

Five Principles to Hold Onto (Even When Your Boss Says the Opposite)

LAURENCE B. SIEGEL, M. BARTON WARING,
AND MATTHEW H. SCANLAN

LAURENCE B. SIEGEL
is the director of research in
the Investment Division of
The Ford Foundation in
New York, NY.
l.siegel@fordfound.org

M. BARTON WARING
is the chief investment
officer, emeritus, for
policies and strategies at
Barclays Global Investors in
San Francisco, CA.
barton.waring@barclaysglobal.com

**MATTHEW H.
SCANLAN**
is a managing director and
head of the Americas
Institutional Business at
Barclays Global Investors in
San Francisco, CA.
matthew.scanlan@barclaysglobal.com

What are alpha and beta? Why does an investor need to know the difference? How should an investor decide whether to take alpha risk and, if so, how much should he take? How much beta risk should he take? What rewards can an investor expect from taking each kind of risk? Is it always clear how much of a given investment return is alpha and how much is beta, or are there grey areas? Can an investor avoid paying “alpha fees” (i.e., high fees justified only by high levels of skill) for the beta, or index, component of performance, or will there always be some ambiguity about whether the level of fees being paid is fair?

Enthusiastic young business school graduates and freshly chartered CFAs bring with them a leading-edge set of answers and approaches to these questions. Yet in their first days and weeks on their new job, they are likely to be told to “forget all that theoretical nonsense” and, instead, use street wisdom, rules of thumb, and other tired old heuristics because “we’ve always done it this way.” The investment management industry has proved to be remarkably resistant to learning from what we broadly term *modern portfolio theory*. Here we explore ways that this nexus of theories and concepts can be used to improve the practice of investment management, but a very similar article could be written for investment banking and for many other occupations in finance.

The most basic of the better answers arises from the market model of Sharpe [1963], which says that any return can be decomposed into a part (beta) that is correlated with one or more market indices, and a part (alpha) that is not.¹ The correlated, or beta, part results from the portfolio being exposed to markets, while the uncorrelated, or alpha, part is usually understood as the value added or subtracted through active management.

Hidden in this gem are a number of subsidiary principles that are important to anyone aspiring to do well in the investment business. These principles, which follow, have been teased out by two generations of researchers and were summarized in our previous work:²

Principle 1. Investors should think about the returns (and risks) of any investment in terms of their elemental parts: beta, alpha, and don’t forget fees and costs.

Principle 2. Alpha is a zero-sum game, but beta is not (as a consequence, beta bets are tremendously valuable but almost “free,” while alpha bets are not only typically expensive, but are only valuable if the fund manager has exceptional skill).

Principle 3. The criteria for making alpha and beta decisions are completely different, as are the rewards for getting alpha and beta decisions right.

Principle 4. Alpha, delivered consistently by a truly skillful fund manager, is extraordinarily valuable, and the high fee usually charged for alpha may even be a bargain—or it may not, depending on the manager’s skill.

Principle 5. It is desirable, and at least generally possible, to separate alpha from beta in practice, and to pay appropriate fees for each.

These are the principles that we’d encourage our young graduates to hold onto in the face of opposition from bosses who say “we don’t do things that way here.”

Principle 1. Investors should think about returns (and risks) in terms of their elemental parts: Beta, alpha, fees and costs.

This principle, which implies that investors should think about alpha and beta as separate portfolio characteristics and make alpha and beta decisions independently, follows so directly from the market model just discussed that it barely seems worth repeating, except for the fact that it is so seldom taken advantage of by those who would benefit from it. Given that beta is the part of a return that is correlated with markets, and alpha is the uncorrelated part that comes from active management, it makes sense that investors should use one set of criteria to determine market exposures—that is, which “natural” risk premia to try to collect over the long run—and another set of criteria for determining how to beat those markets through “alpha bets”—that is, active management.

Why, then, do some investors seem to have so much difficulty making this differentiation? One source of confusion is that most portfolios are a mix of alpha and beta exposures. In other words, each active portfolio is a blend of two essentially unrelated funds: 1) an index fund or portfolio of index funds delivering the beta (or market) part of the return,³ and 2) various active bets, delivering (or failing to deliver) the alpha part of the return. The active bets are over- or underweights relative to the contents of the index or blend of indices, and may be thought of as composing a market neutral hedge fund.

Thus, any active portfolio can be thought of as an index fund plus a market neutral hedge fund. The index “fund” contains the beta exposures, and the market neutral hedge “fund” contains the alpha exposures, or active bets.⁴

Because these conceptual “funds” are typically blended in practice into traditional (long-only) active funds or directional (that is, incompletely hedged) long/short hedge funds, it requires some thought and effort to make alpha and beta decisions independently. The investor must decide whether to invest in the “index fund” on its own merits, and whether to invest in the active bets on *their* own merits. One of the principal goals of this article is to show how to do this, by explaining how the criteria needed to make an alpha decision differ from those needed to make a beta decision. In this later discussion, which comprises Principle 3, we will also make the crucial distinction between the potential rewards for taking alpha risk and those that accrue from taking beta risk.

In addition, fees—which make up a large fraction of total costs with the remainder being mostly transaction costs—are almost always charged on the whole of the return, that is, on the “alpha beta chowder.”⁵ But the beta part of the return, the part that could have been obtained with an index fund, is only worth an index fund fee, which is very low (at least for major publicly traded asset classes such as equities and bonds). An active management fee should, ideally, only apply to the alpha part of the return, the part that comes from the manager having skill in active management. Keeping fees reasonable is a key ingredient of investment success because fees can be very high relative to the total return earned; and, unlike return and risk, fees can be controlled by the investor. We discuss fees in detail in Principles 4 and 5.

When investors are hiring fund managers (other than index fund and market neutral hedge fund managers), they are always—either consciously or unconsciously—taking on both alpha exposures and beta exposures. And if their combined alpha and beta “bet” is unconscious and unintended, rather than the result of a deliberate decision to take on this particular combination of exposures, they should stop and relearn the principles they learned in business school, or from their study toward the CFA charter, before managing money. These principles form the basic rules of the “game” of active management, and if you are unclear on the rules of a game, you might want to consider not playing it! You can always index.

If the arguments in this section are followed to their extreme, all investment portfolios could be built out of index funds or index derivatives (pure beta) and market neutral hedge funds (pure alpha). Another way to say this

is that all alpha could be sought through portable alpha programs. We are not advocating such a position, however. Active managers who deliver a conventional mix of alpha and beta bets (as most long-only active funds do) have a role—a significant role, if the actively managed funds are priced appropriately—in well-engineered portfolios. But the *mental* separation of alpha and beta, the recognition that these exposures are entirely different in their basic nature, is a mandatory starting point for clear thinking about portfolio construction.

Principle 2. Active management, or alpha, is a zero-sum game—but beta is not.

Almost everyone has heard that active management, the production of alpha, is a zero-sum game. It is, but most people don't believe it. Why not?

The real economy is not a zero-sum game. One likely reason most investors don't believe that active management is a zero-sum game is that the zero-sum argument appears to conflict with the fact that investing overall (alpha plus beta) is not a zero-sum game; investments generally earn positive returns on average over time. They do so because investing is a way to participate in the real economy, and the real economy grows over time. Doesn't capitalism produce an ever expanding pool of opportunities? Yes, it usually does, and not at the expense of someone else. Hasn't the economic flourishing of the last half century, not only in the U.S. and other developed countries, but also in the (rapidly) developing world, made monkeys of those who perceive a zero-sum society, arguing that one person's gain is another's loss? Yes, it has. Don't investors stand to profit from these favorable trends? Yes, they do.

Distinguishing the alpha component of investment return from the beta component. Yet alpha is a zero-sum game, and this is where many investors become confused; if you only think of total returns, and not their component parts, it is pretty easy to miss this point. It is the *beta* part of the return that is not a zero-sum game; the growth of the economy is *in the return of the benchmark; it is beta, not alpha*. Investing in beta (buying index funds or participating in the index-related part of an active portfolio) produces a positive expected return over the riskless rate. Thus, just by being an investor exposed to the overall risk of markets, one participates in the economic dynamism just described.

But active management, trying to *beat* the market index or benchmark (i.e., trying to create alpha) is the part of the investing game that *is* zero-sum. A market, or

asset class, is just an opportunity set of securities. As Sharpe [1991] demonstrated in his classic two-page article "The Arithmetic of Active Management," the return summed across all market participants trading within that opportunity set must, by mathematical identity, equal the return on the opportunity set itself, even though the returns of the individual participants themselves will vary widely.

The amount by which any given investor's return varies from the average return is his alpha, his gain (or loss) relative to the market, from playing the active game. This is the sense in which the alpha part, the *active* part, of investing is a zero-sum game. And it is only zero-sum *before* management fees, trading costs, and other frictional costs; after these costs have been deducted, active management is a *negative*-sum game.

It is very profitable to be one of the winners. This applies both to the investor, who receives a superior after-fee return, and to the fund manager, who receives a performance fee, retains a profitable relationship, or both.

Active management is not like most other professions (Is there an average doctor in the house?). Another reason why many investors can't seem to get a grip on the zero-sum nature of active management is that few other conventional professions are zero-sum. (Athletes playing league sports are an exception; at the end of the season, despite great effort by all, the number of games won exactly equals the number of games lost.)

If you walked out the front door of your building and happened to get hit by a car, an average or even below average doctor would be a welcome sight. An average doctor has a great deal of skill, just not as much as some of his cohorts, and could at the very least keep you alive until you found an above average doctor. An average doctor, or airline pilot, or butcher or baker or candlestick maker has great value, and will have many appreciative customers.

But a fund manager *with only average skill* is almost completely useless! You might learn a good deal about companies and markets over dinner with an average money manager, and such an individual might even have considerable skill, in the ordinary English-language sense of the word. Yet, for the reasons we just reviewed, a fund manager with average skill produces the same before-fee return, on average over time, as that which can be earned in an index fund, even though in some years the manager will do better and in others worse. And after fees, he or she will cost the investor money.

So an *average* money manager, by definition, cannot and does not add any value, before the fact. In fact, value will be lost by the amount of the manager's fees and the trading costs incurred in the effort, which are much higher than they would be in an index fund. A fund manager must have a truly exceptional (that is, a well *above* average) level of skill in order to be expected to beat the market.⁶

After the fact, the average manager will have either won or lost; it is unusual for any active manager to match the market exactly. But this win or loss is pure randomness and can't be counted on ahead of time.

So does this mean that the market is efficient and an investor should just index? The hurdle of skill we've established for being able to beat the market is high enough that it could be surmised that we are saying the market is efficient and indexing is the solution. In fact, we're not saying that at all. While some academic theories propose an efficient market, the description of active management as a zero-sum game is correct—as a matter of mathematical certainty—regardless of whether the market is efficient. If the market is efficient, then there is no such thing as skill in active management and an investor cannot beat the market except by luck. If the market is not completely efficient, then the more skillful active managers will earn positive alphas (on average over time), at the expense of the less skillful active managers, who will earn negative alphas.

Which description of the world is more likely to be accurate? An efficient market depends on all relevant information being incorporated into security prices at all times. The only way this could happen is if a large number of extremely skilled people worked night and day gathering this information and incorporating it (through trading) into the prices of securities. And, with the market being efficient, they would have to do all this work with no hope of earning an extra return! At the same time, the less skilled would have to be aware of their ineptness and drop out of the game, for if they didn't, their influence would cause some security prices to be wrong.

But extremely skilled people aren't fools (almost by definition), so they are not willing to work for free. And we also don't see many of the less skilled dropping out of the game; they continue to influence security prices. (The huge volume of trading is one indication that investors vary in their skill levels, otherwise, they would agree much more than they do on the fair prices of securities and would tend to buy and hold, trading only on new information.)

The logical conclusion is that not all security prices are correct at all times. The market may approach effi-

ciency, but it is by no means *perfectly* efficient. The most skilled of the active managers can be expected to earn alpha, *ex ante*.⁷ Few if any finance professors would disagree with this assertion today.

Selecting winning managers. How does an investor become one of the winners? More specifically, assuming that the investor is in the position of selecting managers, not individual securities, how does the investor pick winners?

Stated succinctly, the investor or plan sponsor has to have above average skill at picking managers with above average skill. At the risk of causing even further confusion, we will say it once more, with feeling: "The investor, as the selector of managers, has to have the same kind of skill at selecting managers that a winning manager has at selecting securities."

Let's explore this proposition in a little more detail. One of us (Waring) has referred several times in the past to the "Two Conditions" that an active manager must meet for his or her effort at active management to make any sense: 1) The manager must believe that the market is, to at least some degree, inefficient, so that successful active management is possible other than by pure luck; and 2) because, even in an inefficient market, active management is zero-sum, the manager must also believe that he or she possesses the skill to be one of the winners by perceiving the inefficiencies better than the rest of the players (in a game where, after fees, more than half of all active managers are losers).

And there is another, completely parallel set of "Two Conditions" that the investor needs to meet before it makes sense to select and hire a fund manager: 3) The investor or plan sponsor must believe that superior managers do exist, which relates back to Conditions 1 and 2; and 4) the investor must believe that he or she possesses the skill to pick the superior managers.⁸

A full discussion of selecting winning managers is outside the scope of this article. A "how-to manual" for making such selections is impossible anyway; any recipe that every investor could follow would necessarily fail. (When judgment is not being exercised, skill is not present.) To assist in getting started, however, Waring and Ramkumar [2008] provided an extensive discussion of the issues involved, setting forth an organized approach to the task of estimating fund manager expected alpha. We realize that the phrase "expected alpha," used to refer to a nonzero quantity, is uncomfortable for some readers, but while the zero-sum nature of active management

implies that the unconditional before-cost expected alpha is zero for any manager, the alpha conditional on meeting both sets of Two Conditions is not zero.⁹

Summarizing Principle 2. Let's see where we've gotten so far. Investors should think about investment returns in terms of their elemental parts: beta, alpha (before fees and costs), and all fees and costs. The market isn't efficient, so the most skillful managers will earn alpha at the expense of the less skillful. Yet the before-cost returns of all managers will sum to the return of the market benchmark. And it is difficult to figure out, in advance, *which* managers have above average skill, so the investor (who is in a position of selecting managers, not securities) must also have above average skill, the skill in this case being whatever is involved in picking superior managers.

It has been said that the test of a readily understood idea is whether it can fit on a bumper sticker. Exhibit 1 is our attempt to summarize the foregoing discussion in a few such pithy words.¹⁰ We'll admit that it is a bit much for a bumper sticker, but it achieves some economy of expression.

Let's now expand on these ideas, seeing how the distinction we have made between alpha and beta bets, and the understanding we have provided on the zero-sum nature of active management, lead to advice that can help investors make good alpha decisions and good beta decisions.

Principle 3. The criteria for making alpha and beta decisions are completely different, as are the rewards for getting alpha and beta decisions right.

One way to clarify the difference between alpha bets and beta bets—a way that may help investors structure their decision making—is to contrast the *rewards* for

taking alpha risk and beta risk, as shown in Exhibit 2, and to contrast the *criteria* for deciding to take each kind of risk, as described in Exhibit 3. The following discussion refers jointly to Exhibits 2 and 3.

Deciding what beta risks to take—the general approach. The reward for taking a given beta risk is an expectation that the relevant market will deliver a risk premium to the investor. This expected reward needs to be weighed against the risks and costs of obtaining it. The risks and costs may be thought of as the “price” the investor must pay for the expected reward (noting that, because of the presence of risk, the reward may fail to materialize).

Classical utility theory offers a framework for balancing these risks and costs against expected return. While a precise description of the theory is beyond the scope of this article, it is taught correctly in business schools and the CFA curriculum, and it forms the basis for Markowitz, or mean-variance, optimization. One key concept of classical utility theory is that the greater an investor's risk aversion, the less attractive any given risk is to that investor.

The risk aversion concept may be uncomfortable for some investors, who have had it drummed into their heads that equity beta risk is just about always desirable—especially for “long-term investors” who can afford to wait “until the expected return is realized.” Investors who believe this notion are on the wrong track, because the expected return may never be realized, and that is why it is called “risk.”

It should never be assumed that any risk is desirable—or undesirable. Because there are many types of beta risk for sale in the market, an investor needs to analyze each type of beta risk on a case-by-case basis, taking into account his or her liabilities, risk tolerance, time horizon, and other factors listed in Exhibit 3. Theory

EXHIBIT 1

Alpha versus Beta: As Simple As Possible, But Not Simpler

	Beta Risk	Alpha Risk
<i>Is risk rewarded with higher expected returns?</i>	Yes; not conditional on skill	Potentially; only conditional on having above average skill
<i>Cost</i>	Very inexpensive relative to comparable active investments ³⁰	Relatively expensive

EXHIBIT 2

Rewards for Taking Beta and Alpha Risks

Reward for Taking Beta Risk	Reward for Taking Alpha Risk
– Equity risk premium (or other equilibrium market risk premium) – All market participants earn same reward (per unit of beta risk taken) – Strongly positive over most long periods of time – Highly variable over shorter periods	– <i>None</i> on average across market participants, but among active fund managers there are always winners and losers – If successful, the investor gets a higher-than-beta return that cushions the portfolio in bad times, supercharges the portfolio in good times – Alpha that is consistent over time adds a cumulative return stream that, in the best of cases, can come close to matching beta return – Alpha creation is NPV creation – Greatest investors have earned huge alphas: Warren Buffett, 8.50% per year; Peter Lynch (and successors at Magellan), 3.58% per year* – But—can you pick the next Buffett or Lynch?

*Alphas calculated over January 1980–April 2000. Magellan's alpha was earned with much less active risk than Buffett's.

then tells us that the investor should add to each kind of beta risk until the expected utility from adding another unit of that type of risk is zero. This calculation sounds daunting until one realizes that an optimizer does exactly that. Thus, because there are multiple beta risks available,

an optimizer is more or less required in order to get the mix of betas that is “right” for a given investor.

Nevertheless, without doing any explicit optimizing, a couple of simplified examples can show how an investor can decide which beta risks to take; let's consider defined

EXHIBIT 3

Criteria for Deciding Whether to Take Beta and Alpha Risks

Criteria for deciding how much beta risk to take	Criteria for deciding whether to take alpha risk
– What are your liabilities? – What risk premium (equity, etc.) do you expect? – What is your risk tolerance? – What is your time horizon?	– Do you believe a given market is sufficiently inefficient that active management can be rewarding? – Do you believe that winning managers exist (other than by luck)? – Do you believe that you are sufficiently skillful as a selector of fund managers that you can pick the winners? – Is the <i>expected</i> alpha large enough to more than cover fees and transaction costs?

benefit pension funds and individual investors who are saving for retirement.

Example: Defined benefit pension funds. First, let's consider a defined benefit pension fund with no cost-of-living adjustment. The liability looks much like a long-term nominal bond, mixed perhaps with some inflation-indexed bonds (TIPS) due to inflation during the period from when the benefit is earned until the employee retires and the payout benefit amount is determined.

For such a pension fund, holding long-term bonds involves very little risk. This is because the fluctuations in long-term bond prices as interest rates change are cancelled out by fluctuations in the present value of the liability caused by the same interest rate changes. In other words, the pension *surplus*, assets minus liabilities, tends to be stable if the fund holds long-term bonds and TIPS that have the same interest rate and inflation exposures as the liability.¹¹

Thus, a defined benefit pension fund should be very comfortable taking on long-term bond beta, as well as some TIPS beta. Given the nature of the liabilities, these betas turn out to be risk reducers, not risk enhancers. They are desirable as a hedge against the liabilities, no matter what the asset-only risk of long bonds and long TIPS happens to be.

For such a fund, however, equity beta is a much tougher sell—the investor has to determine that it is a good investment without any offset from reducing the risk of the liabilities, because pension liabilities do not fluctuate with the equity market to any measurable degree. This observation causes us to question whether the high equity betas observed in many pension plans are a good idea. There is reason to believe that pension fund managers have been persuaded to hold these high equity positions *not* based on their low aversion to risk (these funds should be conservatively managed), but on the false argument that equity risk “goes away” for long-term investors. It is much easier to make the argument that their equity levels should be lower than that they should be higher.¹²

Example: Individual investors. Now, let's consider an individual investor saving for retirement. The “liability” of such an investor is the present value of whatever cash flows are needed to maintain his or her standard of living, presumably in real terms. The beta characteristics of this liability look almost exactly like a ladder portfolio of TIPS. So a wealthy and highly risk-averse older investor who can live on the low yields offered by TIPS might very reasonably decide that TIPS beta is the

only beta exposure he or she needs or wants to take; this is a very low risk position.

But most individual investors are very underfunded (where full funding is defined as the amount one would have to hold in TIPS to provide the *desired* standard of living in retirement).¹³ Younger individuals, cognizant that equity risk might not pay off, may still choose to take that risk, because there is a high likelihood that the risk *will* pay off, and if it doesn't, as Bodie, Merton, and Samuelson [1992] pointed out, they have many years of labor income ahead of them, an additional source of savings and accumulation from which they might replace investment losses. The high equity betas observed in young investors' personal portfolios thus do not hedge any specific beta in the liability, but are consistent with the combined effect of the investor's long investment horizon and his or her potential for future savings from labor income. As the individual investor ages and the potential for future savings from labor income diminishes, approaching zero at the time of retirement, the investment decision becomes strictly one of tolerance for risk.¹⁴

Understanding the equity risk premium. As previously noted, we can't decide for any given investor how much of their portfolio should be exposed to the various risk premia or betas that are for sale in the market. We can, however, briefly review some facts regarding the U.S. equity risk premium, which is the most familiar type of beta.

The U.S. equity risk premium is defined as the expected return premium of the U.S. equity market over default-free long-term Treasury bonds (the latter being a proxy for the riskless asset as perceived by a long-term nominal investor). Ibbotson and Sinquefeld [1982] noted that the realized value of this premium measured over the period 1926–2007 and using the S&P 500 as the relevant index was 5.0% per year with an annualized standard deviation of 19.2%.¹⁵

Given that the realized U.S. equity risk premium is very large,¹⁶ it is also very risky. It was delivered with a great deal of variability, having several very long periods, such as 1929–1950 and 1969–1985, during which the delivered premium was negative (that is, an investor would have been better off in a bond index than in a stock index). More recently, from December 1995 through October 2008, the realized equity risk premium has been negative. Indeed, it requires a very long-term investor, with a stomach of titanium, to shrug off these extended periods of underperformance as short-term fluctuations! And

there was never any *guarantee* that these periods would end, although the likelihood that they would someday end was, of course, quite high.¹⁷

Thus, over *any* time horizon, no matter how long, some risk exists that the expected (positive) equity risk premium might not be realized, or that the realized risk premium will be such a low positive number as to make the risk not worth taking. That is why there is a risk premium—because there is risk! More specifically, holding equities involves taking risks that cannot be eliminated by having a long investment horizon.

Summary: Taking beta risk. Taking a given amount and type of beta risk has the same payoff for everyone; we are all in it together. This is in direct contrast to taking alpha risk, in which one investor's gain is another's loss.

If there is just one kind of beta risk (and some of our other work holds that there is; for example, exposure to portfolio Q, the world market portfolio of risky assets), then the decision of how much beta risk to take is easy. The investor simply determines his or her own risk tolerance and holds an appropriate amount in portfolio Q, with the remainder in the riskless asset.¹⁸ We acknowledge, however, that many, if not most, investors prefer the view that there are various kinds of betas for sale in the market and that different investors want to hold different mixes of them. The beta allocation decision thus becomes a classic optimization problem, requiring estimates not only of the investor's risk aversion, but also of the expected return and risk of each beta exposure and of the correlations among the beta exposures.

Determining the investor's risk tolerance is an art, not a science, but we can offer some pointers—a young investor with substantial expected future labor income (human capital) can probably afford to take more risk than an older investor whose human capital account is depleted. Institutions face a similar picture—a financially strong and growing institution can afford to take more risk than a weak, shrinking, or highly indebted one.

Deciding what alpha risks to take. Deciding what alpha risks to take is so different from deciding what beta risks to take that it almost requires a second brain. Because two heads are better than one, that's not a bad thing.

Seriously, the same person or small group of people is often asked to make both sets of decisions, and one must be able to switch between two very different mind-sets to make these decisions effectively. With beta bets, the investor is analyzing *markets* and ignoring the managers that might be hired to implement the decision. To make an

alpha bet, the investor must analyze *managers* and, for the most part, ignore the market. And, instead of a positive-sum game (we are all in this together), the investor is playing a zero-sum game, that is, trying to choose a manager who will beat the market—and, as explained in Principle 2, that means beating the other managers.

Recalling our earlier discussion of skill, playing the alpha game requires that the investor not just ask whether the manager has skill in the ordinary sense, but whether he has *exceptional*, or well above average, skill. And not just above average among all people, but well above the average of *all* money managers, especially among all the *large* money managers. This is a tough crowd and thus the bar of skill is high. And the manager has to be *above* this high bar, because fees and other costs eat up the performance advantage generated by managers who are merely *a little* above average.

And an alpha bet should not be placed just because a given market is believed to be inefficient. (“Index the efficient markets and hire active managers in the inefficient ones”—no, no, no!) Even in inefficient markets, the returns of the players in that market sum to the market's benchmark. To make an alpha bet, one or more managers with superior skill have to be identified. Never hire a manager just because the manager is active.

Take time to review the two sets of Two Conditions in Principle 2. Not only is it a useful mental exercise, but *all* four of them have to be effectively met to justify placing an alpha bet by hiring an active manager.

Principle 4. Alpha, delivered consistently by a truly skillful fund manager, is extraordinarily valuable and worth a high fee.

So far, we haven't been too encouraging to those who would seek to add value to their portfolios through active management. But that is not really the view we want to leave the reader with. As with many questions, the right answer is not black and white; in fact, there can be a good deal of room for active management.

Great active managers exist. Who are the greatest active managers? How good are they? How much value would have been added by identifying them ahead of time? Siegel, Kroner, and Clifford [2001] asked these questions, studying the 20 years and three months of data that they had available to them at the time. They noted that Warren Buffett's Berkshire Hathaway Corporation, which

is properly regarded as a portfolio rather than a single stock, earned a CAPM alpha of 8.5% a year over this long period. A dollar invested in Berkshire Hathaway for the entire period grew to \$173.28, compared to \$27.43 for its benchmark, the S&P 500. Because Berkshire Hathaway's returns were quite volatile, the highest information ratio was achieved not by this manager, but by the Fidelity Magellan Fund, which had a CAPM alpha of 3.6% a year over the same period, causing a dollar invested in the fund over the entire period to grow to \$72.61.¹⁹

These are not isolated examples. A great many managers have added tremendous value for their clients over long periods of time. Given these results, one could be forgiven for suggesting that hard-core believers in efficient markets, who argue that these managers were all merely lucky and that investment skill is not possible, are out of their minds.

Alpha creation is net-present-value creation. Alpha creation may be regarded as net-present-value creation. If expected returns are boosted by boosting the beta of a portfolio, the discount rate must be increased correspondingly, with the result that no additional present value has been created. A dollar is worth a dollar, whether it is invested in high-beta, low-beta, or zero-beta securities. But a dollar invested in a portfolio with a positive expected alpha is worth more than a dollar. The portfolio's expected future cash flows include an increment from the positive expected alpha, but its discount rate, which should properly reflect only the cash rate plus the premium for any beta risks in the portfolio, does not. Thus, alpha is very much worth seeking, as long as the investor is aware of the risks and costs and is satisfied that the Two Conditions have been met.

Why such high fees? Now, what kind of fee should an investor pay for truly skillful management? We would answer by turning the question around and looking at it from the point of view of the fund manager. Let's say you are a manager who believes you have real skill in beating a benchmark, which is the proper definition of investment skill.²⁰ But the skill is useless without substantial amounts of capital—capital that you don't have. So you go to the capital markets and find out how much you have to pay to get the capital that you need.

You start by trying to borrow money from a bank. But the bank will not lend you any, even at a high interest rate that compensates them for your credit risk. The bankers don't even understand your business plan!

Discouraged, you work to attract equity investors. You find to your dismay that, on average, they expect you to give them 80% of your profits in exchange for their provision of risk capital, allowing you to keep only 20%. You argue that you also need a small flat fee, say just 2% of capital invested, per year, to carry on your operations while waiting for your meager profit share to pan out. You win this part of the argument because the flat fee gives assurance to the capital provider that you can stay in business. A market-clearing price has been established; you accept the deal. You will give away most of the value of your skill in order to have capital upon which you can ply that skill.

Stated this way, 2-and-20 sounds like a rip-off of earnest managers by greedy capital providers (investors). We've exaggerated the dynamics a little, by turning the table upside down, but we did so to make a point. If a manager has real skill, his or her expected superior returns are not a lucky fluke or statistical illusion, this logic explains the high fees charged by, and willingly paid to, some managers. And the turn-about logic would apply just as well to an ordinary *ad valorem* fee (e.g., 60 basis points, or 0.6%) as it does to the 2-and-20 fees of hedge fund and private equity managers.

The answer to the question of "why such high fees," then, comes down to whether the manager has real skill. We cannot forecast with certainty the future track record of any manager, but we know that there have been managers who generated truly extraordinary returns over the time periods for which their performance has been measured, with their investors receiving a better-than-fair return even after high fees have been deducted. Almost tautologically, these fund managers are worth their fees after the fact; the real question is whether they were worth their fees *before the fact*. Was it luck, or was it skill? We do not have particularly effective ways of distinguishing one from the other, but if it was skill, then the manager was worth the high fee before the fact. We'd caution, however, that no manager is so good that he is still good after an *indefinitely* high fee has been deducted!

And, as we've already argued, the *average* manager cannot be expected to add any alpha even before fees, much less after them. Below average managers do not even belong in the discussion. But note "the hiring conundrum." When an investor hires a fund manager, she unavoidably signals her agreement that she thinks the manager will deliver positive alpha. And if the investor and the manager agree that the manager is skillful, then it almost goes without saying that the investor also agrees

that the manager is entitled to a fee that shares the profits to be gained from that skill! Viewed this way, no agreed fee is “high” at all—the level merely reflects some combination of the agreed level of manager skill and the agreed method of dividing the rewards between the provider of skill and the provider of capital.

For investors, the question gets most interesting when considering better-than-average, but not truly extraordinary, managers. At some fee level, such a manager could conceivably offer value to the investor, but because the manager’s fees are subtracted directly from his alpha, the manager needs to beat his benchmarks by more than his fees (and transaction costs incurred in the active management efforts) to be worth hiring.

Principle 5. It is desirable, and at least generally possible, to separate alpha from beta in practice, and to pay appropriate fees for each.

5a. It is possible to separate alpha from beta.

This principle is a little less universally applicable than the others. Conceptually, it is precisely correct that *any* return—the return on any asset, no matter how exotic or hard to evaluate—consists of a part that is correlated to some market or set of common factors (beta) and a part that is uncorrelated (alpha). And when analyzing most publicly traded liquid assets or portfolios that have suitable benchmarks, it is merely a statistical chore to run the market-model regression needed to put a historically driven estimate on the expected alpha and beta—a pretty good solution.²¹ Practically, rather than conceptually, it is often difficult to perfectly separate alpha from beta for many private equity, venture capital, real estate, and traditional (non-risk-controlled) hedge funds, for a number of reasons, including lack of fund transparency and lack of an ideal benchmark.

Investing separately in alpha and beta. In asset classes for which index funds and index derivatives are available, not only can alpha and beta be measured quite accurately, but they can be invested in separately. Taking this thought to its logical extreme, an investor can use *only* index funds and/or index derivatives to achieve the desired beta exposures, and *only* market neutral hedge funds to add (or subtract!) alpha. This latter structure has come to be known as portable alpha although it would probably be better named *portable beta*, since the investor typically invests

close to 100% of available capital in the alpha sources and then uses index derivatives, which require almost no capital, to salt and pepper the portfolio to taste with beta exposures.

We do not think that such an extreme structure is warranted. Some traditional long-only active managers, who blend alpha and beta, have the kind of exceptional skill we discussed in Principle 2, and these managers should be part of an investor’s portfolio. We would say the same about hedge funds that are not market neutral (thus blending alpha and beta) and about 130/30 or other fund structures. Skill can be found in many places and the investor should take advantage of it wherever he finds it. But at the same time, the investor should be cognizant of the aggregate of all the beta positions purchased from various managers or otherwise acquired, and work to ensure that this aggregate is consistent with the intended strategic asset allocation policy.

Alpha and beta for illiquid alternative investments. Most of the assets held by the majority of investors are publicly traded stocks and bonds, in which it is possible to invest separately in alpha and beta, so Principle 5a is mostly true. But many of the most *interesting* investments are in alternative asset classes for which it is difficult, and occasionally even impossible, to accurately distinguish alpha from beta. These include private equity, real estate, some hedge funds, and various other types of deals, such as long-term partnerships organized to invest in energy, other natural resources, infrastructure, agriculture, and timber. And it is more interesting to analyze these exceptions than to keep repeating the mantra that alpha and beta are separable.

Investing in some of these alternatives has the feel of investing as a venture capitalist in a privately held operating company in contrast to buying a portfolio of securities. There is no continually updated market price. There is only the accounting provided by the alternative investment manager, using conventions that are very similar to those that govern operating company accounting. The investor may have to wait 10 or more years to get a full payout of cash and/or marketable securities. At that point, it is possible to look back and calculate the internal rate of return, so the investor can figure out how well he did, but *not* the correlations with asset class benchmarks that would make it possible to calculate reasonably accurate alphas and betas.

A continuum from “very easy” to separate alpha and beta to “very difficult.” We can, therefore, draw a con-

tinuum between assets for which alpha and beta are easily measured and those for which alpha and beta cannot accurately be determined. This continuum is illustrated in Exhibit 4.

A market-neutral, style-neutral, and everything-else-neutral hedge fund consisting of liquid securities is at the “very easy” end of the continuum. The proper benchmark is cash, so the return in excess of cash is all alpha. We can debate whether a given alpha produced in such a manner is due to skill or luck, but it is alpha nonetheless. An index fund is also at the “very easy” end. The whole return is beta, with no alpha.

Traditional portfolios that blend alpha and beta exposures, but that are amenable to regression on a benchmark or mix of benchmarks, are also easy to analyze in terms of alpha and beta. There may be some debate about which benchmark or mix of benchmarks is the right one to use, but the analysis can still be satisfactorily performed (although not perfectly, because there may be a somewhat different average beta exposure in the next period of time than was measured in the study period). This comment applies to long/short (hedged) portfolios as well as long-only portfolios.

Traditional hedge funds that consist primarily of publicly traded securities but that do not follow a fixed strategy (or that have a multi-strategy architecture) and that tend to lack transparency are harder to analyze in terms of alpha and beta exposures. The investor must rely


on a mix of the 1) description of the strategy offered by the manager, 2) regression on various benchmarks, and 3) subjective analysis. Using this blend of tools, it may or may not be possible to achieve alpha and beta estimates in which one has some forward-looking confidence.

The menagerie of private investments in the far right column of Exhibit 4 has betas (and thus alphas) that are conceptually *there*, but little or no data exist with which to estimate them. Empirical estimation, however, is not the only tool we can use. We can also use subjective judgment, asserting that, for example, a given private real estate partnership has a beta that is closer to one—measured relative to a published real estate benchmark—than to zero. (A zero beta would imply that the real estate partnership is benchmarked properly by the risk-free rate, without any expected risk premium for the asset class and without any correlation to the variability of the asset class. This claim can easily be dismissed with the recognition that just about any real estate investment is affected by real estate market fluctuations.)

We are not even completely confident that we have the asset classes in the right side of the diagram in the correct order; many hedge funds may be harder to analyze, in terms of their beta and alpha exposures, than many private investments. But we are sure that the betas—and hence the alphas—are “there” regardless of whether we have the data (or subjective analysis) to put a number on

EXHIBIT 4

Continuum of Difficulty in Separating Alpha and Beta



them. The important thing is to have this intuition and to try to use it even when it isn't easy.

Some things are just too hard to do yourself. It's tempting to attribute the whole return of a strategy to alpha when the alpha and beta components cannot be measured, but it would be wrong to do so. Let's look at this question briefly from the perspective of the Nobel Prize-winning economist Ronald Coase's [1937] theory of the firm. Firms exist because consumers want things done that they could do themselves, but would rather not, because the transaction costs are too high. Stated another way, consumers don't want to acquire the needed skills or perform the needed work. A firm bundles together various desired goods and services into convenient, marketable packages.

For example, although most of us know what a dimmer switch and a windshield wiper are, we probably could not wire the two together to create a variable-speed windshield wiper for our cars. Only a firm with access to large amounts of capital and large pools of skilled labor can do so. Thus, a firm creates value, even if it is just using parts that anyone can buy at the hardware store.

On hot dog stands and carbon credits. Isn't that exactly what some exotic hedge funds and limited partnerships are doing? Let's consider some examples, one silly, one serious.

Some time ago, about 1964, there was a craze for hot dog stands in Istanbul, Turkey.²² At first this foreign influence was opposed, so that teenagers had to sneak out of their parents' sight to get a taste of this new treat; after a while, even the parents were eating hot dogs. This new business made a high return on capital; at least it did for a while. Is this high return alpha or beta? Clearly beta, because one hot dog stand is about as good as another and all the hot dog vendor had to do was "be there." We don't know of any funds that were organized to participate in these lofty returns, but investors should have been willing to pay quite a high fee to get into such a fund, properly viewing the fee as an "assembly charge" for making hot dog stands in Turkey into an investable asset.

More recently, the investment firm Chalkstream Capital, created out of Peter Muller's family office, has helped to finance a strategy of buying carbon credits earned by free tillers in India and selling them to industrial polluters in Belgium. Chalkstream charges a full fee for what obviously could be considered beta, since anyone can enter this business and earn roughly the same return. So why don't we do it and eliminate the "exorbitant"

fee? For one thing, we don't speak the language of free tillers in India, and we are otherwise unfamiliar with this unusual market. As a result, we would have a tough time negotiating a good deal (or any sort of deal) on carbon credits. For another, we don't want to live so far from our family and friends. A hedge fund fee seems like a low price, not a high one, for the convenience of not having to learn this new trade, while still being allowed to keep most of the return from it. This is the same reason we pay the plumber's high prices or the car manufacturer's high prices for combining a dimmer switch and a windshield wiper.

Now, we would argue that most of the return from any of these activities, for which all the investor has to do is "be there," is beta. There is an assembly process that needs to be undertaken to make these cottage industries into investable assets. However, the assembly process, while requiring skill, does not require alpha generation. That is, it does not require the kind of way-above-average skill that it takes to beat other managers, and thus earn alpha, in a zero-sum game. Yet it is worth a high fee.

In other words, there are circumstances in which a beta return is worth an alpha-like fee. The circumstances exist when most people are discouraged, by one thing or another, from performing the required assembly steps. Few people want to move to Turkey to invest in hot dog stands or negotiate with Indian farmers for the carbon credits they've earned.

Alpha and beta are always there, even when they can't be "seen." For any investment, including alternatives, although alphas and betas may not be "seen," they are still there. So there are still separate decisions to be made according to the criteria in Exhibit 3. Do you think all long-term natural resource investments are going to go up, so all you have to do is "be there," with some managers doing better than others but basically all of them earning returns superior to what you expect in other asset classes? If so, that's a beta decision. If not, if you think you need to identify a manager with superior skill to justify making a long-term natural resources investment, then that's a combined beta and alpha decision.

Thus, the alpha-beta construct applies as a concept even in situations where alpha and beta are numerically difficult or impossible to calculate. In earlier work, we argued that Sharpe's observation—that the return on *any* portfolio consists of a market (beta) part and a nonmarket (alpha) part—is "perhaps the most profound insight in

modern finance.”²³ You can run from this insight, but you cannot hide.

Is timing among beta exposures alpha or beta? Some investors seem confused about whether attempts to earn a return by timing among beta exposures (or persistently over- or underweighting a given beta exposure) should be credited to beta or alpha. It’s understandable that they’re confused, because the question contains a surprisingly large number of moving parts.

At first blush, it would seem that such bets should always be credited to alpha. If an investor held the S&P 500 from January 1995 to March 2000, then Treasury bonds until February 2003, then the S&P 500 until October 2007, then Treasury bonds, he would have been a genius (or extremely lucky) at market timing and would have earned a huge alpha.²⁴

But let’s dig a little deeper. Sometimes the answer is not so obvious. Many international (EAFE) managers beat their benchmark in the 1990s by underweighting Japan, which, at its peak, represented almost 60% of EAFE. Whether the return from underweighting Japan is alpha or beta depends on the manager’s *intent*. If the manager underweighted Japan because she thought that Japanese stocks were overpriced, then it’s alpha. If the manager underweighted Japan for risk control reasons (say, to avoid putting 60% of a portfolio in a single country), or because his or her peers were also underweighting Japan, then the return is arguably beta; that is, the return should not be regarded as evidence of manager skill. Unfortunately, it’s very difficult to discern manager intent looking backward over time, because the manager is naturally going to say that he intended to make the bets that had favorable pay-offs. To get the information needed to distinguish alpha from beta in the current example, the manager’s stated intentions in “live action” would have to be known.²⁵ In this case, the positions hurt returns for many years as Japan ascended, then helped returns for many years as Japan descended, but over the whole period of ascent and descent it was about a break-even proposition.

Another example is rebalancing. When a market is in a long uptrend or downtrend, rebalancing asset class weights to the policy mix only after they have drifted appreciably away from the policy weights can add 5% a year compared to the policy mix (which assumes continuous rebalancing). Is this alpha or beta? If it is just a mechanical rule to avoid overly frequent trading, it’s beta. By the way, in a trendless, mean-reverting market, the same rebalancing technique produces a *negative* return relative to the bench-

mark. This, too, is beta. As evidence that it’s beta, note that the two effects cancel each other out; there is no preferred return, over both trending and mean-reverting markets, from the strategy. Stated another way, unless the investor knows, in advance, if the markets are trending, the expected return of the rebalancing strategy is simply zero.

We haven’t answered the general question, when is a beta-timing return really alpha? This question is fodder for an entirely different paper, but in an article on distinguishing alpha from beta, we would be remiss not to bring it up. The main point is that *ex post*, or realized, alpha will result either from intended, skillful “beta bets” or from incidental, skill-free ones just as with any other active position, and it’s similarly hard to distinguish one kind of bet from the other. Our purpose in bringing this issue up is not to suggest that there is a “bright line,” but to raise awareness of the possibility that returns from timing among beta exposures or from holding a beta exposure that is persistently below or above the market weight (which is another form of timing) can produce returns that are fairly credited to either alpha or beta, depending on the circumstances, and that one must look “under the hood” to have any chance at telling the difference. But one thing is sure—if you intentionally make beta-timing bets, it’s about the search for alpha, not for beta.

5b. An investor can and should pay alpha (active management) fees for the alpha part of the return and beta (index fund) fees for the beta part of the return.

So what does Principle 5a imply for fees?

An investor should care about the difference between alpha and beta, because, as noted in Exhibit 3, fees for alpha are very high (thus alpha providers should be selected carefully so the fees are worth it!), while fees for beta from well-established asset classes should be very low. Thus, an investor should try to separate alpha from beta (in a measurement or evaluative sense, not necessarily by investing separately in them), and strive for mental clarity so that she is only paying high fees for positive expected alpha, not for the delivery of beta along with random, unskillful alpha production.

A world in which alpha and beta are *fully* separated and in which investors pay active fees *only* for alpha is a fantasy. By sacrificing many investment opportunities, an investor can achieve the fantasy right now by investing only in index (pure beta) funds and in market-neutral

long/short (pure alpha) funds.²⁶ The investor would then literally be paying active management fees only for true alpha and would be paying the radically lower index fund fee levels for beta exposure. Almost no institutional investor behaves this way, and for good reason. Why not?

First, there's nothing wrong—even in principle—with buying beta packaged with alpha, as long as the investor believes fairly that the manager is skillful and thus that the fees charged are an appropriate blend of the high alpha fees and low beta fees.

Second, to avoid buying beta bundled with alpha, the investor would have to completely avoid asset classes for which the lines drawn between the two types of return are hard to see clearly. These include private equity and debt, real estate, and many of the other alternative investments that we have previously discussed. At any given time, attractive opportunities exist in one or more of these asset classes. Such opportunities should not be overlooked just because benchmarks and index funds for these asset classes do not exist, making it impossible to separate alpha from beta in these investments.

Third, the investor would be ignoring portfolio designs that mix alpha and beta exposures. These portfolio designs include the following:

- Traditional long-only active managers
- Hedge funds or hedge-like funds that take some ordinary beta (stock market or bond market) exposure
- Hedge funds or hedge-like funds that provide exposure to nontraditional forms of beta

There is potential value (potential alpha) in all of these blended vehicles. Any investment medium that attracts individuals on the right side of the “skill bell curve” is a possible source of alpha, and investors should not ignore such sources just because they will inevitably wind up paying some alpha fees for beta exposure.

Therefore, a realistic strategy for implementing a full understanding of the difference between alpha bets and beta bets includes the following steps:

- Separate alpha and beta as cleanly as possible in asset classes that are traded in public markets with transparent benchmarks and liquid index funds, ETFs, futures, or swaps based on those benchmarks. To accomplish this,

- use *more* index funds, enhanced index funds, and market-neutral long-short funds,²⁷ and
- use *less* traditional long-only active management, focusing on managers identified by the investor as having the skill necessary to deliver truly superior results after fees.²⁸

- Take nontraditional beta exposures as desired, recognizing that these exposures involve assembly charges to make the betas available to the investing public. These betas may be worth a fee substantially higher than a traditional index fee. Try your best to negotiate such a fee in these asset classes.
- And, finally, accept that, in some asset classes and types of funds, alpha and beta simply cannot be disentangled, and an alpha fee—beyond what can be justified by assembly costs—is charged on the entire return. Knowing that you will pay an alpha fee for the part of performance that is due to beta, invest in these only if the after-fee return you expect is more than sufficient to compensate for all of the risks taken and fees paid.

Such a structure moves the investor a long way toward the ideal of paying alpha fees only for true, expected alpha and of paying index fund fees for beta. The ideal is not fully achieved, but it gets an investor much closer than does current practice. We hope and expect to see investors adopting this structure in the future. Some of the current trends are in this direction, and we are tremendously encouraged by them, but we also see investors paying very high fees for certain alpha sources. Investors will benefit from controlling fees and other investment costs, while following the principles of alpha-beta separation as best as can be done in a complex world.

CONCLUSION

The folk wisdom developed by investment practitioners tells them to do many things that simply do not maximize investment returns in conjunction with a well-informed effort to control investment risk. It is more worthwhile to pay attention to the lessons learned in business school or in the CFA curriculum, overriding the folk wisdom where appropriate.

The five principles we have discussed here are just a start. We could have expanded our list to 10 or 15 prin-

ciples (and aren't you glad we didn't?), all of which would be derived from the fundamentals of portfolio theory as developed by Harry Markowitz, William Sharpe, Richard Grinold, and their many worthy successors.²⁹

Every development in finance—the proliferation of alternative investments, the appearance of previously undiscovered risks in markets, and so forth—shows the necessity of clear thinking about risk and return. This clarity can best be achieved by fully understanding the fundamental principles of investment finance and applying them as we have begun to do in this article.

Investors who follow these principles will have taken a giant step in the direction of using science to inform their investment decisions. They will have higher returns, less risk, and lower costs. They will be better investors.

ENDNOTES

¹Bowing to current usage, what we are referring to here as “beta” is really the *part of the return* (or risk) explained by reference to any market, submarket, or even set of markets and submarkets. Beta in its original (Sharpe [1963]) sense is, of course, just β , the sensitivity of the asset's return to the *overall* market benchmark.

²Waring et al. [2000]; Waring and Siegel [2003, 2005, 2006]; Kahn, Scanlan, and Siegel [2006]; Waring and Ramkumar [2008]; and Waring and Whitney [2008].

³The “index fund” could be cash, if the overall fund is market neutral; that is, with all of its market-related exposures or betas hedged out.

⁴We have made something of a big deal out of the beta part of an actively managed fund being a *blend* of index funds, not a single index fund. This point is important because, by design or by accident, most actively managed funds contain a mix of “styles” or submarkets (so that the relevant benchmark is a portfolio of style indices) or multiple asset class exposures (for example, stocks plus cash). See Waring and Siegel [2003], and in particular, Waring et al. [2000], which deals in depth with style factors, especially as they apply to controlling misfit risk.

⁵*Alpha Beta Chowder* is, originally, the title of a humorous children's book about language (Steig and Steig [1992]). Balder [2004] of State Street Global Advisors picked up on the phrase as being useful in an investment context.

⁶A money manager with below average skill may be useless to its customers, but is not useless to the market as a whole. By losing money relative to the benchmark, he is a “supplier” of the alpha that above average managers achieve. You do not want these players to recognize their inferiority and drop out of the game. Fortunately, they don't do so very often!

⁷This general line of argument is sometimes called the student's proof of market inefficiency, because every finance

class seems to contain one bright student who intuitively it when presented with the case for market efficiency. It has been formalized by Grossman and Stiglitz [1980]. See also Ibbotson and Brinson [1987], pp. 57–60. The idea that the volume of trading is evidence of trader overconfidence, and thus of market inefficiency, is attributable to Odean [1998, 1999].

⁸Waring et al. [2000]; see also Waring and Ramkumar [2008].

⁹See Waring and Siegel [2003] as well as Waring and Ramkumar [2008].

¹⁰Our regular readers, if there are any, will notice that in Exhibit 1 we have used Albert Einstein's wisecrack that things should be made “as simple as possible but not simpler” for the third time in five years of essay writing. We think this phrase sums up nicely our sense that, as much as we try to simplify investment concepts for the general reader, there is still hard work the reader must do to get the concepts right. Many things about investment finance can be oversimplified to the point where they are flat-out wrong, and we are making every effort to avoid doing that.

¹¹While some pension funds are in deficit, it is conventional to refer to assets-minus-liabilities as the surplus, recognizing that the surplus could be a negative number.

¹²For a much more detailed version of this discussion see Waring [2004a, 2004b] and Waring and Whitney [2008].

¹³This statement presumes that the investor has no other source of pension income and cannot live on the income provided by Social Security.

¹⁴More precisely, the part of the investment decision that is *strictly* one of risk tolerance is the allocation to risky assets versus the allocation to the liability-matching asset. The construction of the liability-matching asset is complex and involves understanding the time dimension of the retiree's liability. See Waring and Whitney [2008].

¹⁵We calculate the premium by subtracting the geometric, or compound annual, bond total return from the geometric stock total return. (There are other ways to calculate this premium, but we find the foregoing to be the most readily understood way.) The standard deviation is calculated using the monthly arithmetic differences of stock and bond total returns, then annualizing by multiplying by $\sqrt{12}$. Total returns consist of capital appreciation plus dividends. The S&P 90 Index was used prior to March 1957, when the S&P 500 Index was initiated.

¹⁶The historical risk premium is not, by itself, a forecast of the future risk premium. Numerous authors (Siegel [1997]; Goetzmann and Jorion [1999]) have noted that the estimate of the equity risk premium based on historical returns may be upward biased; one reason is that countries other than the U.S. have had either lower returns or actual interruptions in the operation of markets (Dimson, Marsh, and Staunton [2002]), and an investor could not have known in advance which country

to invest in. Another source of bias is that if, say, the stock market rises without a change in fundamental values, the historical return increases, but the expected return decreases. Grinold and Kroner [2002], among many authors, use a modified dividend discount model to eliminate this bias, essentially discarding historical return data in favor of future-oriented information. See Siegel [2004] for a more detailed discussion.

¹⁷See Siegel [1997] and Goetzmann and Jorion [1999] for a discussion of the biases caused by only observing markets that survived.

¹⁸See Waring and Whitney [2008]. Portfolio Q was first described in Roll [1977]. Waring and Whitney start with Roll's observation and then argue that there is no single riskless asset, because different investors regard different assets or portfolios as riskless. A portfolio of TIPS, laddered so that the cash flows correspond to the consumption plans of the investor, is a good proxy for the riskless asset (although distortions in the Consumer Price Index, used to adjust TIPS coupons and principal amounts, introduce some degree of risk).

¹⁹The information ratio is the annualized alpha divided by the annualized standard deviation of alpha, and is one commonly used method for quantifying the "goodness" of an active manager. It is a function of skill.

²⁰If the benchmark is cash, as it should probably be for many hedge funds, our point still holds. It is a special case (called market neutral or zero-beta) of the general principle that all active investing is about beating some sort of benchmark.

²¹A suitable benchmark, for this purpose, is one that is calculated at least monthly and preferably daily, using up-to-date prices, with the weights of each security in the benchmark a matter of public knowledge, and with index funds or index derivatives (preferably both) available and liquid so that the benchmark is "investable." The S&P 500, MSCI EAFE, the Lehman Aggregate, and the S&P GSCI (a commodity index) are good examples of such benchmarks. See Siegel [2003]. One caveat is that for funds that reflect a good deal of market timing across the components of beta within them (styles, sectors, and factors), it may be difficult to forecast the beta policy simply by looking backwards at the historic data.

²²Pamuk [2007].

²³Waring and Siegel [2006].

²⁴It is a little tricky to determine the proper benchmark for such a strategy, but our argument does not depend on doing so. This timing strategy would have beaten the S&P 500, any Treasury bond benchmark, or any mix of the two by a very large margin.

²⁵See also Waring and Siegel [2005].

²⁶This statement ignores any flat fee charged by the market neutral fund. For the statement to be literally true, the market neutral funds would have to charge only a money market fund fee for the part of their return that is due to the riskless rate and a performance fee on the remainder of their return.

²⁷Enhanced index funds do not just enter the equation solely because of their typically low fees (but not as low as true index funds). They also have a structural advantage over other long-only portfolios that have higher levels of active risk. Their advantage is that the long-only constraint is less binding, so that more information about overpriced securities is used in the portfolio than is possible in a traditional, less-well diversified fund. Not as much as in a market-neutral long-short fund, but a good deal. See Kahn [2000] and Grinold and Kahn [2000].

²⁸Some readers will recall the same advice in Waring and Siegel [2003], referring to Grinold and Kahn [2000]. However, this time we are emphasizing the fee perspective instead of the information-ratio perspective. It is good advice in both dimensions.

²⁹Our logic for identifying these three contributors to portfolio theory, among the many who could have been named, is as follows: Markowitz laid the entire foundation, showing how risk as well as return is critical to the investment decision. Among Sharpe's many contributions, he was the first to distinguish alpha from beta, the aim of this article. And, unlike many other later researchers who focused on beta, Grinold delved into the source and composition of alpha, which is, of course, the locus of most of the effort that goes into the practice of investment management.

³⁰For the largest and most liquid betas, index funds and futures contracts provide exposure to beta at prices that can fairly be characterized as "almost free." For smaller and less liquid betas (emerging market stocks; lightly traded sector funds) the cost increases, but even so they are always far less than those of active investment in the same asset class or sector.

REFERENCES

Balder, J. "Alpha Beta Chowder: Thinking 'Outside the Box.'" State Street Research & Management, May 7, 2004.

Bodie, Z., R.C. Merton, and W. Samuelson. "Labor Supply Flexibility and Portfolio Choice in a Life-Cycle Model." *Journal of Economic Dynamics and Control*, Vol. 16, Nos. 3/4 (July/October 1992), pp. 427-449.

Coase, R.H. "The Nature of the Firm." *Economica*, Vol. 4, No. 16 (1937), pp. 386-405.

Dimson, E., P. Marsh, and M. Staunton. *Triumph of the Optimists*. Princeton, NJ: Princeton University Press, 2002.

Grinold, R.C., and R.N. Kahn. "The Efficiency Gains of Long-Short Investing." *Financial Analysts Journal*, Vol. 56, No. 6 (November/December 2000), pp. 40-53.

Grinold, R.C., and K.F. Kroner. "The Equity Risk Premium." *Investment Insights*, Barclays Global Investors, July 2002.

Grossman, S.J., and J.E. Stiglitz. "On the Impossibility of Informationally Efficient Prices." *American Economic Review*, 70 (June 1980), pp. 393–408.

Ibbotson, R.G., and G.P. Brinson. *Investment Markets*. New York: McGraw-Hill, 1987.

Ibbotson, R.G., and R.A. Sinquefeld. *Stocks, Bonds, Bills, and Inflation: The Past and the Future*. Charlottesville, VA: CFA Institute (formerly AIMR), 1982. Updated annually by Morningstar, Inc.

Jorion, P., and W.N. Goetzmann. "Global Stock Markets in the Twentieth Century." *Journal of Finance*, 54 (June 1999), pp. 953–980.

Kahn, R.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, 2000.

Kahn, R.N., M.H. Scanlan, and L.B. Siegel. "Five Myths About Fees." *Journal of Portfolio Management*, Spring 2006, pp. 56–64.

Odean, T. "Volume, Volatility, Price, and Profit When All Traders Are Above Average." *Journal of Finance*, 53 (December 1998), pp. 1887–1934.

———. "Do Investors Trade Too Much?" *American Economic Review*, Vol. 89, No. 5 (December 1999), pp. 1279–1298.

Pamuk, O. "Forbidden Fare." *The New Yorker*, July 9 and 16, 2007, p. 48.

Roll, R. "A Critique of the Asset Pricing Theory's Tests." *Journal of Financial Economics*, 4 (March 1977), pp. 129–176.

Sharpe, W.F. "A Simplified Model for Portfolio Analysis." *Management Science*, Vol. 9, No. 2 (January 1963), pp. 277–293.

Sharpe, W.F. "The Arithmetic of Active Management." *Financial Analysts Journal*, Vol. 47, No. 1 (January/February 1991), pp. 7–9.

Siegel, L.B. "Are Stocks Risky? Two Lessons." *Journal of Portfolio Management*, Vol. 23, No. 3 (Spring 1997), pp. 29–34.

———. *Benchmarks and Investment Management*. Charlottesville, VA: CFA Institute (formerly AIMR), 2003.

———. "A Herd of Cats." Unpublished manuscript, 2004, accessed at <http://www.qwafafew.org/?q=filestore/download/107> on August 28, 2008. Forthcoming in Barclay L. Douglas, ed., *Investing for Institutions: Thoughts on the Bottom Line*. Hoboken, NJ: John Wiley & Sons.

Siegel, L.B., K.F. Kroner, and S.W. Clifford. "The Greatest Return Stories Ever Told." *Journal of Investing*, Summer 2001.

Steig, J., and W. Steig. *Alpha Beta Chowder*. New York: Harper-Collins, 1992.

Waring, M.B. "Liability-Relative Investing." *Journal of Portfolio Management*, Vol. 30, No. 4 (Summer 2004a), pp. 8–20.

———. "Liability-Relative Investing II." *Journal of Portfolio Management*, Vol. 31, No. 1 (Fall 2004b), pp. 40–53.

Waring, M.B., and S.R. Ramkumar. "Forecasting Fund Manager Alphas: The Impossible Just Takes Longer." *Financial Analysts Journal*, Vol. 64, No. 2 (March/April 2008), pp. 65–80.

Waring, M.B., and L.B. Siegel. "The Dimensions of Active Management." *Journal of Portfolio Management*, Vol. 26, No. 3 (Spring 2003), pp. 90–104.

———. "Debunking Some Myths of Active Management." *Journal of Investing*, Vol. 14, No. 2 (Summer 2005), pp. 20–28.

———. "The Myth of the Absolute Return Investor." *Financial Analysts Journal*, Vol. 64, No. 2 (March/April 2006), pp. 14–21.

Waring, M.B., and D. Whitney. "A Three-Fund Strategic Asset Allocation Theorem for Investors with Liabilities." Working Paper, 2008.

Waring, M.B., D. Whitney, J. Pirone, and C. Castille. "Optimizing Manager Structure and Budgeting Manager Risk." *Journal of Portfolio Management* (Spring 2000), pp. 90–104.

To order reprints of this article, please contact Dewey Palmieri at dpalmieri@iijournals.com or 212-224-3675.