" The answer has to go back to our motivating theme: Policy bets are expected to be rewarded unconditionally, and proportionally to risk taken. Active bets are only rewarded conditionally on skill, and in a declining proportion to risk taken at that; they aren't rewarded at all on average. Thus, one kind of risk is more worth taking than the other. Put this way, it just makes sense that investors would give a higher risk budget to unconditionally expected *market* returns than to highly conditional and proportionally declining expected *alphas*.

We should almost certainly implement this observation by using a lambda risk aversion term that is higher when we are optimizing in the active risk dimension that when we are optimizing in the policy or SAA dimension. Thus, in Equation 1, we indicated with the subscript α that the lambda is specifically an active-risk lambda, different from the lambda expressing aversion to policy risk. At the end of the day, the specific lambda used sets the "active risk budget," and the active risk budget has to be comfortable to the investor given that investor's perception of the alpha that he or she expects in return for taking on that risk.

CORE-SATELLITE?

In the late 1980s many investors started using the "core-satellite" concept. It was a way to temper the risk of an all-active portfolio. An indexed "core" was added to moderate the overall active risk of the portfolio. The concept was even used by some to justify holding extremely high risk, concentrated managers as the "satellites."[27] It made some sense at the time, and represented some first stabs at controlling active risk.

Today's technology, particularly the ability to manage and optimize active risks and returns, allows us to reject the notion of the core-satellite portfolio as an all-too-rigid relic.

But today's technology, particularly the ability to manage and optimize active risks and returns, allows us to reject the notion of the core-satellite portfolio as an all-too-rigid relic, once useful but which we've now grown beyond. For the same reasons that we are rejecting style buckets as insufficiently reflective of the continuous and scalar nature of managers' actual market exposures, the two buckets of "core" and "satellite" can be set aside also.

An optimally constructed portfolio of active managers is likely also to hold index funds. But as any active risk is contemplated, and if "good" low-risk active funds are available, the proportion held in index funds will rapidly shift to low-risk

active funds. This shift will happen smoothly, not abruptly, with increasing risk. Further, the other active managers that are held will be more likely to have moderate risk levels than high risk levels. And absent extraordinary skill levels, not really imaginable, the concentrated portfolio of "just 20 good stocks" is a dinosaur for the reasons discussed immediately above. The core-satellite paradigm is no longer useful when one conceives of manager selection as an optimization problem.

To illustrate this smooth shift among types of managers, examine the allocations at different risk levels in the example case in the Appendix. It shows the allocations across a set of candidate managers including an equity index fund, a risk-controlled active manager, two traditional active managers and a concentrated manager.

VI. What does a portfolio that reflects these principles look like?

THE "NEW COMMON SENSE"

Having built the case for regarding manager structure and selection as an optimization problem, reflecting the reality of pure active return and risk as the real dimensions of active management, we can state what a portfolio reflecting these principles might look like. Moreover, a portfolio reflecting the principles we've advanced should have pretty much the same characteristics whether one formally optimizes or not. Not everyone has the time, focus, or patience to run an optimizer, and we're all equipped with pretty good "fuzzy optimizers" above our shoulders. For investors who use an actual optimizer and for those who don't, the following can inform one's intuition, providing common sense that is directed by the nature of the underlying optimization problem:

1. Be disciplined in forming expected alphas, and in giving the greater weights to managers with higher ratios of expected alpha over active risk squared. Be rigorous when examining *historic* alphas, looking for clues to the future. If a manager's alpha is not statistically significant, why are you looking at it? Your final estimate of expected alpha should be strongly supported by fundamental analysis of the manager and its process.

2. The best portfolio is one that balances the two key dimensions of active management—maximizing pure active return and controlling the total active risk, summed across managers.

   a) The portion of the portfolio that moderates its overall risk will consist of some combination of "good" risk-controlled active funds (sometimes called enhanced index funds) and traditional index funds, which together will likely comprise roughly one-third to two-thirds of the total fund. If an investor has little tolerance for active risk, this portion will lean more towards index funds; if there is more tolerance for active risk, it will lean toward risk-controlled active funds.

   b) Risk-controlled active funds will be more heavily weighted than equally skillful traditional active funds at most risk levels chosen by institutional investors.

   c) Among traditional active managers, prefer skillful lower-active-risk managers over higher-risk, concentrated managers. Bias toward diversified portfolios, away from concentration (unless completely carried away by the concentrated manager's extraordinary forecasting skill).

   d) "Good" market-neutral long-short funds will receive a substantial weight for investors not limited to long-only managers.[28]

3. Set the overall active risk at a comfortable level, your "risk budget." For US equities, a typical investor seems to be most comfortable (in the experience of the authors) at an overall active risk level of 1.5 to 2%, with the very largest investors preferring even less active risk (between 0.75 and 1.25%).

4. Keep a careful eye on misfit risk, trying to minimize it while still maximizing expected alpha. If you don't use an optimizer designed especially for this purpose, you'll have to use a style map or "effective asset mix" table as a supporting tool.

VII. Conclusion: Managing in the right dimensions

We recalled two "first principles" that emerged decades ago from the basic academic work in finance. The first is Harry Markowitz' observation that investors should be concerned about risk as well as return, which he shaped into the mandate that investors build mean-variance efficient portfolios, using an optimizer. The second is William Sharpe's demonstration that the total risk of an investment can be broken into policy risk, which is rewarded by the existence of an equity risk premium, and active risk, which is a zero-sum game when summed across managers and thus not rewarded on average.

We added the observation that managers (being people) have skill levels that really do differ from one another. In addition, capital markets aren't completely efficient. As a result, someone is going to win the active game due to real skill—not just luck or random variation—even while someone else is losing. It pays for the skillful investor to try to discern who will be the winners, giving active management a vital role in the portfolio.

We noted that while the payoff to market-related (policy) risk is linearly related to the amount of policy risk taken, that is not true of pure active risk, even conditional on the manager having real skill. Without skill there is no payoff for pure active risk. But in the presence of skill there is a payoff, one that declines as a proportion to pure active risk as the amount of risk taken increases (it does not decline for unconstrained, long-short strategies).

We then tried to identify the logical consequences of this view of the world.

First, investors should build efficient, or optimized, portfolios of *managers* just as they should of asset classes (or securities within an asset class). Market or policy risk and return are commonly managed through the strategic asset allocation process; we are suggesting that the active-risk decision across the portfolio of managers, being uncorrelated with policy, can and should be managed through its own separate optimization process.

To do so, investors must estimate the expected alpha (after fees and costs) for each manager. Let's spend a moment recalling the importance of this step.

An investor must meet two conditions if he or she is to hire active managers.[29] First, one must believe that superior managers really do exist. That's easy, if one accepts that managers differ in their skill levels. Second—this is the hard one—one must believe that he or she can identify which ones will be the winners. To accomplish this, one needs to be able to make specific alpha forecasts for managers, forecasts that are somewhat more right than wrong. If you can't do that, you should just index.

At ordinary risk levels, and with manager skill levels being generally equal, the optimization will give us a mix of index funds, risk-controlled active managers, equity market-neutral long-short funds,

and moderate-risk traditional active managers. Most likely, it won't give us more than light allocations to *concentrated* managers. At typical risk levels it will favor lower active risk managers over higher risk managers.

Second, investors should disentangle "pure alpha," the part of the active return that is the unique creation of the manager, from the various market factors that are in the manager's customized benchmark or "normal portfolio." The manager's job is to beat this customized benchmark, not the naïve asset-class return. Moreover, when investors make alpha forecasts for managers, it is this pure alpha that they should be forecasting. We can let the optimization process reduce the net misfit risk across all managers, balancing value managers appropriately against growth, and large-cap managers against small, so that the portfolio of managers looks as much like the investor's benchmark as is sensible. This is important; if the collection of the managers' market-related risks doesn't look like the benchmark, the investor's strategic asset allocation policy has surreptitiously been changed.

Managing active managers in these dimensions is simpler than in conventional practice—we are dealing with risk and return, the basic building blocks of finance. Yet it is, at the same time, more complex: The difficult task of specifying scalar values for pure alpha and for misfit risk replaces the easier task of just filling out discrete style boxes. But the scalar approach is the only one that reflects reality; we don't want to pretend that the world is simple where it is complex and there are fine gradations. As Einstein said, "Everything should be made as simple as possible but no simpler."[30]

Active management offers the skillful an opportunity to influence the portfolio meaningfully by adding pure alpha. For those confident of their skill, the question is not *whether* to use active management, but *how.* Hard-working people have entrusted us, the community of investors and asset managers, with trillions of dollars—basically, with much of their worldly goods. We owe them nothing less than the best application of sound financial knowledge that we can deliver. We know that the first principles synthesized here are valid, at least to the extent that today's understanding of economic science allows us to make such a claim. Let's apply these principles the best that we realistically can.

At ordinary risk levels, and with manager skill levels being generally equal, the optimization will give us a mix of index funds, risk-controlled active managers, equity market-neutral long-short funds, and moderate-risk traditional active managers.

PART TWO:
Debunking Some Myths of Active Management

In our travels we speak to hundreds of large investors. There are a number of questions about active management and active managers that come up over and over again. They come up because what we are saying is not always consistent with the "prevailing wisdom" long taught 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 old saying that "an untruth repeated often enough becomes the truth," and whether unchallenged myths assume the status of fact.

We strongly believe that, in the presence of skill, active management can be successful. But we also believe that 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, sowing confusion, misunderstanding and, ultimately, distrust of healthy management disciplines.

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

A: Inefficiency means that there is *some* information out there, available to *some* investors, that is not yet impounded into security prices. If among those that have access to this information there are people with special skill at using it to forecast prices, they will generate a positive pure active return at the expense of those without the information or that are unable to skillfully interpret it and act on it. Having, and exercising, this forecasting skill is the only way to truly "beat the market."

Active management in a market, 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. At its 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 less than a dozen stocks in its equity index (and yes, the major index providers and index fund managers do work with such small countries). Many of these countries by any definition would be considered to be grossly inefficient: there is no equivalent to the US Securities 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).[31]

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 mid- and late 1990s, most US-based international equity managers beat their benchmarks by 5%, 7%, and more, 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 (in their presumably ivory towers) have not yet squarely faced? The winning managers often attribute their success to the inefficiency of international equity markets.

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.

It is now widely acknowledged that the US managers gained this unusual performance by underweighting Japan relative to MSCI EAFE, during a period 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 US 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 others' stocks or holding their own "treasury stock,"[32] and likely also individ-

Exhibit 3

ACTIVE MANAGERS VERSUS EAFE

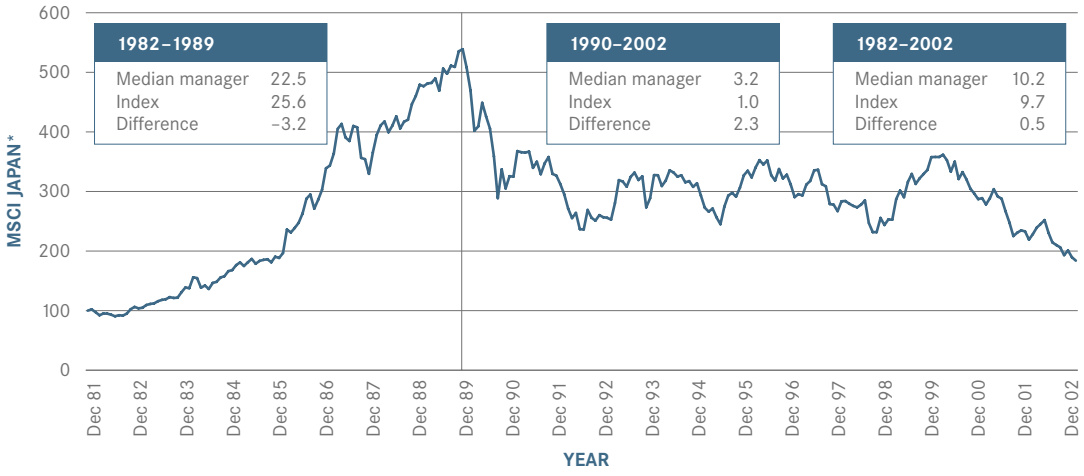

* *Gross dividend yen returns.*
Source: InterSec, MSCI, BGI.

ual and institutional Japanese investors. Most countries' investors hold a disproportionate percentage of their own local securities, the so-called "home country bias."

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

During the last two decades, Japan grew from a modest component of the non-US equity market to the dominant component (about 60% at the high), and since then it has fallen back to a more modest 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 3.

So, looking at the decision to underweight Japan, over the entire history of that bet by US 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 relationship between inefficiency, alpha opportunity and transactions costs, but there is some intuition that the other side of the trade will always try to protect itself against the trader that 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?

A: We've heard this ability claimed by traditional active managers. This story is *sooo* appealing. We all *want* to believe that 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 (over and above the risk-free rate):

$$r_p = \beta_p r_{bm} + \alpha \qquad (1)$$

In text: "The excess returns for a manager's portfolio can be explained as the sum of the excess return for taking market risk through the benchmark (beta-adjusted), 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, then the parallel risk model for this portfolio has to be:

$$\sigma_p^2 = \beta_p^2 \sigma_{bm}^2 + \omega^2 \qquad (2)$$

In English instead of Greek, "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 one,[33] 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 $\beta_p^2 \sigma_{bm}^2$) will quite naturally differ from the natural risk of the benchmark (the benchmark of course *does* have a beta of one).[34] 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. However, stare for a minute 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 one 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 large. But if that manager really did have the ability to deliver such a big alpha, wouldn't it be touting it directly, promising to deliver it *on top* of a proper benchmark exposure?[35]

Just as important, if you the investor really wanted 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 which 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 correlated positively or negatively, it simply means that the manager has a beta that is different from one.

No matter how high the manager's true mean alpha, with some predictable probability that manager will have periods of underperformance.

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. As already discussed, 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, hard-headed 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 that they simply match the benchmark in up markets, but that they beat it in down markets. How can I test this claim?

A: This is a claim often made by traditional active managers, particularly value managers. (But it seems like we only hear it during up markets.)

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 that 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? I would like to hire a manager that only outperforms. Is this possible?

A: Position-by-position, day-by-day, the alpha distribution for any manager will very much be 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 (much too 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.[36]

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. So long as fresh new decisions are always being made on fresh new information, the average alpha of a manager's returns will approach being normally distributed over many periods (or log-normally, 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 even histrionic or platitudinous![37]

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 that 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. Thus the difficulty in separating luck from skill.

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

A: 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 (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, that they predictably beat their benchmark over time.

We can find some small evidence that directionally supports this claim, but the magnitude of this evidence isn't sufficient to make the case that professionals should automatically be preferred over index funds.

What is that evidence? Gary Brinson and a series of others[38] have conducted empirical studies of the returns achieved by professional investors (that is, pension plans and other institutions and their managers, and mutual funds). The average alphas (after all fees and costs) are slightly different in each study, but it isn't unfair to summarize the results across these studies 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 this 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 nonprofessionals 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.

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.

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

A: We can do some very quick and dirty statistical stuff to answer this question. We'll use the rule of thumb for calculating the statistical

significance "t-statistic," noting that it is in essence just the pure information ratio (IR) multiplied by the square root of time.[39] If it required 80 years to get a significant two-tailed t-statistic of 2, and calling the square root of 80 an even 9, that suggests this 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.[40]

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 for such a manager. 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.

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, most of which aren't significant, we will at least caution you to ignore our advice at your own risk. And even when historic alphas are statistically 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.[41]

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

A: It is evident from the main body of this article that active portfolios of anything, either of securities or of managers, if not informed by skill (positive information coefficient, or IC) are necessarily just poorly constructed index funds. The forward-looking expected return of such a portfolio, before costs, is the same as that for the properly constructed index benchmark for the portfolio (i.e., the market return plus a zero expected pure active return).

And of course such a portfolio won't be a very good index fund. It will experience high tracking error—but that's OK because the costs will be high as well! And if it is not carefully constructed, it may have misfit risk as a result of being mis-weighted relative to the benchmark and its components, 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 which isn't in fact informed by skill, and so it will only out-perform (or under-perform!) 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.*

And 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 that 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?

A: The ability to move a fund to a lower beta position, either by going to cash or bonds or otherwise going "more conservative," is a frequently asserted advantage claimed by traditional active managers when comparing their approach to many more modern portfolio construction approaches—index funds, enhanced index funds, in fact over 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 that have special skill and insight into making market timing bets will win persistently; any other winners will be just lucky and are likely to be unlucky 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 for active decisions is that they add risk. It turns out that market timing adds a great deal of risk. Mark Kritzman does an excellent job of explaining this.[42]

Here is the *extra* caveat that we don't need to bother with when discussing security selection approaches, and it's a subtle one: Market timing is what we call a "low breadth" active management discipline. This means that there are very few investment decisions being made. To be successful, even a skillful investor needs a lot of bets across which to diversify his or her 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 the manager can achieve a higher information ratio with whatever skill level they do possess.

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 done, 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 a manager 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 that skill has been adequately demonstrated.

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

A: It really isn't "our" benchmark-relative framework. The term "benchmark" is in one aspect misleading. The term 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 of saying 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 simply by reference to the benchmark index 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 only positive in expectation *conditional* on the manager having skill.

A manager's task is to generate pure active return without taking too much pure active risk.

The long-only manager that 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 among which they rotate. Yet there is some "normal portfolio" representing the manager's benchmark level of market risk hidden in there somewhere, of necessity. In extreme cases it is best for the analyst to simply default and use the closest general benchmark as a good-enough proxy—the S&P 500 for example. But a better job can usually be done simply by understanding the manager's articulation of its process. A value vs. growth style rotator, for example, might acknowledge favoring value; the investor might estimate his 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 "falls out," 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 require if they are to deliver similar information ratios. It

often seems that managers that most strongly resist the concept of benchmark relative investing are those that simply desire to avoid clarity of accountability for producing "pure alpha" and thereby hope to maintain the mandate for long periods without rigorous review.

Q: What makes a manager a "good" manager?

A: A manager's task is to generate pure active return without taking too much pure active risk. As discussed in the main body of the article, active managers are too expensive to hire just for the purpose of 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 SAA 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.

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

a) Good managers "get it." They know the difference between true alpha and market returns. Igor Sikorsky, reflecting on his early years developing helicopters, is reported by his son Sergei 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."[43] You don't want managers that represent ignorance squared.

   Most (this may not be too strong a word) traditional managers still don't understand that they aren't hired to deliver persistent misfit risk relative to their benchmarks. Examples of this are found in the so-called "core-plus" strategies in fixed income (high yield debt persistently and intentionally held in a portfolio benchmarked to the Lehman Aggregate Bond Index), 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 properly stylized 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 that "get it."

b) The good manager's active risk levels are low or moderate (for long-only funds) and are carefully budgeted and managed. The long-only constraint is a powerful force reducing 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.

c) 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. If it is a process contributed to by many people, the process must harness all contributors to the effort to deliver high information ratios in the final portfolio, when all their inputs are aggregated together. But the good manager today is less likely to be a single smart person, or even a group of smart people managed by a single smart person. Good managers are finally moving from the preindustrial period of building one handmade musket at a time using individual, nonstandard crafts, to Eli Whitney's industrial model where smart people organize the 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 purify 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 teams of people trying to pick one security at a time.

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

A: We touch on this in a footnote in the main article. The question deserves a full article of its own, which we hope to provide in the future. Suffice it for us to repeat that an investor is a portfolio manager, managing a portfolio of managers who in turn manage a portfolios of stocks. So an investor must have skill at selecting managers that have skill at picking stocks. 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 is possible, at the end of the day skill can only be acted on 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 through forecasting returns and holding appropriate portfolios reflecting 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 require your thoughtful and educated judgmental inputs.

Appendix: Manager structure optimization example

This example shows a stylized manager structure optimization case. The alpha assumption is entirely artificial, assuming a constant pure information ratio of 0.05.

At a typical 2% active risk level, the optimal portfolio would be about half traditional active and half risk-controlled active, with no index fund. The declining holdings of index funds and their replacement by risk-controlled active funds demonstrates the point that we made in our section discussing core-satellite investing. Likewise the concentrated manager receives only a small allocation contrary to some core-satellite interpretations. To make the best possible case for using concentrated managers, we used the same information ratio of pure alpha to active risk for the concentrated manager as for the others. In the real world, and assuming equal skill, the long-only constraint would cause the concentrated manager's information ratio to be much lower than the others. So these small allocations are in reality much overstated.

SPONSOR BENCHMARK: RUSSELL 3000 BROAD CAPITALIZATION US EQUITY

MANAGER ASSUMPTIONS **Manager (with style betas)**	**Expected alpha (%)**	**Expected active risk (%)**
Index Fund *(Russell 3000)*	0.00	0.00
Risk-Controlled Active Manager *(Russell 3000)*	0.75	1.50
Growth Fund *(80% R3 Growth 20% R3 Value)*	2.50	5.00
Value Fund *(20% R3 Growth 80% R3 Value)*	2.50	5.00
Concentrated Fund *(Russell 3000)*	9.00	18.00

CAPITAL MARKET ASSUMPTIONS*					
		Correlations			
Style index	**Expected total risk (%)**	**R1 G**	**R1 V**	**R2 G**	**R2 V**
Russell 1000 Growth	19.39	1.00			
Russell 1000 Value	14.11	0.70	1.00		
Russell 2000 Growth	25.44	0.78	0.49	1.00	
Russell 2000 Value	14.24	0.55	0.70	0.76	1.00
Equity Benchmark (Russell 3000)		**46.6%**	**46.6%**	**3.4%**	**3.4%**

OPTIMAL MANAGER ALLOCATIONS							
	Active risk budget (%)						
	0.0	**0.5**	**1.0**	**1.5**	**2.0**	**2.5**	**3.0**
Index Fund	100.0	72.2	44.4	15.8	0.0	0.0	0.0
Risk-Controlled Active Manager	0.0	16.5	33.0	50.0	52.5	38.7	17.2
Traditional Active: Total	0.0	11.3	22.6	34.2	47.5	61.3	82.8
Growth Fund	0.0	5.0	9.9	15.0	20.7	26.4	35.2
Value Fund	0.0	5.0	9.9	15.0	20.7	26.4	35.2
Concentrated Fund	0.0	1.4	2.8	4.2	6.2	8.6	12.4

Based on 10 years of data (11/92–11/02).

Endnotes

1 While there hasn't been much academic research focused on managing active decisions, there is an excellent (if abbreviated) literature on this topic, the most authoritative, general and up-to-date example being Grinold and Kahn (2000a). The technology described therein can be recast in a form that applies not just to the security selection problem, where their work was focused, but also to the manager selection and structuring problem, as previously demonstrated by one of the authors of this paper. See Waring, et al. (2000). These works provide a substantial bibliography of prior research in this area.

2 We'll often refer to securities as "stocks," although our comments apply to securities in any asset class or across any group of asset classes. Also, while we'll frame this discussion in the familiar turf of security selection, it applies equally to the alpha-generating efforts of tactical asset allocation, market timing, sector rotation, and other methods of actively managing across groups of securities.

3 In addition to the weights of the securities held, for reverse optimization one needs a reference portfolio (the benchmark), estimates of the pure active risk (standard deviation) of each stock, the correlation of each stock with every other; and the expected alpha-for one of the stocks (to get the scale right). For a general discussion, see Sharpe (1964).

4 A very simple case: If three active managers hold half the portfolio (the other half in index and enhanced index), an 8% expected alpha may be implied for them!

5 See Sharpe (1964).

6 While the difference in returns between equity and fixed-income markets (the "equity risk premium") may vary in size over time, logic pretty much dictates that its expected value must always be positive for markets to clear. See Grinold and Kroner (2002) and Leibowitz et al. (2001).

7 The "Evans-Archer" diagram (see Evans and Archer [1968]) is the classic tool used to introduce this point, showing risk declining asymptotically down to some irreducible amount as the portfolio becomes more and more diversified. Although Evans and Archer were only contrasting policy and active risk, without saying what risks were rewarded, one can deduce that diversifiable risk need not be rewarded since investors can avoid it, almost for free, by indexing; but the undiversifiable risk does need to be rewarded. Sharpe (1964) shows that, under the stringent assumptions of the CAPM, policy risk is associated with an expected reward while active risk is not. Many subsequent observers have criticized the CAPM as relying on unrealistic assumptions or as being empirically unsupported. However, Ross (1976) demonstrated in the context of developing his Arbitrage Pricing Theory that, even if the CAPM conditions do not hold, policy risk is still rewarded while active risk is not. Thus, while acknowledging the limitations of the CAPM, we can proceed with confidence in dividing investment risk into rewarded and unrewarded components.

8 This version of the argument is developed most eloquently in Sharpe (1991).

9 The magnitude of the component of variance that is explained by luck rather than skill is surprisingly large. Since the IC is just a correlation coefficient, for a process with correlated signals we can convert it to an R-squared by just squaring it. Grinold and

Kahn (2000a), p. 292. A manager with a very respectable IC of 0.10 would have an R-squared of 0.01, meaning that just 1% of active variance was explainable by the manager's skill, the remaining 99% of active variance being just random luck. Astounding. Small wonder that it is so difficult to separate skill from luck over short time periods.

10 We aren't really sure who was first to classify investments or managers based on regression factors, but certainly Barr Rosenberg and his successors at BARRA have been the most complete at making this practice into a science. See *www.barra.com/research/barrapub/risk_models.asp* for a good introductory-level description. Sharpe (1988, 1992) saw that style and size factors could be productively and intuitively used in such regressions (called returns-based style analysis), and did much work in the area involving special types of regressions designed to make the output more intuitive to lay audiences. Sharpe's approach is certainly the most commonly used in practice. For determining which styles or common risk factors are most relevant for use in returns-based style analysis, one authoritative source is Fama and French (1993), drawing on the much earlier observation by Banz (1981) and Reinganum (1981) that small-cap stocks had historically outperformed larger-cap stocks, and the observation by Basu (1977, 1983), among others, that "value" stocks (having low price/book, price/earnings, or other valuation ratios suggesting that the stocks were cheap) had outperformed "growth" stocks (having expensive valuations).

11 Realized pure alpha may be determined by an actual multivariate regression or by a constrained optimization technique that mimics a regression such as Sharpe's style analysis method. The purpose of the optimization technique is to allow for a no-shorting constraint; that is, to require all factor betas to be between zero and one. We generally prefer ordinary regression, for its greater ability to accurately describe "deep" value and growth managers.

Pure active return, of course, is properly thought of *after* fees. One should also incorporate manager-transition costs into the pure active return; these must be amortized over the time period for which the manager is likely to be held, so that the cost (which is paid only when the manager is hired or fired) is properly converted into annualized return form.

12 The CAPM uses just a single factor to capture the market. We are simply trying to control risk relative to the asset allocation policy better, by dividing market risk into more granular subcomponents. Either way, regression is a useful model for sorting out the market and idiosyncratic components of risk and return.

13 In our experience, some investors don't find that the term "information ratio" conveys much intuition about its meaning. So, observing that IR measures the consistency with which the active return is delivered, perhaps "consistency ratio" would be a better term for this concept, and one more likely to be commonly appreciated. IR is a key measure of historic "goodness" for a manager, and on a forward-looking basis is a key input into the manager's role in the portfolio. It is very useful—it incorporates the two key dimensions of active management into one measure.

14 At least among the more sophisticated, scientifically oriented managers. See Markowitz (1952, 1991). Grinold and Kahn (2000a), oft-cited throughout this paper, is a complete and up-to-date text on this topic.

15 Quoted in Bell (1952), at *www-gap.dcs.st-and.ac.uk/~history/Quotations/Einstein.html*.

16 See, for example, Kahn and Rudd (1995), which contains an extensive review of the literature on the topic up to that time, which happens to include most of the relevant literature even today. Further work on the topic has been done by Carhart (1997) and Wermers (2000) and their results do not differ much from those summarized in Kahn and Rudd (1995). Most of these performance studies cover mutual funds because of easy data availability, but the findings are extremely likely to apply in roughly the same way to institutional funds, which are largely managed by the same managers.

17 To calculate a manager's alpha t-statistic, divide the realized historic monthly, quarterly or annualized arithmetic alpha (regression or pure alpha) by the standard deviation of that alpha (expressed at the same frequency), and multiply the result by the square root of the number of periods represented in the data. Note that if the data is expressed in annualized terms, this result is just the information ratio times the square root of the number of years of data. Every regression software package, including Microsoft Excel™, provides this t-statistic automatically whenever a regression is conducted. By using a regression alpha, we eliminate any accidental market return effects that might otherwise distort the manager's actual non-market-related returns. It is pure active return, or alpha.

18 Indeed, BGI enjoys significant t-statistics in the majority of its active strategies, measured since inception. We note this, of course, only to show that successful active management is indeed possible (we would hate to taint our scholarly argument with a commercial!).

19 One can use either returns-based or holdings-based style analysis, or both, to calculate these weights. There is no clearly best single method.

20 Correcting an error in the optimization formula at Equation 2 in Waring et al. (2000), the first term, representing "misfit return," should have a beta in it, as follows:

$$(\mathbf{h}^{\mathrm{T}}\mathbf{X} - \beta \mathrm{h}_{\mathrm{b}}^{\mathrm{T}})\mathbf{r}_{\mathrm{k}}$$

If the return and risk assumptions are all estimated in such a way that they fall onto a common security market line, this is a zero term. In this case the whole term can be dropped.

21 At least for US equity managers. For international equity managers and for fixed income managers, correlations between pure alphas are empirically nonzero (because these asset classes have fewer degrees of freedom—in fixed income there are only a couple of major bets to be made, and for international equity managers, there appears to be more emphasis on regions and industries than on individual security selection). Thus one might wish to make specific estimates of manager cross-correlations when dealing with these asset classes.

22 There are few factors if one uses style analysis, but not necessarily so few if one were to use a more detailed factor model such as BARRA's. The good news is that the highly detailed factor models are commercially developed and conveniently come with capital market assumptions for each factor.

23 See *www.larry.denenberg.com/predictions.html* for a list of the many personalities to whom this quote has been attributed, along with related Web links.

24 As Michaud (1989, 1998) has pointed out, optimization outputs are no better (and may in fact be worse) than the quality of the inputs, which are statistical estimates and thus subject to natural estimation error. Kritzman (2003) has responded by noting that, for any given set of inputs, optimization is still the tool that gives the best outputs. He also makes many valuable suggestions about forming good forward-looking optimizer inputs.

25 To get started in forecasting alphas, investors might consider using the forecasting relation:

Alpha = IC x Volatility x Score

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). This formula is further elucidated, in a security selection context, in Grinold and Kahn (2000a). Additional improvements to recognize that most managers are subject to the long-only constraint, and to include the effect of fees, can be incorporated.

The Score is the key input variable. A two-standard deviation manager, with a score of 2, is the type of unusually skillful manager one might hope to find, but of course is quite rare. For IC, if you have no skill at manager evaluation, put in a zero (you're only right about managers half the time) and quit. Those truly blessed with selection skill might try a 0.3 (you're right about managers 65% of the time) or if you're rakishly overconfident, a 0.5 (you're right about managers 75% of the time!).

26 See Grinold and Kahn (2000b).

27 As stated above, for a given level of skill, the information ratio must be lower for a concentrated high-risk active manager then for a lower-risk manager. And, of course, if the concentrated manager really is unskillful, its high risk level will mean that downside realization will be especially painful. On average, the highly concentrated manager perfoms at the average (before fees and costs—just like any other active manager.

28 And alphas delivered by high information-ratio managers sourced in one asset class can, at least theoretically, be ported to another asset class (where, perhaps, high information-ratio managers are scarce) by the use of futures or other derivatives. Such a "portable alpha" strategy is most frequently used to add alpha, generated in hedge-fund programs, to an equity or fixed-income account.

29 See Waring, Pirone, Whitney, and Castille (2000), appendix C; and Waring, Harbert, and Siegel (2001).

30 *Reader's Digest* (October 1977). Peter Neumann, who knew Einstein personally, writes (see *www.csl.sri.com/users/neumann/neumann.html*), "I was probably just one of many people who heard him say something like, 'Everything should be as simple as possible but no simpler' or 'Everything should be made as simple as possible but no simpler' … After his death, I asked Mrs. Dukas (his long-time assistant) if there was any record of such a statement in any of his writings (she was the one responsible for organizing all of his papers); she eventually responded that she was unable to find any written record attributable directly to AE."

31 We can highlight the difficulty with the "go active in inefficient markets" myth by noting that we might require the dubious skill of selecting the manager with the best access to inside information.

32 Part of the fault lies in the benchmark, which until recently was not float-adjusted, so that illiquid shares (such as crossholdings, common in Japan) were counted as part of the market when they shouldn't have been. The float-adjusted version of the benchmark, had it existed then, likely would have had a smaller relative weight than did the non-float-adjusted version actually then in use. The relative size of the US managers' underweight positions would have been much smaller by reference to a float-adjusted benchmark.

33 Grinold and Kahn (2000a) identify the portion of the beta that is different from 1 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 normal portfolios, rather than the variable value associated with market timing strategies.

34 Remember, we're 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 (2000a), Appendix to Chap. 4.

35 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 it. This is discussed on page 30.

36 OK, the underlying distributions must have a finite variance, but that's not an issue here.

37 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.

38 For these data, see Brinson, Hood, and Beebower (1986) and Brinson, Singer, and Beebower (1991). For generally similar results and an updated and improved interpretation, see Ibbotson and Kaplan (2000) and Surz, Stevens, and Wimer (1999).

39 Any standard regression package will readily calculate the precise t-statistic for the analyst. But 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. The quick and dirty method used in the text works with an annual periodicity, as that is how information ratios are usually stated.

40 Grinold and Kahn (2000a), page 130.

41 We're necessarily 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 more rigorous work should be done here than is currently done. We'll leave it to others to explore the fine points of that rigor.

42 Kritzman (2000) Chap 5.

43 Sikorsky (2003).

References

Banz, Rolf W. (1981). "The Relationship Between Return and Market Value of Common Stocks." *Journal of Financial Economics* (March).

Basu, Sanjoy. (1977). "Investment Performance of Common Stocks in Relation to Their Price-Earnings Rations: A Test of the Efficient Market Hypothesis." *Journal of Finance* 32 (3): 663–81.

Basu, Sanjoy. (1983). "The Relationship Between Earning's Yield, Market Value and Return for NYSE Common Stocks." *Journal of Financial Economics* 12: 129–56.

Bell, Eric Temple. (1952). *Mathematics, Queen and Servant of the Sciences.* Reprint, Washington, DC: Mathematical Association of America, 1987.

Brinson, Gary P., L. Randolph Hood Jr., and Gilbert Li. Beebower. (1986). "Determinants of Portfolio Performance." *Financial Analysts Journal* (July/August): 39–44.

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

Carhart, Mark M. (1997). "On Persistence in Mutual Fund Performance." *Journal of Finance* (March).

Evans, L. John, and N. Stephen Archer. (1968). "Diversification and the Reduction of Dispersion: An Empirical Analysis." *Journal of Finance* 23 (5): 761–67.

Fama, Eugene F., and Kenneth R. French. (1993). "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics* 33 (1): 3–56.

Grinold, Richard C. (1990). The Sponsor's View of Risk. *Pension Fund Investment Management: A Handbook for Investors and Their Advisors,* edited by Frank J. Fabozzi. Chicago: Probus.

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

Grinold, Richard C., and Ronald N. Kahn. (2000b), "The Surprisingly Large Impact of the Long-Only Constraint." *Investment Insights* 3 (2). San Francisco: Barclays Global Investors.

Grinold, Richard C., and Kenneth F. Kroner. (2002). "The Equity Risk Premium." *Investment Insights* 5 (3). San Francisco: Barclays Global Investors.

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

Kahn, Ronald N. (2000). Most Pension Plans Need More Enhanced Indexing. In *Enhanced Indexing: New Strategies and Techniques for Investors,* edited by Brian R. Bruce. New York: Institutional Investor, Inc. Investment Guides.

Kahn, Ronald N., and Andrew Rudd. (1995). "Does Historical Performance Predict Future Performance?" *BARRA Newsletter* (Spring). Available online at *www.barra.com/research/BarraPub/hpp-n.asp.*

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

Kritzman, Mark P. (2003). "Value In–Garbage Out." *Economics and Portfolio Strategy,* New York: Peter L. Bernstein, Inc.

Leibowitz, Martin L. et al. (2001). At the *Equity Risk Premium Forum* (November 8, 2001) sponsored by the Association for Investment Management and Research, Charlottesville, VA. Available online at *www.aimrpubs.org/ap/issues/v2002n1/pdf/v2002n1.pdf.*

Markowitz, Harry M. (1952, 1991). *Portfolio Selection: Efficient Diversification of Investments,* 1st ed. 1959; 2d ed. 1991. Oxford: Basil Blackwell, 1991. First published in Journal of Finance (March 1952).

Michaud, Richard O. (1989). "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal* (January/February).

Michaud, Richard O. (1998). *Efficient Asset Management.* Boston: Harvard Business School Press.

Reinganum, Marc R. (1981). "Misspecification of Capital Asset Pricing: Empirical Anomalies Based on Earnings' Yields and Market Values." *Journal of Financial Economics* 9: 19–46.

Ross, Stephen A., (1976). "The Arbitrage Theory of Capital Asset Pricing." *Journal of Economic Theory* 13 (3): 341–60.

Sharpe, William F. (1964). "Capital Asset Prices: A Theory of Market Equilibrium Under Conditions of Risk." *Journal of Finance* (September): 425–42.

Sharpe, William F. (1974). "Imputing Expected Security Returns from Portfolio Composition." *Journal of Financial and Quantitative Analysis* (June): 463–72.

Sharpe, William F. (1988). "Determining a Fund's Effective Asset Mix." *Investment Management Review* 2 (6): 59–69 (corrected version).

Sharpe, William F. (1991). "The Arithmetic of Active Management." *Financial Analysts Journal* (January/February).

Sharpe, William F. (1992). "Asset Allocation: Management Style and Performance Measurement." *Journal of Portfolio Management* 18 (2): 7–19.

Sikorsky, Sergei. (2003). "Igor Sikorsky." *AOPA Pilot* (February): 138.

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

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

Waring, M. Barton, Duane Whitney, John Pirone, and Charles Castille. (2000). "Optimizing Manager Structure and Budgeting Manager Risk." *The Journal of Portfolio Management* (Spring).

Waring, M. Barton , Lee D. Harbert, and Laurence B. Siegel. (2001). "It's 11 P.M.—Do You Know Where Your Employees' Assets Are?" *Investment Insights* 4 (2). Barclays Global Investors, San Francisco.

Wermers, Russ. (2000). "Mutual Fund Performance: An Empirical Decomposition into Stock-Picking Talent, Style, Transaction Costs, and Expenses." *Journal of Finance* (August).

Investment Insights

Published by BlackRock

Please direct story ideas,
comments, and questions to:
Marcia Roitberg, editor
Telephone 850 893 8586
Facsimile 415 618 1455
marcia.roitberg@blackrock.com

BLACKROCK
55 E. 52nd Street
New York, NY 10055

www.blackrock.com

As of December 1, 2009, the former Barclays Global Investors entities referred to in this publication are owned and controlled ultimately by BlackRock Inc. Neither Barclays nor any Barclays affiliate is responsible for, or associated with, the contents of this publication.