InvestmentInsights

The Investment Research Journal from BARCLAYS GLOBAL INVESTORS

12.02

DECEMBER 2002
VOLUME 5 ISSUE 5

THE DIMENSIONS OF ACTIVE MANAGEMENT

Why Alpha and Active Risk are the only things that matter

by BLAKE GROSSMAN, BARTON WARING *and* LAURENCE SIEGEL

EXECUTIVE EDITORS

Barton Waring
415 597 2064 phone
415 597 7168 facsimile
barton.waring@barclaysglobal.com

Gail Radzevich
415 597 2073 phone
415 717 9581 facsimile
gail.radzevich@barclaysglobal.com

EDITOR

David Kurapka
415 597 2705 phone
415 618 1434 facsimile
david.kurapka@barclaysglobal.com

CONTRIBUTING WRITERS

Blake R. Grossman

Global Co-Chief Executive

Blake oversees the strategic development of BGI's global business, a role he shares with Andrew Skirton. He also has primary responsibility for BGI's active and enhanced investment strategies globally and the institutional businesses based in the US. Blake was previously Global Chief Investment Officer with overall responsibility for the investment management, research and business activities of the Advanced Active Strategies Group, which is the unit focused on BGI's systematic active and enhanced strategies globally. He joined BGI in 1985, with responsibility for the launch and management of the US Alpha Tilts enhanced equity funds. He graduated Phi Beta Kappa from Stanford University, where he received BA and MA degrees in economics.

M. 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 B. Siegel

Director, Investment Policy Research, Ford Foundation

Guest co-author Laurence B. 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 Ken Kroner.

THE DIMENSIONS OF ACTIVE MANAGEMENT

Table of Contents

IN THIS ISSUE

Executive summary 2

Introduction 4

Active management 5

Building portfolios of managers: Current practices and their costs . .10

Building portfolios of managers: Best practices 13

What will a portfolio that reflects these principles look like? 16

Conclusion: Back to basics 19

Endnotes 21

SIDEBAR

You don't really expect me to forecast manager alphas, do you? 17

The end of the debate between index and active 18

EXECUTIVE SUMMARY

Just about all of us in the investment management business either hire, or are, active managers. But what exactly is the role of an active manager in the plan sponsor's larger portfolio? Can plan sponsors measure the risk and return of active managers in a way that is fair, and at the same time useful for building portfolios from the selected manager array? Are current practices for building portfolios of managers the best that can be designed, or is there a better way?

This issue of *Investment Insights* examines these questions by exploring the real "dimensions" in which active managers vary. We discuss how building a portfolio of managers, which is the plan sponsor's job, is like building a portfolio of anything: it's about maximizing return subject to a penalty for risk. In other words, it's an optimization problem.

We hope to convey two fundamental insights. The first is that one hires active managers to generate "pure" active return, knowing that in their efforts to do so they will generate "pure" active risk. These are the real dimensions of active management that must be managed. By defining the dimensions of active management in this way, in contrast to style boxes and other devices, one creates a "level playing field" that enables investors to compare the contributions of managers with different levels of risk and return or even in totally different strategies and asset classes. The use of these precise measures makes it possible to construct more sensible portfolios of managers, portfolios that are more likely to "deliver the goods."

The second insight is that in building a portfolio of managers one is concerned with maximizing total portfolio return, and with controlling total portfolio risk. Fortunately, one can parse this problem into parts that make it easier to solve and to construct useful estimates or inputs. By breaking total risk into its elemental pieces—policy or market risk (controlled by the strategic asset allocation policy and the asset class benchmarks), and active or idiosyncratic risk (which arises from implementation through active managers, and which is generally uncorrelated with policy risk), one can use optimization techniques to solve the policy construction problem in practice.

Existing practice for building portfolios of managers is not ideal. Historical performance comparisons, style boxes, and heavy reliance on traditional active management—the cornerstones of current practice—are not only costly and inefficient but deflect attention from the real problem of maximizing active return and controlling active risk.

Optimization, in which expected active return is formally or informally balanced against expected active risk, is a better way to form a structure of good managers. Long used to build portfolios of stocks, and to build whole investment programs out of asset-class benchmarks, optimization can easily be adapted to the problem a sponsor faces when assembling a portfolio of active managers. 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.

There is no doubt that active management is difficult, and only generates a positive expected return when and if skillful managers are skillfully chosen. Plan sponsors who don't have or can't develop the skill to estimate expected alphas for managers should just hold index funds and not play the active game.

However, most sponsors do hire active managers, asserting (at least by implication) that they have skill at manager selection. For such sponsors, the following manager structure rules of thumb can inform one's intuition:

1. The best portfolio is one that balances active risk and active return summed across managers, taking into account their correlation structure and keeping overall active risk at a comfortable level. For US equities, a typical sponsor seems to be most comfortable (in the experience of the authors) at an overall active risk level of 1.5–2%, with the largest sponsors preferring less active risk (between 0.75% and 1.25%).
2. The "core" of a portfolio consists of some combination of risk-controlled active funds and traditional index funds, which together likely will comprise roughly one-third to two-thirds of the total fund. This core is the principal source of risk control for the whole portfolio.
3. Risk-controlled active funds will be more heavily weighted than equally skillful traditional active funds.
4. Equally skillful traditional long-only active management will be less heavily weighted than lower risk managers, and will become even less so as the active risk level goes even higher (no highly concentrated portfolios). Hold strong active managers with a bias towards more diversified rather than more concentrated managers.
5. While a core and satellite structure may be a helpful way to think about manager structure, core and satellite does not mean core and concentrated active. Concentrated active means high risk, but instead you want a combination of low risk and high expected information ratio.
6. Market-neutral long-short funds, if skillfully designed and managed, will receive a substantial weight, particularly at higher risk levels of sponsor risk tolerance. If you have identified skillful managers in this category, hold enough of them to make a difference.

To implement an optimization approach, one must isolate the pure alphas of managers, which means controlling for risk factors such as the market factor and style exposures. Each manager has his or her own unique set of style exposures, measured along a continuum of growth to value, and large to small capitalization. Style "buckets" make no sense, since they clump together managers whose styles typically have only the most superficial resemblance to one another.

Thus, plan sponsors should estimate the style exposures of their managers, make alpha forecasts, and then run an optimizer. If sponsors prefer a more informal approach to building an efficient portfolio of managers, they should at least think about manager structure as an optimization problem. The sponsors should then hold a portfolio they can defend as maximizing expected return while holding expected risk to a level with which they are comfortable.

Introduction

Just about all of us in the investment management business either hire, or are, active managers. But what exactly are active managers? What is their role in the sponsor's larger portfolio? Can one measure the risk and return of active managers in a way that is fair, and at the same time useful for building portfolios from the selected manager array? How? 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 think they know the answers to at least a few of these questions: Active managers select securities. They are there to add "alpha," an incremental return over their benchmark. But by looking at these questions a little more critically, and by applying insights from technologies—such as optimization—that are central to the science of investment management, one can gain additional perspectives that have the potential to improve returns, control risk, and help investors build better portfolios.

In this article, we hope to convey two fundamental insights. The first is that one hires active managers to generate "*pure" active return*, knowing that in their efforts to do so they will generate *pure active risk.* We use the term "pure" active return to signify that we're carefully differentiating between market returns and active returns in a regression sense. Many of our readers will see that we're referring to an alpha sometimes described as "selection" alpha or "specific" alpha, but the bottom line is that it is the "real" or "pure" alpha. These are the real dimensions of active management that must be managed. By defining the dimensions of active management in this way, in contrast to the conventional dimensions implied by some consultants and many sponsors' style boxes and other devices, one creates a "level playing field" that enables investors to compare the contributions of managers with different levels of risk and return or even in totally different strategies and asset classes. The use of these precise measures makes it possible to construct more sensible portfolios of managers, portfolios that are more likely to "deliver the goods."

The second insight is that building a portfolio of managers is just like building a portfolio of anything: one is concerned with *maximizing total portfolio return*, and with *controlling total portfolio risk*. In other words, it is an optimization problem. Fortunately, one can parse this problem into parts that make it easier to solve and to construct useful estimates or inputs. By breaking total risk into its elemental pieces—*policy or market risk* (controlled by the strategic asset allocation policy and the asset class benchmarks), and *active* or *idiosyncratic risk* (which arises from implementation through active managers, and which is generally uncorrelated with policy risk), one can use optimization techniques to solve the policy construction problem in practice.

Along the way, we will point out that existing practice for building portfolios of managers is not ideal. Historical performance comparisons, style boxes, and heavy reliance on traditional active management—the cornerstones of current practice—are not only costly and inefficient but deflect attention from the real problem of maximizing active return and controlling active risk. Optimization, in which expected active return is formally or informally balanced against expected active risk, is a better way to form a structure of good managers. While familiar as a general tool for balancing risk and return, a slightly modified optimization approach can be used specifically to address the problem of active manager selection.

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 job of managing that tension, and doesn't even explicitly focus on it, but there are new technologies and ways of thinking that do.

Active management

Active managers, of course, 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 make stock-by-stock forecasts of expected alpha, and then translate these forecasts (plus information about risk and correlation) into portfolios.[1]

Well, not exactly. In practice, most active managers don't make explicit forecasts of stock-specific alphas. Instead, they may establish "price targets" or express forecasts in some other numerical space that does not translate directly to expected alphas. Most, of course, simply assemble their portfolios with no formal views other than that the stocks they have chosen will go up.

What are active managers? What are they for? How can we properly measure the risks they take?

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."[2] For traditional active managers with a limited number of holdings, we have observed that implied expected alphas are surprisingly large for the stocks that are held, and often also surprisingly large (but negative) for stocks with a large weight in the benchmark that are *not* held.

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

Likewise, the sponsor has made alpha forecasts for the *managers* it holds. These forecasts are rarely explicit in current practice, but they are always implicitly there. These can also be backed out through reverse optimization. Sponsors, too, should only hire active managers if they have believable positive alpha forecasts for the managers. 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.

POLICY 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 the Nobel Prize-winning economist William F. Sharpe, we note that the total risk of any investment can be broken into: (1) *policy risk*, also called market, systematic, undiversifiable, or beta risk; and (2) *active risk*, also called specific, unsystematic idiosyncratic, diversifiable, or alpha risk.[3] Policy risk is usually managed by sponsors through their adoption of a strategic asset allocation policy, with its attendant asset-class benchmarks. Active risk comes into the sponsor portfolio through the use of active managers.

Exhibit 1 is a schematic diagram showing the separation of policy and active risk. Business schools teach that policy risk is the only risk that is worthy of consideration, on the ground that if markets are efficient active risks are just random bets that cancel each other out. In practice, however, asset classes are "staffed" with active managers, whose returns can differ widely from benchmark returns, so active risk needs to be managed as well.

Since investors can participate in the overall profits of the corporate sector of the economy simply by buying the market, they are rewarded for taking policy risk (on average over the long run).[4] Active risk, in contrast, cannot be rewarded *on average* since active managers are betting against each other in a zero-sum game.[5] 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,[6] so they must earn the return on the market—minus fees, transaction costs, and other expenses, which can be substantial.[7]

Does this mean that sponsors should only index, shunning active managers entirely? Of course not. As long as the market is not completely efficient (and we believe that it is not) and differences in human intelligence and skill levels exist, some managers will outperform through real skill, not just random variation. The search for such managers is, arguably, the sponsor's highest calling.

DEFINING "PURE" ACTIVE RETURN AND RISK: WHAT MANAGERS SHOULD BE HIRED TO DELIVER

Having defined active risk and return, we now must go one step further by separating these "gross" components into finer parts: misfit risk (and return), and pure active risk (and return).

Managers' active returns are not uncorrelated with one another, but have many common factors that go a long way toward explaining the active returns achieved. The most intuitive of these factors are

Exhibit 1
SEPARATION OF POLICY AND ACTIVE RISK

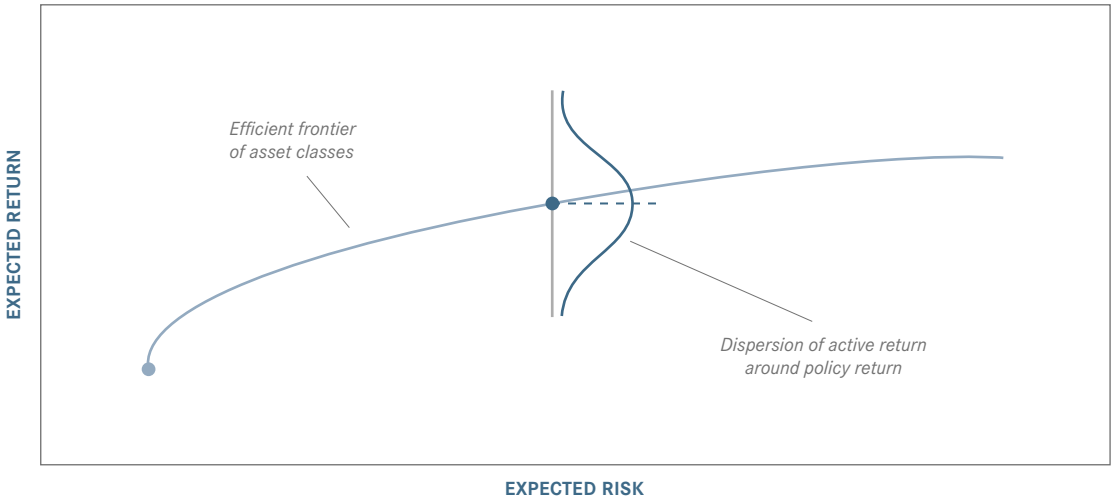

called *styles*—the familiar large-capitalization, small-capitalization, value, and growth categories.[8] Additional factors have been identified that more completely explain the variation in manager returns, but some of these are hard to describe in plain English so we'll stick with large, small, value, and growth for the moment.[9]

The existence of styles—and style index funds, which deliver the returns on the styles at very low cost—implies that active managers should not get "credit" for returns that are caused merely by the manager being exposed to a given style or mix of styles. If a sponsor just wants a particular mix of styles, he or she can buy the style index funds. The value added by the manager should, instead, be regarded as the part of return that is *beyond* that delivered by the manager's style or mix of styles. This part of the return, that is the unique contribution of the manager, is the manager's *pure active return*, or *pure alpha*. Note that an investor can achieve market returns quite inexpensively with index funds. You don't need to pay an active manager to give you market exposure, only to try to *beat* the return of the market.

Historical pure alpha is measured by regression analysis.[10] The regression determines the effective style mix of the manager, or mix of style benchmarks that has the "best fit" to the manager's actual returns. This best-fit benchmark may be regarded as the manager's individualized or custom benchmark. The pure alpha is then calculated as the manager's actual return in excess of the return on the custom benchmark.[11] It is only this pure alpha for which a manager should be hired (and paid an active fee), since the custom benchmark can be built out of index funds at low cost.

Pure active return, of course, is properly measured *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.

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.[12] These two parameters,

- Pure active return or pure alpha, α
- Pure active risk, σ_α

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

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

representing the amount of active return delivered (or expected) per unit of active risk taken (or expected). The term "information ratio" isn't very intuitive, so we'd like to note that *IR* measures the consistency with which the active return is delivered; perhaps "consistency ratio" is a better term for this concept, and one more likely to be used.

Style is scalar. The view of style that we've been describing, by the way, is scalar—that is, a manager can have any amount of exposure to a style benchmark or to multiple style benchmarks. That's the way style works in the real world. In contrast, the style-box world pretends that style is a binary, in-or-out, yes-or-no characteristic; each manager fits into a given style box, or else fits into no box and gets excluded from searches. We discuss the problems with style boxes, and the usefulness of a scalar view of investment style, in greater detail in beginning on page 10.

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) to compare any manager with any other, across asset classes, styles, and risk levels, creating a "level playing field" for all managers. Of any two managers, even if in different asset classes or displaying different levels of active risk, the one with the higher expected or forward-looking information ratio is the better manager.[13]

Even more importantly, we can use these measures to properly separate investment results that are the sponsor'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 *sponsor*, and have nothing to do with the manager. The pure alpha should be credited to the *manager*, without regard or reference to the performance of the benchmark.

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

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 sponsors 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.

At this point we hope we've achieved the first of our two conceptual goals: to define the dimensions of active management as expected pure active risk and return—in contrast to the conventional dimensions implied by style boxes—and to begin to show how these real dimensions can be used to measure and compare managers scientifically.

WHAT ACTIVE MANAGERS ARE FOR: MAXIMIZING RETURN MINUS A PENALTY FOR RISK (MAXIMIZING EXPECTED UTILITY)

So why hire active managers? They provide the possibility of adding active return, of course, but they also add risk (that is, they add tracking error, which is a component of total risk). Exhibit 2 shows how policy and risk combine to form a portfolio's total risk; since active risk is typically uncorrelated with policy risk, the relationship is Pythagorean, so total risk is less than the ordinary sum of policy and active risk but it is still larger than policy risk alone.

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.

In general, the expected utility (or usefulness, or desirability) of an active portfolio is equal to its expected alpha, minus a risk penalty for active risk:

$$E(U_\alpha) = E(\alpha_p) - \lambda \times E\sigma_\alpha^2)] \qquad 1)$$

where

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

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

λ is the risk aversion parameter for the investor, or rate at which he or she translates risk into disutility;

and $E\sigma_\alpha^2$ is the active risk (expected variance) of portfolio p.[14] (Variance is the square of standard deviation, σ.)

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 is adding enough expected alpha to justify the extra risk? How,

Exhibit 2
THE RELATIONSHIP BETWEEN POLICY AND ACTIVE RISK: IT'S A PYTHAGOREAN THING

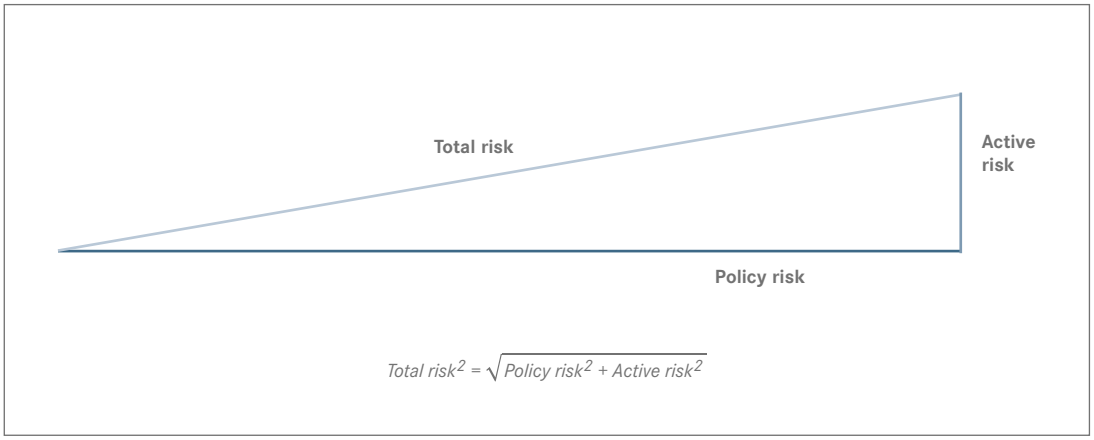

taking this one step further, does one *maximize* expected utility? Through optimization. Although Harry Markowitz created this tool a half century ago, and although it's the dominant practice for building efficient portfolios of asset classes and, increasingly, of securities within an asset class,[15] optimization is just beginning to be used for building efficient portfolios of *managers* within the sponsor's portfolio.

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.

Now we hope that we've achieved the second of our two conceptual goals, to show that building a portfolio of managers is an exercise in utility maximization (that is, maximizing return while controlling risk) and therefore an optimization problem. Later, when we discuss manager structure optimization, we'll show how pure active risk and pure active return fit into the optimization framework, bringing together our two sets of insights.

Building portfolios of managers: Current practices and their costs

Before going more deeply into our prescriptions for better approaches, let's note the chief features of current practice for building portfolios of managers:

- Intense focus on historical performance comparisons
- Style boxes as a tool for making the manager selection problem more manageable
- Heavy use of traditional active management

Each of these practices has flaws that contribute to the building of portfolios that are not only suboptimal in terms of investment performance and risk management, but costly as well.

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. Despite massive evidence that past performance is at the very best a weak predictor of future results,[16] plan sponsors devote great effort to analyzing it.

That's probably because analysis of past performance is something that sponsors can do with a sense of objectivity and confidence. It's hard to argue with actual historical returns.

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 sponsor 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.

A simple statistical test, the t-test, tells us whether a set of historic alphas from a given manager are different enough from the benchmark to provide real evidence of skill. The t-test measures the number of standard deviations, adjusted for the length of the time period, that the historical return on a fund differs from that of the benchmark.[17]

Statisticians generally agree that if a manager's performance is more than two standard deviations above that of the benchmark (that is, its t-statistic

is greater than two), then we can say the manager has a "statistically significant" alpha. This simply means that there is a very high probability (more than 97.5% if alphas are normally distributed) that the manager was skillful rather than lucky in the period over which the test was conducted. A high likelihood of true skill, of course, is not proof (which is impossible) and doesn't necessarily justify projecting that history into the future—but we can at least contemplate the possibility that the historical alpha will be repeated to some degree if the t-statistic of the alpha exceeds two and if the process that was used to create that alpha remains in place.

Unfortunately, significant t-statistics on managers are quite rare—although such rarity is exactly what theory predicts. Indeed, BGI enjoys significant t-statistics in the majority of its active strategies. 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).

Without a strong t-statistic for a particular manager, it makes no sense to even consider whether the manager's historical performance will repeat going forward. The sponsor will have to make his or her alpha forecast based on a more qualitative or fundamental type of analysis.

If one cannot usually rely on past performance to select active managers, then how can one 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 plan sponsor 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 skill. 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," beginning on page 15, we touch on other issues related to forecasting alphas for managers.

If one cannot usually rely on past performance to select active managers, then how can one select them at all?

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.

Since a cap-weighted combination of all the style boxes gives you back the market portfolio, there's nothing really wrong with style boxes *per se*. However, the implementation of the style-box system is usually badly flawed.

First, the system of style boxes lacks sufficient granularity to fully reflect the style characteristics of actual managers. As we've already noted, style boxes encourage investors to think of style as a binary, yes-or-no characteristic: a manager is either a value or a growth manager; a stock is either a value or a growth stock. Realistically, however, style is a continuum, or "scalar," and the style exposures of any given manager (or

stock) are best expressed as a blend—with a growth manager consisting of, say, 70% in large growth, 15% in large value, 10% in small growth, and 5% in large value.[18]

Moreover, sponsors who use style boxes often unwittingly pretend that all managers in a given style box have exactly one unit of exposure to that style. That, too, is unrealistic: a concentrated value manager might have twice as much exposure to the "value factor" as a broadly diversified value manager.

Finally, although any style box could be filled using an index fund that tracks the style, such funds are rarely used, with traditional (high-risk and high-cost) active managers usually being selected instead. Little attention is paid to active risk in this approach, which, as we demonstrated earlier, is a hurdle that active managers should be expected to clear (by providing a lot of alpha to compensate for the added risk) to meet the criterion of adding utility at the total portfolio level.

Style maps. Rather than simply using boxes to represent styles, some consultants and sponsors 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. The sponsor's total portfolio is also mapped, so that the misfit risk of the fund becomes visually obvious. One can also simulate changes in managers, or in manager weights, to see the effect on misfit risk and other portfolio characteristics.

Style maps are better than style boxes, since they recognize the scalar nature of styles and the importance of managing misfit risk. Nevertheless, they do not directly address the fundamental problem of active management, which is to *maximize expected active return* subject to a penalty for active risk. For that, one must use some kind of optimizer, a topic to which we turn later. It isn't enough simply to use a "completion fund" to gain exposure to styles not fully captured by the chosen active managers; we have to be willing to take some misfit risk, balancing this risk against the expected active return from the managers, if we are to have a utility-maximizing portfolio of managers. This calculation can be done only with an optimizer.

Style maps are better than style boxes, since they recognize the scalar nature of styles and the importance of managing misfit risk.

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

As we've already noted, many of the sponsors 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 sponsors often express disdain for risk-controlled active strategies, saying, in effect, "I want more juice."

Such sponsors 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. Remember that active risk is, on average, a zero-sum game so you can't expect outperformance simply by choosing to go with active management. Active risk will only be rewarded if you have real skill in picking the managers.

Second, sponsors fail to realize that, as Ronald Kahn and Richard Grinold have shown, a sponsor's manager allocation should be dominated by managers who have a combination of the highest forecast information ratios and the lowest levels of active risk.[19] Moreover, 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.[20] These considerations speak quite clearly to the opposite of the "more juice" argument so often heard. We discuss them in detail beginning on page 16 in "What will a portfolio that reflects these principles look like?"

CURRENT PRACTICES: CONCLUSION

The conventional dimensions of active management, as defined by historical alphas, style boxes, and other labor-saving but ultimately harmful devices, describe a nonexistent world that is almost exactly the opposite of the finely filigreed, quantitatively determined world of real active managers. Yet the conventional dimensions are invoked to make investment decisions that affect millions of plan participants and trillions of dollars. The portfolios thus constructed just *have* to be suboptimal. We would like to change that.

Building portfolios of managers: Best practices

Increasingly, sponsors are beginning to think of optimization as the basis for their approach to building *portfolios* of managers, in place of the conventional approach of looking at each manager in isolation. Long used to build portfolios of stocks, and to build whole investment programs out of asset-class benchmarks, optimization can easily be adapted to the problem a sponsor faces when assembling a portfolio of managers—the unit of selection that we're dealing with in this essay. 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.

FOUNDATIONS OF OPTIMIZATION

First, let's briefly review what makes optimization such a valuable tool in the investment world.

Maximizing utility involves trade-offs between risk and return that are complex and that take advantage of the correlation structure of every asset with every other. For example, a high expected return on a given security causes it to be attractive, and to have a relatively large weight in the portfolio, even if its risk or its correlations with other securities are not all that low. Likewise, a security will be attractive if it has very low correlations with other securities, diversifying the portfolio effectively, even if its expected return is not all that high. Optimization is the *only* method available for resolving all these trade-offs at once.

As most readers will recall, the inputs to a generalized mean-variance optimization (MVO) problem are (this fits for both policy decisions and manager structure or decisions:

- Expected return on every security, manager or asset class
- Expected risk (standard deviation) of every security
- Expected correlation of every security with every other

The output is a set of efficient portfolios, each of which represents the combination of securities that gives the highest expected return at each

level of risk. The locus of all efficient portfolios is called the efficient frontier, the shape of which is familiar to most investors. The question of which efficient portfolio to hold is decided according to the utility function, or risk tolerance, of the investor.

FROM SECURITIES TO MANAGERS: TREAT EACH MANAGER LIKE A STOCK

Now, let's implement the optimization technique using managers as the "securities" or "stocks" from the above formulation. As we noted earlier, one can use regression analysis to decompose the risks and returns of the various candidate managers into policy, misfit, and pure active components, making an otherwise unwieldy optimization problem doable.[22] The pure active returns and risks, taken by themselves, likewise describe a pure-active efficient frontier, similar to the generalized one and shown in Exhibit 3. (Note that the upward slope of the curve is conditional on the managers being skillfully selected.)[23] The pure active returns, or pure alphas, of managers tend to be uncorrelated with each other, since the common factors have been removed (the policy and misfit returns capture these factors) and because these pure alphas arise from the manager's unique approach or set of skills.

Thus, following Waring et al,[23] who call their method "manager structure optimization (MSO) |in homage to Markowitz's MVO, we can set up an optimization problem with a greatly reduced set of inputs. Rather than estimate the correlation of each manager with every other manager, as would be required in MVO and which would be nearly impossible, we need only to estimate:

1. Expected return and risk of each factor, and the expected correlation of each factor with every other factor
2. Factor loadings of each manager on each factor
3. Expected pure active return (alpha) and pure active risk of every manager

Exhibit 3
EFFICIENT FRONTIER OF ACTIVE MANAGERS

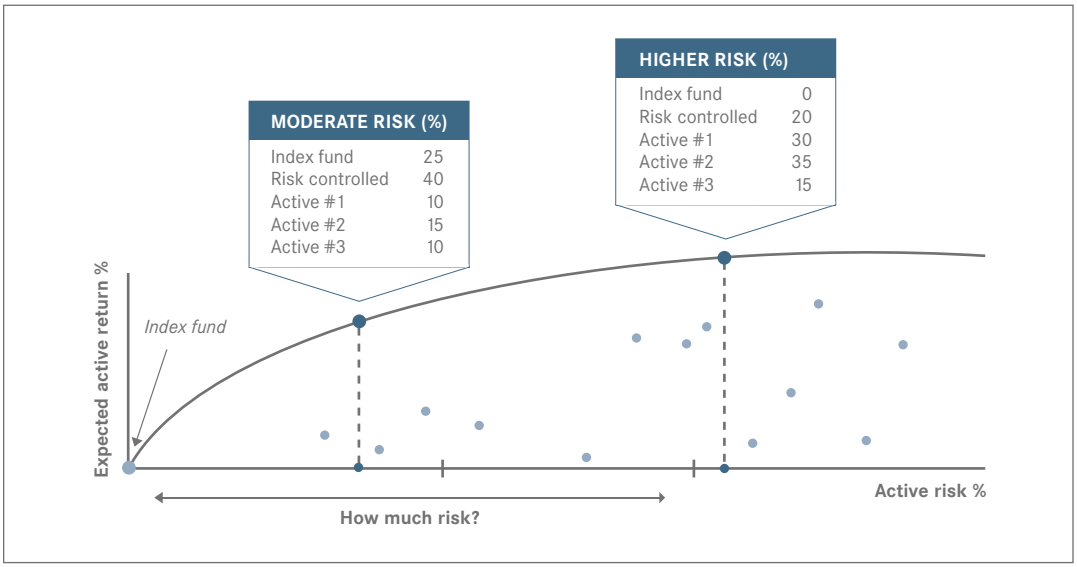

4. Expected correlation of every manager's alpha with that of every other manager and with each factor

Whew! This looks more daunting than the usual list of MVO inputs. However, it's actually much less so. The first two items are not difficult because there are so few factors,[24] and the fourth item is typically a matrix of zeroes (since managers' pure alphas are uncorrelated with the factors by construction, and are also generally uncorrelated with one another).[25]

The difficult part is the third item, and in particular the return component of it, expected alpha. We'll come back to this in a minute. Exhibit 3 is a representative example of an optimizing across managers in the active dimension. Notice how, 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 increases.

THOSE PESKY EXPECTED ALPHAS

Estimating the expected alpha—of anything—is the hardest problem in finance. Yet to justify holding active managers, sponsors must have such forecasts, and must be able to defend them with intellectual integrity. (In any portfolio that includes active managers, the forecasts are there, implied by the holdings weights, whether or not they are explicitly acknowledged.)

In our experience, most sponsors quail at the task of making specific, numerical estimates of the expected alpha for each manager. Sponsors typically defend their reluctance to make such forecasts on the ground that they've been taught that markets are efficient and that active management is a zero-sum game. Yet they almost always hire traditional active managers. Maybe being required to make alpha forecasts, for use in an optimizer, encourages sponsors to think more clearly about whether they have real evidence that a given manager is going to beat his or her properly style-adjusted benchmark. We strongly believe that such managers exist, but no investor should assume that a given active manager is going to be a winner just because he or she is active. It takes real insight about the manager to justify an expectation that he or she will add value to the portfolio, other than by mere chance.[26]

Estimating the expected alpha—of anything—is the hardest problem in finance. Yet to justify holding active managers, sponsors must have such forecasts, and must be able to defend them with intellectual integrity.

Order-of-magnitude forecasts. At the very least—and we think this is asking too little of sponsors—one could run an optimization using order-of-magnitude alpha forecasts. Managers can be grouped into, say, five buckets: exceptional, good, average, bad, terrible. An expected alpha or expected information ratio can be assigned to each bucket. (Only managers in the first two buckets are going to become candidates for hiring, of course.) One must also generate the rest of the estimates—active risk, factor loadings, and so forth. A portfolio of managers that is built using these rough estimates in an optimizer is likely to

have a higher risk-adjusted return (and handle misfit risk better) than a portfolio of managers that is built entirely without the benefit of optimization.[27]

Are optimizers "error maximizers"? There is a thread of thinking in the investment profession, reflected in the views of some sponsors and consultants, that says that optimization is an "error maximizer" that turns necessarily imprecise inputs (statistical estimates of expected return, risk, and correlation) into outputs (portfolio weights) that are just as imprecise as the inputs and potentially more so. This critique of optimization is technically correct.[28]

What will a portfolio that reflects these principles look like?

The methods and modes of thinking that we've outlined here all point toward the conclusion that active management is difficult, and only generates a positive expected return when and if skillful managers are skillfully chosen. We can state the negative case very clearly:

- If you don't have or can't develop the skill to estimate expected alphas for managers, you should just hold index funds, not playing the active game.

The best portfolio is one that balances active risk and active return summed across managers, taking into account their correlation structure and keeping overall active risk at a comfortable level.

However, style boxes and other heuristics have the same problem, only much worse. Visualize style boxes, for a moment, as botched optimizations (with expected alphas and so forth as the inputs, and portfolio weights as the outputs). Clearly many of the inputs—identical style exposures for all managers within a style box, and inordinately high alphas for managers that are selected—are ridiculous. There is nothing about style boxes that makes bad inputs into good outputs, and there is nothing about optimization that does that either. You have to be careful when estimating inputs for any decision-making tool, and you have to acknowledge their imprecision.

What tilts the balance in favor of optimization, then, is the lack of a suitable alternative. We are reminded of Winston Churchill's observation that "democracy is the worst form of government, except [for] all those other forms that have been tried."[29] The same might be said of optimization: it remains our most useful tool for balancing expected returns against expected risks.

However, most sponsors do hire active managers, asserting (at least by implication) that they have skill at manager selection. For such sponsors, the following can inform one's intuition, whether one actually uses an optimizer or just thinks about manager selection and structure as a conceptual optimization problem:

1. The best portfolio is one that balances active risk and active return summed across managers, taking into account their correlation structure and keeping overall active risk at a comfortable level. For US equities, a typical sponsor seems to be most comfortable (in the experience of the authors) at an overall active risk level of 1.5–2%, with the largest sponsors preferring less active risk (between 0.75% and 1.25%).
2. The "core" of a portfolio consists of some combination of risk-controlled active funds and traditional index funds, which together

likely comprise roughly one-third to two-thirds of the total fund. This core is the principal source of risk control for the whole portfolio.

3. Risk-controlled active funds will be more heavily weighted than equally skillful traditional active funds.

4. Equally skillful traditional long-only active management will be less heavily weighted than lower risk funds, and will become even less so as the active risk level goes even higher (no highly concentrated portfolios). Hold strong active managers with a bias towards more diversified rather than more concentrated managers.

5. While a core and satellite structure may be a helpful way to think about manager structure, core and satellite does not mean core and concentrated active. Concentrated active means high risk, but instead you want a combination of low risk and high expected information ratio.

You don't really expect me to forecast manager alphas, do you?

In our experience, practically everyone sees our point that building portfolios of managers is an optimization problem—in theory. They follow along and nod their heads in agreement. But when they find out they're going to have to estimate expected alphas for managers, they freeze up. The efficient-market and zero-sum-game alarm lights go on in their heads, and they can't see their way to estimating alphas so that the optimization problem for managers can be solved in practice.

But sponsors—the same ones who object to estimating expected alphas on efficient-market and zero-sum-game grounds—almost always hire active managers. The resulting portfolios contain *implied* alpha forecasts, which can be calculated through reverse optimization. Thus, investors are forecasting expected alphas whether they resist the notion or not!

How can one reconcile these conflicting impulses? If you don't think you can forecast alphas, 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. 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.

The end of the debate between index and active

While active managers typically see index funds as a threat—and characterize the discussion of the relative merits of active versus indexed as a debate that should have a clear winner or a clear loser—the reality is entirely different. This framework for evaluating manager structure reveals that there is an important and positive role not only for the truly superior active managers, predictably capable of adding value to a portfolio, but also for index funds that make no effort to do so. The relative proportions of active and passive managers are driven not by the merits of the old arguments on an "either-or" basis, but instead simply by the investor's desire to control the maximum amount of active risk it is willing to accept in its portfolio. It is a classic risk budgeting question.

The twenty-seven-year-old war begun by the indexing revolution is reconciled in this framework, and the investor comes out as the clear winner, with an improved set of tools for risk control, return control and cost control. This approach integrates managers of all levels of active risk, including both indexed and active managers, into appropriate roles in the overall portfolio. The bottom line is investor-focused, not manager-focused. Active and passive funds should be mixed together as appropriate to meet the risk budgets and to generate the return expectations of the investor. This approach benefits the investor, not the protective concerns of active and index managers.

6. Market-neutral long-short funds, if skillfully designed and managed, will receive a substantial weight, particularly at higher risk levels of sponsor risk tolerance. If you have identified skillful managers in this category, hold enough of them to make a difference.[30]

WHY LOW ACTIVE RISK MANAGERS ARE PREFERRED

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? Because, as Grinold and Kahn have shown, the weight that one should hold of a manager is inversely proportional to that manager's active risk:[31]

$$h_p \sim IR_p / \sigma_\alpha \qquad 2)$$

where

h_p is the holdings weight of portfolio or manager p,

IR_p is the information ratio of manager p, defined as σ_α

Exhibit 4

THE IMPACT OF THE LONG-ONLY CONSTRAINT ON PORTFOLIO EFFICIENCY

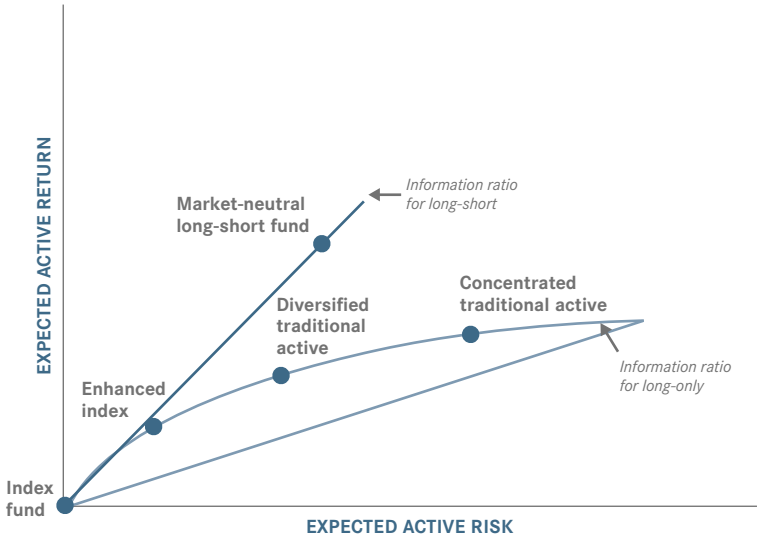

Source: Ronald Kahn, Barclays Global Investors.

σ_α is the expected tracking error of manager p to the properly constructed benchmark. (All return variables are expected, not historical, values.)

If this seems to punish for risk twice (because standard deviation is in the denominator of the equation and also in the denominator of the expression used to construct *IR*), it does. Active risk is to be avoided in proportion to its *square* (active variance).

THE SURPRISINGLY LARGE IMPACT OF THE NO-SHORTING CONSTRAINT

Greater active risk does not just mean a smaller holding, keeping the information ratio constant. As Ronald Kahn has shown, with long-only portfolios the information ratio itself tends to decline at higher levels of active risk.[32] Exhibit 4 illustrates this concept. Even though a skillful alpha forecasting process across securities will generate negative alphas (sell signals) for some securities that are of the same quality as the positive alphas (buy signals) estimated for others, few of the negative alphas can be acted on in a 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 strength of the negative signal.[33]

The amount of alpha per unit of active risk thus declines 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 study to have roughly twice the expected information ratio as efficient as their long-only, traditional active counterparts (given equal skill).[34]

Conclusion: Back to basics

We noted at the outset that practically everyone in the investment business is an active manager of one sort or another. Just about everyone in our field has also gone to business school. In school, one learns from Harry Markowitz that one should build mean-variance efficient portfolios (or, more precisely, that one should build portfolios that maximize return minus a penalty for whatever the investor perceives as risk). One learns from William Sharpe that under stringent conditions the market-cap portfolio is the only portfolio[35] that maximizes investor utility. That may or may not be true, but after Sharpe the burden of proof lay squarely on those who would disprove the inefficiency of cap-weighted benchmarks by beating them.

For the sponsor, Sharpe's observation can be recast as what we call the Two Conditions of Active Management—one should hire active managers if, and only if, he or she believes:

1. That some good managers really do exist, with positive expected alphas (in other words, they are skillful, not lucky); and
2. That the investor in turn has the skill to identify the good managers: that he or she can predict with some accuracy what the managers' alphas will be.

Thus, the lesson of Markowitz and Sharpe is not to avoid hiring an active manager, or being one, but to meet the Two Conditions.[36]

So, should we apply what we learned from Markowitz and Sharpe, and from their intellectual heirs Barr Rosenberg, Richard Grinold and Ronald Kahn, who did for securities what Markowitz and Sharpe did for portfolios? Or should we bow to "practical reality" and use simplifying techniques such as style boxes and historical performance comparisons in the hope that the complexities of the world will go away and we won't have to deal with them (and that no one will notice the difference)?

Well, we should do what is right. We are fiduciaries and are bound to do no less. What is right is to apply the lessons of the masters as best we can.

First, do not pretend that the world is simple when it is complex and there are fine gradations. (Albert Einstein wisely and wittily said, "Everything should be made as simple as possible but no simpler."[37]) Investment style is a risk factor to which anyone can have any degree of exposure they want: it is a *continuum*. No active manager lines up perfectly with growth or value, or large or small styles; he or she has a unique mixture of exposures to these risk factors. Thus, style "buckets" make no sense, since they clump together (for the convenience of the sponsor) managers whose styles typically have only the most superficial resemblance to one another.

Estimate the style exposures of your managers and then make alpha forecasts, which represent the return you expect them to earn after all systematic factors have been accounted for. Then, run an optimizer, as we've described. If you can't or won't do the work necessary to get all the required inputs, at least think about manager structure as an optimization problem. It will help you look at the manager selection question with clear eyes.

Hold a portfolio that you can defend as maximizing expected return and minimizing whatever kind of risk you think it is sensible to be averse to.

Hard-working people have entrusted us, the community of plan sponsors 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 realistically deliver.

Endnotes

1 For brevity, we refer to all securities as "stocks," although our comments could be generalized to apply to any asset class.

2 In addition, for the reverse optimization to work, one needs estimates of the risk (standard deviation) of each stock and of the benchmark, the correlation of each stock with every other and with the benchmark; and the expected alpha for one of the stocks (to get the scale right).

3 Sharpe, William F. "Capital asset prices: A theory of market equilibrium under conditions of risk." *Journal of Finance*, September 1964, pp. 425–442.

4 The expected return on the overall equity market must be higher than that on bonds, cash, and other nonrisky investments in order to get investors to buy equities. 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 it must always be positive for markets to clear. See Grinold, Richard C., and Kenneth F. Kroner, "The equity risk premium," *Investment Insights*, Barclays Global Investors, San Francisco, July 2002; and Leibowitz, Martin L., et al., *Equity Risk Premium Forum*, November 8, 2001, sponsored by the Association for Investment Management and Research, Charlottesville, Va., full text online at *www.aimrpubs.org/ap/issues/v2002n1/pdf/v2002n1.pdf*.

5 Sharpe showed 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 being unrealistic. However, Stephen A. Ross, in "The arbitrage theory of capital asset pricing," *Journal of Economic Theory*, vol. 13, no. 3 (December 1976), pp. 341-360., 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.

6 They are the market along with other market participants, such as index funds and individual investors. It is worth noting that, if one accepts the findings of Brinson et al. (see below) regarding the returns achieved by professional investors (that is, pension plan sponsors and their managers) versus other categories of investors, the professionals beat the other market participants over the time period studied, although not by enough to give them a return higher than that of their passive benchmark. See Brinson, Gary P., L. Randolph Hood, Jr., and Gilbert L. Beebower, "Determinants of portfolio performance," *Financial Analysts Journal*, July/August 1986, p. 39–44; and Brinson, Gary P., Brian D. Singer, and Gilbert L. Beebower, "Determinants of portfolio performance II: An update," *Financial Analysts Journal*, July/August 1986, p. 39–44.

7 This concept is developed further in Sharpe, William F., "The arithmetic of active management," *Financial Analysts Journal*, January/February 1991.

8 We aren't really sure who first classified investments or managers into styles, but William F. Sharpe set forth the best-known formal approach to style analysis (called returns-based style analysis) in "Determining a fund's effective asset mix" (corrected version), *Investment Management Review*, November/December 1988, Volume II, Number 6, pp. 59–69; and in "Asset allocation: Management style and performance measurement," *Journal of Portfolio Management*, Volume 18, Number 2, Winter 1992, pp. 7–19. For determining which styles or

common risk factors are most relevant for use in returns-based style analysis, one authoritative source is Fama, Eugene F., and Kenneth R. French, "Common risk factors in the returns on stocks and bonds," *Journal of Financial Economics*, Vol. 33, No. 1 (February 1993), pp. 3–56, drawing on the much earlier observation by Rolf Banz and Marc Reinganum that small-cap stocks had historically outperformed larger-cap stocks, and the observation by Donald Keim and Sanjoy Basu, 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).

9 Barr Rosenberg and the BARRA organization have led the way in building multifactor (that is, beyond large-small-value-growth) models of securities and portfolios. See *www.barra.com/research/barrapub/risk_models.asp* for a good introductory-level description.

10 The pure alpha may be determined by an actual multivariate regression or by a constrained optimization technique that mimics a regression. The purpose of the optimization technique is to allow for a no-shorting constraint, that is, to require all factor exposures to be non-negative.

11 The beta of the portfolio also goes into this stew, as do additional, non-style factors if one considers them important (as BARRA does). We leave these details out to achieve simplicity of expression. To define the pure alpha a little more precisely than we have done in the body of the paper, it is the return of the manager in excess of the return on a "normal portfolio" that incorporates all of the beta, style, and non-style factors that one puts into the regression analysis, and that may be regarded as the manager's individual benchmark that he would "go home to" if he hadn't generated any particular insights.

12 We would also point out that while non-style factors certainly exist, and help explain managers' active returns, to be practical and fair to the manager one should evaluate him against a custom benchmark or normal portfolio that is built out of available, low-cost index funds of one kind or another (including, say, bond funds or international funds).

Also, the pure alpha may be thought of as a regression alpha with more than one regressor or factor, rather than a simple CAPM alpha or (what is worse) the simple difference between the manager's return and that of the benchmark.

It is of course necessary to convert the pure alpha and pure active risk to sensible units (usually annualized percentage rates of return).

13 And alphas delivered by high information-ratio managers 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.

14 Markowitz, Harry M., "Portfolio selection," *Journal of Finance*, March 1952. Published in book form as Markowitz, Harry M., *Portfolio Selection: Efficient Diversification of Investments*, first edition, 1959; second edition, Basil Blackwell, Oxford, 1991.

15 At least among the more sophisticated, scientifically-oriented managers.

16 See, for example, Kahn, Ronald N., and Andrew Rudd, "Does historical performance predict future performance?" BARRA Newsletter, Spring 1995; also available on the World Wide Web at *www.barra.com/research/BarraPub/hpp-n.asp*. The Kahn and Rudd article 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 Mark Carhart and Russ Wermers and their results do not differ much from those summarized in Kahn (1995). See Carhart, Mark M., "On persistence in mutual fund performance," *Journal of Finance*, March 1997; Wermers, Russ, "Mutual fund performance: An empirical decomposition into stock-picking talent, style, transaction costs, and expenses," *Journal of Finance*, August 2000. Most of these performance studies cover mutual funds because of easy data availability, but the findings are extremely likely to apply in about 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 annualized alpha by the annualized standard deviation of that alpha, and multiply the result by square root of the number of years represented in the data. (We are referring to the pure alpha, or pure active return, and the standard deviation thereof.) Note that this result is just the information ratio times the square root of the number of years of data.

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

19 Kahn, Ronald N., "Most pension plans need more enhanced indexing," in Brian R. Bruce, ed., *Enhanced Indexing: New Strategies and Techniques for Plan Sponsors*, Institutional Investor, Inc. Investment Guides, New York, 2000. Grinold, Richard C., "The sponsor's view of risk," in Frank J. Fabozzi, ed., *Pension Fund Investment Management: A Handbook for Sponsors and Their Advisors*, Probus, Chicago, 1990.

20 Grinold, Richard C., and Ronald N. Kahn. "The surprisingly large impact of the long-only constraint." *Investment Insights*, Barclays Global Investors, San Francisco, volume 3, number 2 (May 2000).

21 No portfolio can be constructed that is above and/or to the left of the frontier; inefficient portfolios are below it and to the right. Note that the portfolio with the highest expected return (at the top right in Exhibit 3) is composed only of the security with the highest expected return. All other portfolios are at least somewhat diversified.

22 The optimization problem is unwieldy without the decomposition because the managers all invest in the same few hundred or few thousand stocks, making the raw returns highly correlated. One has to isolate the pure active returns (by removing the effect of policy and misfit returns) to optimize effectively.

23 The familiar asset-class efficient frontier (which is drawn in total return/total risk space) is upward sloping because capital market theory strongly suggests that riskier assets will provide higher long-term rates of return. There is no corresponding theory suggesting that active return should be rewarded on average; it's a zero-sum game. Thus a pure-active efficient frontier drawn to include managers selected without special skill should not be upward sloping; that is, Exhibit 3 is correctly drawn only for a population of managers that is skillfully selected.

24 Waring, M. Barton, Duane Whitney, John Pirone, and Charles Castille, "Optimizing manager structure and budgeting manager risk," *The Journal of Portfolio Management*, Spring 2000.

25 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.

26 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.

27 Sometimes we hear the assertion that any large collection of active managers in a sponsor portfolio is "just a closet index fund," apparently solely because the number of managers is large. As can be gathered from the balance of this discussion, if the true dimensions of active management have been well-attended to by the sponsor in building this structure, the portfolio of managers would have a positive expected active return without regard to the number of managers. If the portfolio has a positive expected active return, it is clearly not a closet index fund. Of course, any structure of managers that hasn't been skillfully chosen will tend to be a closet index fund—and an expensive one at that—whether it is composed of one manager or many, since the expected active return if unskillfully managed must be zero less fees and costs.

28 To get started in forecasting alphas, sponsors might consider using the forecasting relation;

$$\text{Alpha} - \text{IC} \times \text{Volatility} \times \text{Score}$$

where IC (information coefficient) = expected correlation between one's forecasts and the subsequent realizations of those forecasts (investors' manager selection skill); Volatility = standard deviation of the return being forecast (active risk); and Score = strength of the signal embedded in the forecast, 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) (the forecast of the manager's skill). This relation is further elucidated in Grinold, Richard C., and Ronald N. Kahn, *Active Portfolio Management*, second edition, McGraw-Hill, New York, 1999.

BGI's Client Advisory Group has built a worksheet (descriptively called AlphaBuilder) that implements this calculation, including the additional effects of the typical long-only constraint and fees. The equation shows one way of breaking up the alpha-estimation problem into component parts that may be easier to estimate than alpha itself, and it puts all your managers on a level playing field with respect to their relative alpha.

29 See Michaud, Richard O., "The Markowitz optimization enigma: Is 'optimized' optimal?," *Financial Analysts Journal*, January/February 1989; and Michaud, Richard O., Efficient Asset Management, Harvard Business School Press, Boston, 1998.

30 House of Commons, November 11, 1947, quoted in The Oxford Dictionary of Quotations, Oxford University Press, Oxford and New York, 1979.

31 Our approach leads to some additional intuitive conclusions, particularly on misfit risk and on tactical asset allocation.

Misfit risk can be managed productively. We can put the decision to take (or not take) style misfit risk on a level playing field with policy-risk and pure-active-risk decisions by properly separating these risks. It

just takes a bit of extra work in a formal optimization, and a strong awareness of the separateness of these risks if one is adopting an informal approach, to be confident that the misfit risk taken is the misfit risk one intends to take. While some investors may wish to avoid misfit risk entirely, others will rationally take some misfit risk—that is, have some preferential exposure to large-cap, small-cap, value, or growth styles—in the pursuit of active return that is, in practice, inseparable from the style in which a given manager earns that return.

See Waring, Whitney, Pirone, and Castille (2000) for an extensive discussion of misfit risk and the use of manager structure optimization as a tool for managing it, placing the detailed approach in the public domain. The Client Advisory Group of BGI has implemented this paper in a software program called PortfolioWorks, to demonstrate the ease and usefulness of this approach, and has used it with dozens, perhaps hundreds, of major US and overseas institutional investors.

Tactical asset allocation (TAA) can also fit into the optimization framework, as can any strategy that adds expected utility, as Barton Waring points out in a work in progress. TAA, for example, can be viewed as just another active manager, making bets around a 60–40 or other multi-asset-class benchmark. The TAA manager's expected alpha can then be balanced against the active risk taken by deviating from the benchmark weights.

32 Our equation (2) is adapted from Grinold (1990), page 156, equation (6).

33 Grinold and Kahn (2000).

34 Additionally, there is a second order effect: by not selling short, no extra cash is raised that would otherwise allow further strengthening of the long positions to gather returns from the strong positive alpha forecasts.

35 Where efficiency is measured by the information ratio. If portfolio A is twice as "efficient," in this sense, as portfolio B, it uses the same information to achieve twice as high an information ratio, i.e. twice as high an alpha at the same level of active risk. Portfolio efficiency is thus defined as the ability of the portfolio to use information efficiently to produce risk-adjusted performance.

36 Setting aside combinations of this portfolio and long or short positions in the riskless asset.

37 We have discussed the Two Conditions in pervious work; see Waring, Whitney, Pirone, and Castille (2000), Appendix C; Waring, M. Barton, Lee D. Harbert, and Laurence B. Siegel, "It's 11 pm—do you know where your employees' assets are?," *Investment Insights*, Barclays Global Investors, San Francisco, October 2001.

38 Reader's Digest, October 1977. Peter Neumann, who knew Einstein personally, writes in *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."

AUSTRALIA

BARCLAYS GLOBAL INVESTORS
ABN 33 001 804 566
Level 1, 111 Harrington Street, Sydney, NSW 2000
P.O. Box N43, Grosvenor Place, NSW 1220
Telephone 61 2 9272 2200

CANADA

BARCLAYS GLOBAL INVESTORS
BCE Place, 161 Bay Street, Suite 2500
P.O. Box 614, Toronto, ON M5J 2S1
Telephone 416 643 4000

BARCLAYS GLOBAL INVESTORS
1000 Sherbrooke ouest, Suite 1730
Montréal, QC H3A 3G4
Telephone 514 843 9595

HONG KONG

BARCLAYS GLOBAL INVESTORS
One International Finance Centre, Suite 2901
1 Harbour View Street, Central, Hong Kong
Telephone 852 2295 5288

JAPAN

BARCLAYS NIKKO GLOBAL INVESTORS
Barclays Global Investors Japan Investment Trust
Ebisu Prime Square Tower
1 1 39 Hiroo, Shibuya-ku, Tokyo, 150 8402
Telephone 81 3 5469 4100

Barclays Global Investors Japan Trust & Banking
Telephone 81 3 5469 4300

Barclays Global Japan Investment Trust
Telephone 81 3 5469 4500

SINGAPORE

BARCLAYS GLOBAL INVESTORS
23 Church Street
#13 04 Capital Square, Singapore, 049481
Telephone 65 6395 3320

THE NETHERLANDS

BARCLAYS GLOBAL INVESTORS
World Trade Center
Tower D, 6th Floor
Strawinskylaan 667, 1077 XX Amsterdam
Telephone 31 20 798 0060

UNITED KINGDOM

BARCLAYS GLOBAL INVESTORS
Murray House
1 Royal Mint Court
London, EC3N 4HH
Telephone: 44 207 668 8000

UNITED STATES

BARCLAYS GLOBAL INVESTORS
45 Fremont Street
San Francisco, CA 94105
Telephone 415 597 2000

www.barclaysglobal.com

BARCLAYS GLOBAL INVESTORS