The Dimensions of Active Management

Why alpha and active risk are the only things that matter.

M. Barton Waring and Laurence B. Siegel

Every reader of this article likely either employs active managers or is one. But what exactly are active managers? What role do they play with regard to the investor's portfolio? What should their role be? What are the real dimensions on 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 pose, 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).

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 trade-off 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.

M. Barton Waring is a managing director and heads the Client Advisory Group at Barclays Global Investors in San Francisco (CA 94105).

Laurence B. Siegel is director of investment policy research at The Ford Foundation in New York (NY 10017).

Current practice, to the contrary, 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: balancing the hope for gain against concern about risk. The current state of practice does not do a good job of managing that tension, and doesn't even explicitly focus on it, but there are new technologies and ways of thinking that do.

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 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, or optimize those forecasts to construct portfolios (although 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.

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 you do agree with.

Likewise, investors make alpha forecasts for the managers they hold. These forecasts are rarely explicit in current practice, but they are also always implicitly there. They also can be backed out through reverse optimization—and they can be embarrassingly high.[4]

Investors should hire active managers only if they have skillfully formed positive alpha forecasts for the managers. And if they have these forecasts they should use them (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: for managers looking at stocks, and for investors 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 total risk into its component parts. Following Sharpe [1964], 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.

Policy risk is usually managed by investors through their adoption of a strategic asset allocation (SAA) policy, with its attendant asset class benchmarks, all fully diversified and therefore subject only to 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.

Academicians have focused much of their finance research in the last few decades on market risk, observing that if markets are efficient, then we can't beat the market. In practice, however, investors don't seem to completely believe this. Investors still routinely hire active managers, whose returns can differ widely from benchmark returns. In fact, the all-passive institutional portfolio is a rarity. So active risk is in fact taken in the search for pure active return, and this trade-off needs to be managed too, besides the trade-off between market return and risk that is at the heart of asset allocation policy.

Capital markets must function so that the expected return on the overall equity market must be higher than that on fixed-income investments.[5] 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.[6]

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. Active managers (and other market participants) in aggregate *are* the market, however, so they must again in the aggregate earn the return on the market—minus fees, transaction costs, and other expenses, which can be substantial.[7]

Where's the Beef?

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

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, some managers will outperform through real skill, not just by virtue of random variation. Under these conditions, the notion that there can be an expected alpha or pure active return makes perfect sense. It is important to notice that while 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, so there will always be risk left over, the 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 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.[8]

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 strategic asset allocation 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—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.

DEFINING PURE ACTIVE RETURN AND RISK

We agree on the separation of market or policy risk and pure active risk, in principle. Let's go one step further and apply that 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 (not pure active returns). 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 capital asset pricing model 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.[9] 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 investibility 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 the same only 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 so that, at least in sum across the managers, they are consistent with the benchmark's factor weights.

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 established 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 that you assign 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 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 it 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 it is redundant to say pure alpha, but we're trying to call attention to this precise definition, one that doesn't entail a market (or style) risk component. It is also what Sharpe means by the 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.[10] The regression determines the effective style weights of the manager, or the 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 the 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 or by any other information that is useful to characterize the normal style biases of the manager. Even a tactical asset allocation manager or a style rotator has such a home position.

The view of market risk that 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.[11]

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 (e.g., 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 exposure). 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.

The Information Ratio

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, and pure active risk or omega, ω, can be combined to arrive at a single measure of manager achievement (either historical or expected), the pure information ratio: $IR = \alpha/\omega$.

The IR represents 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.[12]

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 would wish 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.

Level Playing Field for Evaluating Managers

Now that we have 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 important, 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 the 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 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 achieves). 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 it fails to beat the value benchmark. Which one should the manager really be held against? If a manager persistently chooses to exercise its expertise in one domain of market risk such as deep value, isn't that the domain against which its value-added should be measured?

WHAT ARE ACTIVE MANAGERS FOR?

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. It 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 strategic asset allocation 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:

$$E(U_{\alpha}) = E(\alpha_p) - \lambda_{\alpha} \cdot E(\sigma_A^2) \tag{1}$$

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_a 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 ω 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 farther, does one maximize expected utility? Through optimization.

Harry Markowitz [1952, 1991] developed this tool a half century ago. 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, optimization has just begun to be used for building efficient portfolios of *managers* within the investor's portfolio over the last 6 or 7 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 trade-off. Optimization is the technology that explicitly calculates these trade-offs in search of the highest-utility portfolio (of anything) for a given investor.

Increasingly, if they think 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, rather than looking at each manager in isolation. To accomplish this 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.

BUILDING PORTFOLIOS OF MANAGERS

Let's first hold up to the mirror some 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" (quoted in Bell [1952, p. ___]. But is this good enough?

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 depending 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.[13]

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 be true even 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, or, more precisely, to test whether one can conclude with reasonable confidence that any given alpha, positive or negative, has been achieved through skill rather than luck.

There is general agreement (with only minor dissent) 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.[14]

Here is why we digress on statistics. If a manager's historical alpha is not statistically significant in this way, 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 the alpha forecast for that manager according to a more qualitative or fundamental type of analysis, which of course can be done—it is just hard.

If a performance history does display statistical significance, it is fair to include it among one's other inputs when evaluating a manager. Yet one still doesn't just extrapolate it into the future without thought. A high t-statistic 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 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, 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. The key ingredient is the tough one—one has to have great insight and ability. This is 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.

Unfortunately, significant t-statistics on managers are quite rare, although possible—and this is exactly what theory predicts. 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 provide no useful information?

Style Boxes and Style Maps

A practice developed over the past couple of decades or so is to divide the equity universe into *style boxes* (such as large value, large growth, small value, small growth, and perhaps mid-cap or core categories) and then to staff each style box with managers. This practice seems useful for breaking 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 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 neatly in these boxes. Some are stronger (deep value) or weaker (growth at a reasonable price) than others. The amount of "valueness" or "growthiness" does not always come in units of one.

Most 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 large value.[15]

Particularly for traditional active managers, style exposures usually do 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 we aren't going to formally optimize as we suggest in this article, style maps and tables of the effective asset mix are the next best thing.

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 investors to just fill them up, without critically examining their own skill in assessing expected manager alpha or the relative expected alphas of the candidate managers. For investors who aren't strongly disciplined, style boxes and maps invite taking on active risk without a real expectation of pure active return.

Reliance on Traditional Active Management

As we've already noted, many 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 equate greater active risk with higher active return. Such investors often express disdain for risk-controlled active strategies and a preference for concentrated high-risk active strategies.

These 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 rewarded only conditionally, depending on skill.

Second, such investors don't realize that their optimal portfolio of managers, at whatever active risk level they are comfortable with, will be constructed from good low-risk active managers in preference to higher-risk managers, for two reasons. The first is that the mathematics of optimization dictate that an investor's manager allocation should be dominated by managers with 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 allocations to concentrated managers.

We conclude that sophisticated investors don't really want more risk-laden ventures. What they really want, or should want, is higher alphas, and less risk.

Current Practices Overall

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 risk to produce higher returns—do not generally support the investor's goal to add 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?

BUILDING PORTFOLIOS OF MANAGERS

Tomorrow's best practices have a lot to do with portfolio theory.

EXHIBIT 1
Efficient Frontier of Active Managers

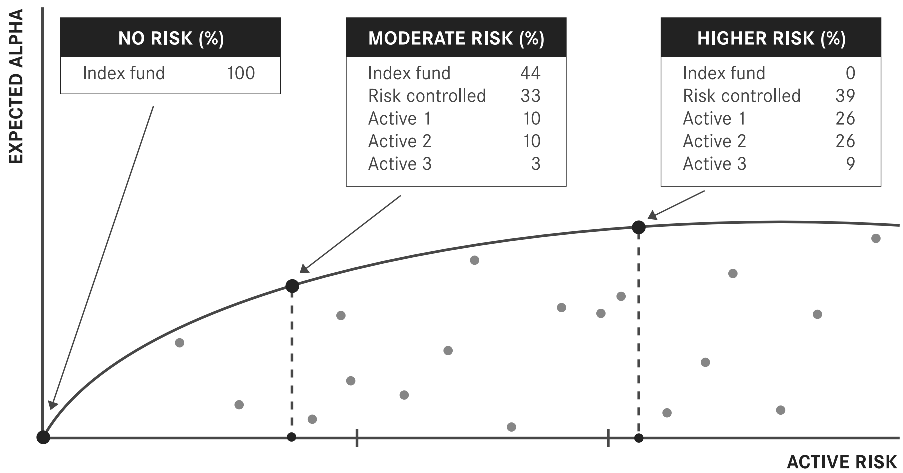

Treat Each Manager Like a Stock

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, with many advantages over traditional methods of constructing active stock portfolios.

We suggest treating each manager as one would a stock, giving investors the same powerful tools for managing their portfolios of managers that the best managers have for managing their portfolios of stocks. 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.

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 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.[16]

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.

This may look more daunting than the usual list of MVO inputs, but it is easier than it looks. The correlations of the pure alphas are often simply estimated as zeros for two reasons.

First, we have eliminated the market-related risk factors by calculating a regression alpha, the pure active return that is uncorrelated with the market factors. While this does not necessarily mean that all common factors have been removed so that the alphas are necessarily uncorrelated with each other, the alphas by virtue of the process will 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 these 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.[17]

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 accustomed 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).[18]

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

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. A more detailed example is in the appendix.

Expected Alphas

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 changes. The efficient market and zero-sum game alarm lights go on in their heads. They can't see themselves estimating alphas so that they can solve the optimization problem for managers in practice, and they mentally start moving back to their comfort zone—filling out style boxes.

But the 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 incorporate *implied* alpha forecasts, which, as we have pointed out, 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, 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. This is more honest and productive than simply selecting some active managers and hoping that the alpha forecasts implied by the holding weights turn out to be accurate.

Having done this, run an optimizer to select the portfolio. 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.[19]

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.[20]

Why Low Active Risk Managers Are Preferred

All these prescriptions embed the idea that active risk is in and of itself bad—that of two equally skillful managers (managers with the same information ratio), the lower active risk manager is to be preferred. But why? Grinold [1990] and Kahn [2000] analyze the utility function for active management and get the result:

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

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
- ω_{mgr} is the expected volatility of the manager's pure alpha around a properly established benchmark.

Expressed in words, the size of a manager's allocation is directly proportional to a higher expected information ratio and inversely proportional to a lower level of active risk. If this seems 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.

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*. You want managers with more alpha, and a lot less risk.

There is another effect to consider as well. It tends to diminish the extent of allocations to managers that have greater active risk.

Most managers are prohibited by their mandates from taking short positions, and thus hold long-only portfolios. This limits a manager's ability to make bets on its insights, and the limitation is exacerbated as active risk is ratcheted up and requires larger and larger active positions.

As a result, even with skill levels held constant, the

EXHIBIT 2
Impact of Long-Only Constraint on Portfolio Efficiency (assuming equal skill levels)

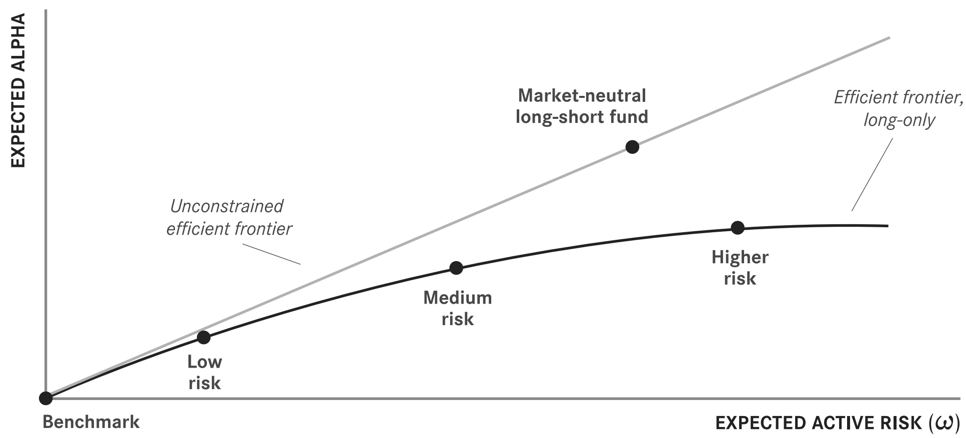

information ratio that can be achieved is reduced as a manager constructs a portfolio with greater concentration and active risk (see Grinold and Kahn [2000b]). Exhibit 2 illustrates this concept.

What are these limitations? A skillful alpha forecasting process is equally 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 be sold down from the benchmark weight to only a zero weight, and not below, no matter how strong the strength of the negative signal. The manager's insights are wasted, and the effect is greater as greater risk is taken on (larger 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 [2000] to have roughly twice the expected information ratios as their long-only, traditional active counterparts that have moderate risk.

How Much Active Risk is Enough?

Since pure active risk 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 greater than policy risk alone. For example, if policy risk is 9% and the active risk is 3%, the total risk will be

$$\sqrt{9\%^2 + 3\%^2} = \sqrt{0.009} = 9.5\%$$

So total risk goes up only by half a percentage point 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, determines how much risk an investor will take in search of the available expected return. We have to look at sources other than theory to get a proper sense for the appropriate ranges of active λ for investors.

One source to look at is our own behavior. According to Brinson, Hood, and Beebower [1986], Brinson, Singer, and Beebower [1991], and related studies, about 90% of the variance of a typical portfolio's returns is attributable to strategic asset allocation decisions (market risk); 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%, the investor has a 0.0009 active variance, lower than the policy variance in a 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.

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 goes back to our motivating theme: Policy bets are expected to be rewarded, unconditionally, and proportionally to risk taken. Active bets are rewarded only 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. Putting it this way, it 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 λ risk aversion term that is higher when we are optimizing in the active risk dimension than when we are optimizing in the policy or SAA dimension. Thus, in Equation (1), we indicate with the subscript α that the λ is specifically an active risk λ, different from the λ expressing aversion to policy risk.

At the end of the day, the specific λ used sets the active risk budget, and the active risk budget has to be a comfortable one to the investor, given that investor's perception of the alpha expected in return for taking on that risk.

Core–Satellite?

In the late 1980s many investors started using the *core–satellite* concept as 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 used by some even to justify holding extremely high-risk, concentrated managers as the satellites. This made some sense at the time, and represented a first stab at controlling active risk.

With today's technology, particularly the ability to manage and optimize active risks and returns, we can reject the notion of the core-satellite portfolio as too rigid. For the same reasons that we reject style buckets as not reflective enough of the continuous and scalar nature of managers' actual market exposures, we can also set aside the two buckets of core and satellite.

An optimally constructed portfolio of active managers is likely to also 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 with increasing risk, not in a discrete way. Further, the other active managers 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. The core-satellite idea 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.

WHAT DOES A PORTFOLIO THAT REFLECTS THESE PRINCIPLES LOOK LIKE?

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 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 rather good fuzzy optimizers above our shoulders.

For investors who use an actual optimizer and for those who don't, our basic points 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 *historical* 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.
   - 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 repre-

sent roughly one-third to two-thirds of the total fund. If an investor has little tolerance for active risk, this portion will lean more toward index funds. If there is more tolerance for active risk, it will lean toward risk-controlled active funds.

- 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 over higher-risk, concentrated managers. Bias toward diversified portfolios, away from concentration (unless completely carried away by the concentrated manager's extraordinary forecasting skill).
- Good market-neutral long-short funds will receive a substantial weight for investors not limited to long-only managers.[21]

3. Set the overall active risk at a comfortable level, your risk budget. For U.S. equities, a typical investor seems to be most comfortable (in our experience) at an overall active risk level of 1.5% to 2.0%, and the very largest investors prefer 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.

CONCLUSION: MANAGING IN THE RIGHT DIMENSIONS

Two first principles emerged decades ago from the basic academic work in finance. The first is Harry Markowitz's 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 an equity risk premium, and active risk, which is a zero-sum game when summed across managers and thus not rewarded on average.

We have added the observation that manager skill levels really do differ. 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 note that while the payoff to market-related (policy) risk is linearly related to the amount of policy risk taken, this 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. In the presence of skill, there is a payoff, declining as a proportion of pure active risk as the amount of risk taken increases (it does not decline for unconstrained, long-short strategies).

There are some 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, as uncorrelated with policy, can and should be managed through its own separate optimization process.

To do so, investors must estimate the expected alpha 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 (see Waring, Harbert, and Siegel [2001] and Waring et al. [2000]). First, one must believe that superior managers really do exist. This is easy, if one accepts that managers differ in their skill levels. Second, but harder, investors must believe that they can identify which ones will be the winners. To accomplish this, you need 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 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 naive 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 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 when it is complex and there are fine gradations. As Einstein is reported to have said, "Everything should be made as simple as possible but no simpler."[22]

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. 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, 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.

APPENDIX: MANAGER STRUCTURE OPTIMIZATION EXAMPLE

Our example below 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 demonstrate the point made in our discussion of core-satellite investing. The concentrated manager receives only a small allocation, contrary to some core-satellite interpretations.

To make the best possible case for choosing concentrated

Manager Assumptions

Manager	Expected Alpha	Expected Active Risk
Index Fund *(Russell 3000)*	0.00%	0.00%
Risk-Controlled Active Mgr *(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 *

Style Index	Expected Total Risk	Correlations 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%

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

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 Mgr	**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*

Sponsor benchmark: Russell 3000 broad capitalization U.S. equity.

managers, we use 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.

ENDNOTES

The authors gratefully acknowledge the assistance provided by Peter Bernstein, Bruce Goddard, Blake Grossman, Ronald Kahn, Kenneth Kroner, Gerry Rocchi, Clinton Stevenson, and Linda Strumpf. Many other very helpful comments were received as well from colleagues who shall remain nameless.

[1]While there hasn't been much academic research focused on managing active decisions, an authoritative, general, and up-to-date example is Grinold and Kahn [2000a]. The technology they describe can be recast in a form that applies not just to the security selection problem, where their work is focused, but also to the manager selection and structuring problem (see Waring et al. [2000]). Both these works provide a bibliography of the 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 set this discussion on the familiar turf of security selection, the observations apply 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 [1974].

[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]While the difference in returns between equity and fixed-income markets (the equity risk premium) may vary in size over time, logic largely dictates that it must always be positive for markets to clear. See Grinold and Kroner [2002] and Leibowitz et al. [2001].

[6]The Evans-Archer diagram [1968] is the classic tool used to introduce this point, showing that risk declines asymptotically to some irreducible amount as the portfolio becomes more and more diversified. Although Evans and Archer contrast only policy and active risk, without saying what risks are 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 capital asset pricing model, 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. Ross [1976] demonstrates in 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.

[7]This version of the argument is developed most eloquently in Sharpe [1991]. Managers sometimes argue that they are not on a par with non-professionals, that as professionals they can take advantage of other market participants even if the return across all participants must be a zero-sum game—and that therefore it makes sense to hire them. Given that professionals represent a large portion of the total market, it is difficult to imagine that this could be a large factor; their own substantial activities will discipline prices for themselves as well as for others.

We can find a little evidence to support the claim, but not sufficient to make their case usefully persuasive. Brinson, Hood, and Beebower [1986] and Brinson, Singer, and Beebower [1991] conduct a series of empirical studies of the returns achieved by professional investors (that is, pension plans and other institutions and their managers, and mutual funds). The data are slightly different in each study, but it isn't unfair to summarize the results as saying that active management underperforms its passive benchmark, after costs, by about 0.5 percentage point per year. Sharpe [1991] would justify a prediction that in the absence of a professional effect the average loss would be well over 1 percentage point per year, which is the estimated sum of manager fees, explicit trading costs, and market impact trading costs.

So, given this empirical difference, there is an implication that the professionals did in fact beat the other market participants over the time periods studied, on average—but not by enough to cover fees and costs or to give the professionals a return higher than that of their passive benchmark. For similar results and an updated interpretation, see Ibbotson and Kaplan [2000] and Surz, Stevens, and Wimer [1999].

[8]The degree of the component of variance that is explained by luck rather than skill is surprisingly high. Since the IC is just a correlation coefficient, for a process with correlated signals we can convert it to an R^2 by just squaring it (Grinold and Kahn [2000, p. 292]). A manager with a very respectable IC of 0.10 would have an R^2 of 0.01, meaning that just 1% of active variance is explainable by the manager's skill, the remaining 99% of active variance being just random luck. Small wonder that it is so difficult to separate skill from luck over short time periods.

[9]We aren't sure who was first to classify investments or managers on the basis of 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 return-based style analysis), and developed 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 return-based style analysis, one authoritative source is Fama and French [1993], drawing on earlier observation by Banz and Reinganum [19__] that small-cap stocks had historically outperformed larger-cap stocks and Keim and Basu [19__] that value stocks (with low price/book, price/earnings, or other valuation ratios suggesting that the stocks were cheap) had outperformed growth stocks (with high valuations).

[10]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 exposures to be between zero and one. We generally prefer ordinary regression, because it can more 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.

[11]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.

[12]In our experience, some investors don't find that the term information ratio conveys much about its meaning. Observing that IR measures the consistency with which the active return is delivered, perhaps consistency ratio would be a better term and more likely to be understood. 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

[13]See, for example, Kahn and Rudd [1995], who provide an extensive review of the literature to that time, which happens to include most of the relevant literature even today. Further work on the topic is in Wermers [2000] and Carhart [1997], 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 likely to apply in roughly the same way to institutional funds, which are largely managed by the same managers.

[14]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 are 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.

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

[16]To correct an error in the optimization formula (Equation (2) in Waring et al. [2000]), the first term, representing misfit return, should include a beta, as follows:

$$\left(h^T X - \beta h_b^T\right) r_k$$

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

[17]At least for U.S. equity managers. For international equity managers and for fixed-income managers, correlations between pure alphas are empirically non-zero (because these asset classes have fewer degrees of freedom; in fixed-income, there are only a few 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 in these asset classes.

[18]There are few factors if one uses style analysis, but not necessarily so few if one uses 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.

[19]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] responds 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.

[20]To get started in forecasting alphas, investors might consider using the forecasting relation:

$$\text{Alpha} = \text{IC}(\text{Volatility})\text{Score}$$

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 forecasted (volatility of the pure alpha); and score is the strength of the manager under evaluation, 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 right about managers only 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).

[21]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.

[22]In fact, there is no written record of this statement (see Neumann [1977]).

REFERENCES

Banz, Rolf, and Marc Reinganum.

Bell, Eric Temple. *Mathematics, Queen and Servant of the Sciences.* Mathematical Association of America reprint edition, December 1987.

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

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

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

Evans, L. John, and N. Stephen Archer. "Diversification and the Reduction of Dispersion: An Empirical Analysis." *Journal of Finance*, December 1968, pp. 761-767.

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

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

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

——. "The Surprisingly Large Impact of the Long-Only Constraint." *Investment Insights,* Barclays Global Investors, May 2000.

Grinold, Richard C., and Kenneth F. Kroner. "The Equity Risk Premium." *Investment Insights*, Barclays Global Investors, San Francisco, July 2002.

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

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

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

Keim, Donald, and Sanjoy Basu.

Kritzman, Mark P. "Value In – Garbage Out." *Economics and Portfolio Strategy,* Peter L. Bernstein, Inc., January 15, 2003.

Leibowitz, Martin L. *et al., Equity Risk Premium Forum*. Sponsored by the Association for Investment Management and Research*,* November 8, 2001.

Markowitz, Harry M. "Portfolio Selection." *Journal of Finance,* March 1952.

——. *Portfolio Selection: Efficient Diversification of Investments,* 2nd ed. Oxford: Basil Blackwell, 1991.

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

——. "The Markowitz Optimization Enigma: Is 'Optimized' Optimal?" *Financial Analysts Journal,* January/February 1989.

Neumann, Peter. "________________" *Reader's Digest,* October 1977.

Ross, Stephen A. "The Arbitrage Theory of Capital Asset Pricing." *Journal of Economic Theory,* December 1976, pp. 341-360.

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

——. "Asset Allocation: Management Style and Performance Measurement." *The Journal of Portfolio Management*, Winter 1992, pp. 7–19.

——. "Determining a Fund's Effective Asset Mix" (corrected version). *Investment Management Review,* November/December 1988, pp. 59–69.

——. "Imputing Expected Security Returns from Portfolio Composition." *Journal of Financial and Quantitative Analysis,* June 1974, pp. 463-472.

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

Waring, M. Barton, Lee D. Harbert, and Laurence B. Siegel. "It's 11 p. m.—Do You Know Where Your Employees' Assets Are?" *Investment Insights*, Barclays Global Investors, October 2001.

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

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

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