

Publication


February 2004

Engineering an Alpha Engine

By Dr. Lee R. Thomas, III


Who can resist reminiscing about the long boom of the 1980s and 1990s? We lived in a charmed world where it was actually hard to lose money in stocks or bonds. Each of us has a favorite anecdote, but the averages tell the real story. As stocks roared back from the excessive pessimism of the 1970s, price-earnings ratios soared from 7.4 at the end of 1979 to 32 at the end of 2003, producing 13.65% average annual total returns in the S&P 500. Bond returns were no less extraordinary. A 20-year disinflation, which saw the inflation rate drop from almost 15% in 1979 to 2% by 2003, pulled long-bond yields down from 11.98% to 5.07%¹, and produced for a typical bond portfolio an average annual return of 9.79%.²

**Price-to-Earnings of the S&P 500
1979-2003**



Source: Bloomberg, S&P

**U.S. Inflation
1979-2003**



Source: Bloomberg

¹ Average annual total return based on Lehman Brothers Aggregate Bond Index

² Source: Bloomberg

Publication

February 2004

The chance of another such confluence of events in our lifetimes is not high. Price earnings ratios are unlikely to quadruple again any time soon, since that would drive the market's earnings yield below 1%. And the only way to replicate such sustained disinflation over the next two decades would be to push inflation below zero into deflation—which the U.S. Federal Reserve has vowed it will not tolerate.

So we are condemned to live in the real world, where making money is hard. For many, holding passive positions in stocks or bonds will no longer produce a satisfactory return. What is an investor to do?

One answer is: produce alpha.

During the boom we could all but ignore alpha, that extra few percentage points above the market return that active managers claimed they could deliver. When market indexes were producing double-digit returns, alpha seemed irrelevant. If our active managers failed to deliver—which more than half of them inevitably did—it was unimportant. But now, with market returns to passive stock and bond investments reverting to their norms, outperforming the market has become urgent. Alpha matters.

But producing alpha is not easy. If it were, everyone would do it. In fact, for every investor who outperforms the market, there must be another who turns in sub-par returns. Only half of investors can outperform the average investor, and the average investor defines “the market.” That means roughly half of you reading this article who use active funds managers will suffer a negative alpha. And this is a best-case that neglects to account for transactions costs and management fees. In fact, more than half of you reading this article will underperform indexed management.

Yet producing alpha is not impossible. Financial markets are indeed nearly efficient, but they are not perfectly efficient. The ceaseless quest for alpha by active managers is the invisible hand that pushes the financial markets toward efficiency. If active managers never succeeded, they would abandon the pursuit. Without active managers, the market would be rudderless. Paradoxically, perfect market efficiency would lead to markets becoming inefficient.

But producing alpha often seems a haphazard affair. This paper describes how an investment management firm can produce high quality alpha for its clients, and what that implies for the organization and management of the firm.

We begin by demonstrating how an investment firm can engineer an engine for producing high quality alpha. Like an engineer whose engine designs must obey the laws of physics, we must start from investment theory. Our design is rooted in the assumption that financial markets are very efficient, but not perfectly efficient. When you “bet” against the market, your probability of getting it right can—at best—be only slightly better than 50-50. It can never approach 100%, no matter how hard you try. Our alpha engine is also based on two other relatively recent ideas in investment

Publication

February 2004

theory: the fundamental law of active management and the fundamental law of alpha portability, which we will describe below.

Then we ask what sort of investment management organization is needed to support an alpha engine meeting our specifications.


We conclude that producing alpha requires marginally more skill than the average portfolio manager has. We cannot create alpha from thin air. But it is possible to produce extraordinary outperformance with portfolio managers who are only marginally better than the average. This may sound like magic, but it is not. Like most good engineering, our alpha engine may look like magic to an outsider—hence the many myths about what good money management entails. In fact, our alpha engine relies on careful design and craftsmanship rooted in science, not superhuman powers. Our approach relies on theory, but is designed to be practical and realistic. This is engineering, not pure science. Yet our design requires a much more aggressive approach to portfolio diversification than the norm, and a novel—indeed some might say subversive—management and organization of an investment firm.

Alpha Bets, the Alpha Portfolio, and the Fundamental Law of Alpha Portability

At its simplest, alpha represents the difference between an actively managed portfolio's return and the associated benchmark's return. If we list the securities in the benchmark portfolio we are trying to beat, and compare them to the securities in our actively managed portfolio, each difference represents an "alpha bet." We will call the collection of all these differences, or alpha bets, the investor's "alpha portfolio." If the investor uses more than one active manager, the alpha portfolio consists of alpha bets taken by all the investor's active managers.

If, say, shares in "Company A" represent 3% of the benchmark, but only 2% of your portfolio, then you have an alpha bet: your alpha portfolio is "short" 1% in Company A. This bet is part of the alpha portfolio, as the chart below illustrates. If shares in Company B represent 4% of the benchmark, but 5% of your portfolio, then your alpha portfolio is also "long" 1% in Company B.

Figure 1: The Alpha Portfolio



Generally, active portfolio managers think of the alpha bets in their portfolio as being composed of more than a single, line-item difference between the benchmark and the actual portfolio. Rather, a single bet often is represented by active exposures to many different securities. In our example, the manager may consider the two positions as a single alpha bet that is "long" 1% in Company B and "short" 1% in Company A.

Publication

February 2004

This way of thinking about your portfolio, as the benchmark plus a collection of alpha bets, works for bonds, too. In a bond portfolio, a single representative bet might be to overweight BBB rated bonds, and underweight the same dollar value of BB rated bonds.³ Although the overweight of a number of BBB names and the underweight of a number of BB names would appear as many accounting entries in the alpha portfolio, the portfolio manager thinks of it as a single bet on the future evolution of the yield spread between BBB and BB bonds.

Some of the securities in the actively managed portfolio may not even be in the benchmark. Those securities would be considered overweighted relative to the benchmark weighting of zero.

The general nature of an alpha bet is to overweight one group of securities and simultaneously underweight another. The alpha portfolio is defined to be the difference between the benchmark and the actual portfolio. Figure 1 shows this as merely an accounting identity. But a corollary of portfolio theory makes the concept of the alpha portfolio much more valuable to a practitioner. We call this the fundamental law of alpha portability.⁴

The Fundamental Law of Alpha Portability: *The composition of the optimal alpha portfolio is independent of the benchmark.*

In other words, when you design a portfolio of alpha bets to “beat” a benchmark, you can ignore the contents of the benchmark. When you have finished constructing the alpha portfolio, you simply attach it to the benchmark portfolio. Another way to think about it is this: when we build our portfolio, we can use modular assembly techniques. We can construct a benchmark-replicating portfolio as one module. And we can independently construct an alpha-seeking portfolio as a second module. To assemble the portfolio, just combine the two modules; the aggressiveness of the portfolio, its tracking error, will be determined by how much of the alpha-seeking portfolio you add to the benchmark-replicating portfolio.

This innocuous-sounding idea is actually quite powerful.

First and very practically, it makes it possible to manage many portfolios for different clients with different benchmarks. Suppose you manage funds for 100 different clients, and each client has a different benchmark. What an expensive proposition funds management would be if a manager had to construct 100 different alpha portfolios for 100 benchmarks, requiring costly handcrafting of every portfolio for every client. But because the composition of the optimal alpha portfolio is independent of the

³ More likely, an experienced bond manager would structure the bet as overweight 'x' years of duration in BBB rated bonds, underweight the same number of years duration of BB rated bonds. The idea of an alpha bet is quite general. It can be expressed in terms of dollar values, or parametrically, in terms of duration, beta, or the like.

⁴ See: “A Mean/Variance Analysis of Tracking Error” by R. Roll, Journal of Portfolio Management, Summer 1992

Publication

February 2004

composition of the benchmark, you need only find this single, optimal alpha portfolio, and then attach it to each benchmark. The result is handcrafted quality at mass production prices.

For our purposes, one important implication of alpha portability is that we need not design a whole family of alpha engines. Under ideal conditions, one will do.

Another reason alpha portability is important is that it facilitates diversification. If you are able to add alpha in one particular asset class, then alpha portability allows you to create alpha in a portfolio even if the client's benchmark does not contain this asset class. And if you are able to add alpha in many asset classes, then all these asset classes' alphas can be transported into your client's alpha portfolio, even if the client's benchmark portfolio contains only a few of them.

Once You Have Skill, Diversification Is the Only Thing that Matters

Diversification has been called the only free lunch economics offers. At the price, you should eat heartily.

Many people understand in principle that diversification matters. But as we will see, it's not just that diversification is one of the things you should be thinking about. Unless you are willing to discard the idea of efficiency and assume a super-human ability to beat the financial markets, once you have a little skill, diversification is the only thing that matters.

The One-Bet Approach

Suppose an investment manager makes one alpha bet each year, and this year it is to overweight mortgage-backed securities relative to the aggregate bond benchmark's weight, and simultaneously underweight Treasury bonds. Also, suppose the manager's goal is to add on average 100 basis points of alpha, and further suppose the manager has above-average skill, so that the probability of getting that bet right is 60%. Why 60%? Even a manager with no skill, like a monkey at the dartboard, can expect to get the bet right about 50% of the time. The manager who can't get bets right half the time will have a negative expected information ratio. Simply put, he lacks the skills to remain in this business, and competition will eventually force him elsewhere. So a 60% success rate says our manager has above average skill.

If you believe in perfectly efficient markets, a success rate of 60% sounds high. You would opt, instead, for 50% and market indexing, no matter how skillful the investment house advertised itself to be. Alternately, a success rate of only 60% may sound low to you. A novice portfolio manager might assess the probability of getting her bet right at 80%, or even 90%. But experience disabuses most of us of such faith early on.

The truth is that it is hard to make money in markets that are nearly efficient. The key assumption we are making is that most financial markets are just that: nearly—but not perfectly—efficient. If they were perfectly efficient, all managers would expect to get


Publication

February 2004

their bets right only 50 percent of the time, no one would have the skill to “beat” the market, and designing an alpha engine would be impossible.

But if the market is imperfect, some managers may be able to do a little better than 50%.⁵ However, in nearly efficient markets, the margin between getting it right and getting it wrong will always be narrow. Much is resolved only by chance. Managers who think they can push their success rate substantially higher by more deeply analyzing a bet are fooling themselves. The nature of prices in almost-efficient markets is that most of their future evolution is not predictable, not even in principal. Because market participants are so skillful, chance plays a big role in every financial market outcome.⁶

If an alpha bet is right 60% of the time, then that means it will also be wrong 40% of the time. Assuming that when you bet wrongly you lose whatever you make when you get it right, how big does the single annual bet have to be to produce the targeted average alpha of 100 basis points per year? As the chart below illustrates, the bet has to be sufficiently large so that it makes 500 basis points when it is correct. Why so large? The bet succeeds only 60% of the time, and 60% of 500 equals 300 basis points of alpha. But 40% of the time the bet is going to lose 500 basis points, and 40% of -500 is -200. Add the 300 basis points earned when the bet is correct to the 200 basis points lost when the bet is wrong, and the result is an expected alpha of 100 basis points. In other words, if you repeated this experiment year after year, each year making a single bet and getting it right 60% of the time, this one-bet strategy, scaled to produce (or lose) 500 basis points each time, can expect to average 100 basis points of alpha per year.



Source: PIMCO

This graph is not representative of past or future performance of any PIMCO product.

⁵ Remember, it is a hard job. For each manager who gets it right 60% of the time, there must be another manager who gets it right only 40% of the time; alpha production is, at best, a zero-sum game.

⁶ See: "Proof That Properly Anticipated Prices Fluctuate Randomly" by Paul Samuelson, *Industrial Management Review* Vol. 6: 41-49, 1965. For a more recent entertaining discussion, see "Fooled by Randomness – The Hidden Role of Chance in the Markets and in Life" by Nassim Nicholas Taleb, 2001

Publication

February 2004

How Does the One-Bet Approach Perform?

But what a wild ride it would be.

To measure the success of this one-bet strategy, we divide the alpha it produces by its tracking error, or volatility. This ratio of excess return to tracking error, called the information ratio, can be thought of as the manager's report card.⁷ The information ratio shows how much extra risk the manager had to take to get 1% of extra return. A very successful investment manager probably produces an information ratio of about 0.4, meaning that for every 100 basis points in extra risk, he produces 40 basis points in extra return. A sub-par manager might produce an information ratio closer to 0.2. Conversely, a manager who can consistently produce information ratios of 0.8 or 1.0 would be hailed as one of the great investors of her generation.

So how well did the one-bet, 60%-right strategy perform in terms of both risk and return? The expected alpha is 100 basis points, while the volatility turns out to be 490 basis points, resulting in an information ratio of 0.2. That poor result will make it difficult to retain clients.

How does one improve this information ratio? One obvious answer would be to get the bet right more than 60% of the time. That is the solution most managers concentrate on. In an engineering sense, they try to use brute strength to beat market efficiency. It is the equivalent of a scientist trying to solve the problem of space flight to another star by attempting to design a rocket that travels faster than the speed of light. It won't work, no matter how hard you try, because it violates a fundamental law. No matter how smart you are, the market reflects the combined intelligence of all its players, making it hard to beat much more than half the time. While some investors are more skilled than others, our business is highly competitive. Finding and attracting those people to join your firm is difficult, and it is implausible to suppose that any one firm can succeed just by attracting portfolio managers who are much smarter than average and win their bets more often.

To use another engineering analogy, suppose an architect advertised that his buildings were superior because he had access to better steel than his competitors. Would you find this plausible and hire him? Then why would you hire an investment management firm that made an equally outlandish claim?

We will demonstrate a better approach to improving the information ratio that does not require implausible levels of skill. Instead, it supposes that you, the manager, have access to people of just better than average skill. In a sense, this is really what good management is all about; not finding and recruiting super-humans, but enabling people with just better than ordinary skills to produce extraordinary results.


⁷ See: "The Dimensions of Active Management," by Barton Waring and Laurence Siegel, *The Journal of Portfolio Management*, Spring 2003.

Publication

February 2004

The Five-Bet Approach

This is where diversification comes in. Suppose that instead of making one bet each year, we make five, each independent of the others. As before, we will assume that our objective is to produce 100 basis points of alpha. But with five bets, each must produce just 20 basis points of alpha on average, not 100 as before. The chart below illustrates the implications:



Source: PIMCO

This graph is not representative of past or future performance of any PIMCO product.

By making five bets, each can be scaled so that it is only one-fifth as large as it was in the one-bet strategy, so that when the bet is right, the manager makes 100 basis points from it, and when it is wrong, she loses 100 basis points. It is still possible that this strategy will make 500 basis points in a year—if all five bets are right. Assuming a 60 percent probability of getting a bet right, the probability of getting all five correct is $0.6 \times 0.6 \times 0.6 \times 0.6 \times 0.6$, or 7.8%, as the bar on the far right shows. Whereas the probability of one bet earning 500 basis points in a single year was 60%, the five-bet approach could also lose 500 basis points—if all five bets turn out wrong. That probability is $0.4 \times 0.4 \times 0.4 \times 0.4 \times 0.4$, which, as the far left bar shows, is only about 1%. Recall that the chance of getting the one-bet strategy wrong was 40%.

The much more likely outcome is that some bets will pan out and others will fail. For example, getting four bets right produces 400 basis points and one bet wrong costs 100 basis points, producing a net 300 basis points. As the chart shows, this probability is slightly above one-in-four. Following the same logic, getting three bets right and two wrong results in a net of 100 basis points. The probability of that is about one-in-three.

As we can see, placing just five bets yields a distribution that starts to look like the familiar bell-curve—a normal distribution. What is so striking is that with just five bets, the probability of hitting the targeted 100 basis points exactly is now about one in three. The one-bet strategy never produced exactly 100 basis points in a single year. In some

Publication

February 2004

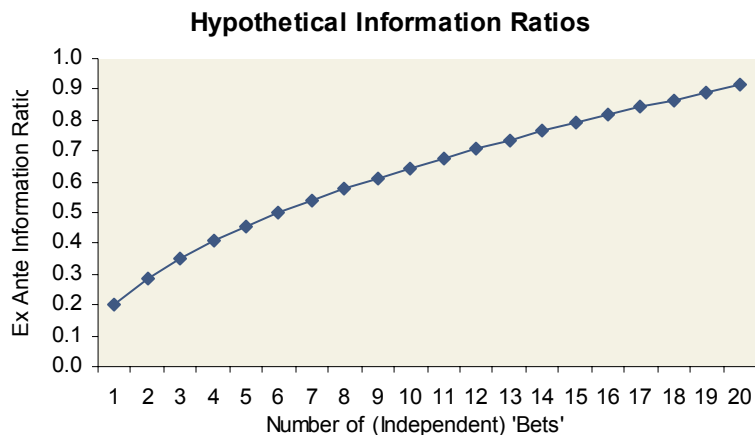
years, the one-bet strategy made 500 basis points and in other years it lost 500 basis points. That is not a pattern clients like to see. The five-bet strategy, by contrast, will produce exactly 100 basis points of alpha in one out of every three years, and in other years returns are much more likely to be within a tight range of the 100 basis point goal (+300 basis points to -100 basis points is the most likely outcome).

How Does the Five-Bet Approach Perform?

Running the same volatility calculation with five bets as we did for one shows how well the five-bet approach performed. Based on our assumptions, volatility in the five-bet portfolio is 220 basis points, resulting in an information ratio of 0.45. Suddenly, this manager has become a player. Rather than an annual result that is up or down 500 basis points, results are much more often in the vicinity of the 100 basis point goal. His average alpha is the same as his one-bet competitor's, but the consistency is far superior. This manager's 100 basis points of alpha is of higher quality.

What Distinguishes a Successful Investment Manager?

We can see that simply increasing the number of bets from one to five a year turns a marginal market participant into a player. This approach obviously can be amplified. The chart below tracks the information ratio as the portfolio moves from one bet (0.2 information ratio), up to five bets, (0.45 information ratio), and so on. Make 20 independent bets each year, each with a 60% probability, and your information ratio grows to 0.9. Suddenly, you are an investment guru.



Source: PIMCO

This graph is not representative of past or future performance of any PIMCO product.

The differences between the outcomes on the left and the right sides of this diagram are striking. An investment manager on the left side of this chart is barely managing to stay in business, while a manager on the right side is a market star, sought for interviews by *Forbes* and CNBC. Yet the difference between failure and stardom merely reflects the difference between taking one investment bet each year and taking

Publication

February 2004

20. In either case the manager's skill—the probability of getting a bet right—was the same, 60%.

Popular media discussions of investment excellence seldom mention “scope” or “diversification.” Instead, success is explicitly or implicitly ascribed to superior skill. Our simple example, where skill is a constant (60% success rate), shows how wrongheaded that idea is. The surprising result is that two managers, equally skilled, produce information ratios ranging from a dismal 0.2 to an extraordinary 0.9. One struggles to survive, the other receives awards. Yet the difference hinges far less on skill, which is hard to raise much above the norm in nearly efficient markets, than it does on increasing the number of independent bets in the portfolio.

That means finding a greater number of high quality, independent alpha bets—or to put it another way, having a broader scope. Yet most investment managers spend most of their time trying to push their success rate from 60% to 70%; that time would be better spent finding more good bets rather than fine-tuning the ones they already have.

So you have two choices. To achieve investment excellence you can find more independent investment bets for your portfolio. That is difficult. Or you can try to become so skilled that you beat nearly efficient markets almost all of the time. That is impossible.

Why Doesn't Everyone Do This?

If taking lots of independent bets is all it takes to produce alpha, why doesn't everyone do it? As it turns out, it's not easy. Finding 20 bets is easy, and even finding 20 bets, each with a success rate of 60%, may be easy. The trick is that each of the 20 bets each year, each with a 60% success rate, must be independent of all the others, which is far more difficult.

A manager might think, “I will make two bets: one will overweight the three-year Treasury note and the other will overweight the five-year Treasury note.” Yet this is not two bets, but rather a single bet—expressed in different ways—that interest rates will fall in the intermediate sector of the U.S. yield curve. Those bets are anything but independent.

Finding these independent, better than average bets each year is easier than becoming a clairvoyant, but it is still hard. And it has profound consequences for how one should organize and manage an investment firm: it must be designed to feed the alpha engine with many independent bets.

Expanding the Concept of Diversification

This requires pushing the notion of diversification much further than modern portfolio theory usually considers. Finding these 20 independent bets every year requires diversifying in many different ways, some much more subtle than those Harry Markowitz may have contemplated.

Publication

February 2004

Diversification Across Instruments

If the key is finding more bets, the first and most obvious aim is to diversify across different instruments. One never knows where an inefficiently priced security will appear – we have no theory of market inefficiency, only theories of market efficiency – but we do know that such inefficiencies are rare. So we need to watch many markets, all the time, so that when a sporadic pocket of inefficiency bubbles to the surface, we can grab it before someone else does. If you manage bonds, then include Treasuries, mortgages, corporates, high yield, emerging and cash instruments – not just one or two of these. If you manage equities, watch growth stocks, value stocks, Big Board and OTC. Preferably, manage both stocks and bonds.

Diversification Across Regions

Diversification across regions provides more opportunities. If we can invest in the U.S. or in Europe, clearly we can identify many more potential bets than if we look only at one. Very simply, investing in two regions offers twice as many potential opportunities to find inefficiency or a mispriced security. But it's much better than that. As soon as you look across regions, you will see not only the bets that you can make within each region, but also potential bets across regions. By doubling the number of regions, it is possible to more than double the number of potential alpha bets.

There is another advantage to cross-regional investing. Most investors place most of their assets in their domestic market; they have what academics call home-country bias. This bias appears to be surprisingly strong and durable. With lots of money chasing inefficiencies within a single market, opportunities are quickly arbitrated away. Significantly less money is invested across regions, making the competition less formidable and creating more opportunities to find high quality alpha bets that span different regions. Home country bias suggests it is quite possible to find that all the alpha bet opportunities within both region A and region B have been arbitrated away, but that region A securities are inefficiently priced compared with region B securities.

Diversification Across Styles

The third area—and where most investment firms fall woefully short—is diversification across styles. The reason is simple: If you are leading an investment management firm, you got there by succeeding as an investor. It is highly likely that you relied on a single style—fundamental analysis, technical analysis, top down macro, bottom up credit analysis. Or maybe you are a quant. You are unlikely to excel at ALL of these styles, unless your name is Clark Kent. Your success at one method probably bred the impression that you know the one true way to make alpha. Not uncommonly, you defend the method that worked for you with religious fervor. Our business is famous for such feuds. Value managers despise growth managers; fundamentalists think technicians are charlatans; quants only respect other quants. But casual observation of the league tables suggests that all these styles can be successful, at particular times and through time.

This point bears restating because it is profoundly counterintuitive and apostasy to most investment firm managers. A chief investment officer is likely to have developed

Publication

February 2004

one particular style, regardless of what it is, to the exclusion of others. He surrounds himself with subordinates who share and reinforce his prejudices. He long ago decided that only one style—his—has any merit. So he emphasizes that style, hires people who share it rather than those who complement it, and seeks a consensus within this homogeneous organization before taking any major alpha bet.

Yet if diversity is the key to investment success—you have to find 20 INDEPENDENT alpha bets each year—this is exactly the opposite of what he should be doing.

If you are a fundamentalist, you should fill your next vacant position with a technician. A growth stock manager should be looking for value stock managers. The complement to a top-down macroeconomist is someone with bottom-up credit analysis capabilities. But the vast majority of successful investors succeeded by exploiting a distinct style, and they are more likely to stick with the style that bred that success. This is a fundamental error.

What we propose is just the opposite: a firm that resembles a collection of boutiques—teams employing different styles, each of which contributes a handful of alpha bets to the overall portfolio. These different team bets are produced independently, based on different approaches to identifying potential alpha, thus generating diversification. The idea of multiple styles in the same portfolio runs counter to the way many investment firms evolve. Most start as boutiques, oriented toward a single style, and they maintain that orientation out of familiarity and prejudice reinforced by success. Diversity also violates conventional thinking among management consultants who see style clarity as a plus and style drift as a sin. But in fact homogeneity and consensus violate a key engineering principal on which our alpha engine is built, the fundamental law of active management.

Diversification Across Skills

Obviously, different styles require different skills, creating a management headache. A top-down macroeconomic style requires an economist; her supervisor will find the job easier if she too is trained as an economist. Bottom-up analysis that appraises financial statements to evaluate companies requires a background in accounting or financial analysis. A quantitative analyst looking for basis, or arbitrage, trades often comes to investment management from mathematics or physics. The manager must manage them all, and understand their disciplines well enough to provide useful guidance, and to assemble the teams' proposed alpha bets into an alpha portfolio.

The organizational challenge of skill diversification is daunting. In a sense, these people don't speak the same language. Bringing them together to draw a conclusion often results in cacophony and acrimony rather than an investment decision. Even if they want to talk to each other, they may not be able to.

Bringing It All Together

So how does one organize an investment management firm to achieve style and skill diversification, and reach conclusions that lead to active positions? It does little good to

Publication

February 2004

gather everybody in a room, each with different training, market specialization and style, to seek a “consensus.” Curiously, this idea, at odds with the fundamental law of active management, is considered enlightened by most investment managers.

The better approach is to organize analysts in small teams—a half dozen or less seems to work best in practice—who are able to talk to each other, have debates, and discuss investment opportunities only within their area of joint expertise. Do not encourage the groups to cross-pollinate—that is a recipe for ending up with one bet in your portfolio, expressed in many different ways. Each group’s objective should be to find the best alpha bets within its specialty.

A word of caution: each independent specialty investment team, looking at its own markets or instruments, must understand that its role is not to be a cheerleader for its sector. A team must be allowed to say, “This quarter, in the markets we’re responsible for, we don’t see any good investment bets to take.” That is a valuable piece of information for the CIO, who can use it to rotate the firm’s active risk budget away from that market sector and allocate it to another team overseeing an alpha-rich environment. That means your compensation scheme (pecuniary and non-pecuniary, including praise) must recognize that recommending no bet at all is a useful piece of information and should be encouraged.

Then the firm must have a process for bringing these alpha bets together, and eliminating bets that do not seem up to standard or which turn out to be duplicative—the same bet, just expressed differently. These bets must then be managed to stay within the firm’s risk budget. After applying risk weights to each team’s proposed alpha bets (optimizing the alpha portfolio’s structure) the combined bets then make up your model alpha portfolio.

Finally, one must ensure that this model portfolio of active bets is expressed in each client’s portfolio, to the extent the client’s guidelines permit. That means transporting the alpha bets from model alpha portfolio into each client’s portfolio. That can require some ingenuity if client guidelines permit some alpha bets to be taken one way, but not another.

Implications for the Investment Management Business

What does our alpha engine imply for investment management? Some cherished conventional wisdom must be turned on its head. The idea of the lone wolf investor, romantic as it sounds, is a model for mediocre performance. The lone wolf who keeps his own council, thinks about the markets deeply and then comes to a profound conclusion on one or two key bets, may actually be much more skilled than the average investor. But the mathematics underlying our alpha engine mean that a team of less-talented investors who can assemble multiple alpha bets into an optimal portfolio will nearly always beat the more-skilled lone wolf.

David Beckham is a great soccer player. But even Beckham could not beat a decent amateur soccer team if he had to play all eleven positions by himself.

Publication

February 2004

About the Author: *Dr. Thomas is a Managing Director and senior member of PIMCO's portfolio management and investment strategy groups. He joined the firm in 1995, previously having been associated with Investcorp, an international investment bank, as a member of the management committee. Prior to that, Dr. Thomas was associated with Goldman Sachs, where he was an Executive Director in the fixed income division of their London office. He has published extensively in the field of international fixed income and foreign exchange management. He has twenty-four years of investment experience and holds a bachelor's degree and a Ph.D. in economics from Tulane University.*

The hypothetical graphs are provided for illustrative purposes only. No representation is being made that any account, product, or strategy will or is likely to achieve profits, losses, or results similar to those shown. Hypothetical or simulated performance results have several inherent limitations. Unlike an actual performance record, simulated results do not represent actual performance and are generally prepared with the benefit of hindsight. There are frequently sharp differences between simulated performance results and the actual results subsequently achieved by any particular account, product, or strategy. In addition, since trades have not actually been executed, simulated results cannot account for the impact of certain market risks such as lack of liquidity. There are numerous other factors related to the markets in general or the implementation of any specific investment strategy, which cannot be fully accounted for in the preparation of simulated results and all of which can adversely affect actual results.

Past performance is no guarantee of future results. This article contains the current opinions of the author and does not represent a recommendation of any particular security, strategy or investment product. Such opinions are subject to change without notice. This article is distributed for educational purposes only and should not be considered as investment advice or an offer of any security for sale. Information contained herein has been obtained from sources believed reliable, but not guaranteed.

Each sector of the bond market entails risk. Municipals may realize gains and may incur a tax liability from time to time. The guarantee on Treasuries and Government Bonds is to the timely repayment of principal and interest, shares of a portfolio are not guaranteed. Mortgage-backed securities and Corporate Bonds may be sensitive to interest rates. When interest rates rise, the value of fixed income securities generally declines and there is no assurance that private guarantors or insurers will meet their obligations. An investment in high-yield securities generally involves greater risk to principal than an investment in higher-rated bonds. Investing in non-U.S. securities may entail risk due to non-U.S. economic and political developments, which may be enhanced when investing in emerging markets. The credit quality of the holdings in a portfolio does not apply to the safety or stability of a portfolio. Diversification does not ensure against loss.

No part of this article may be reproduced in any form, or referred to in any other publication, without express written permission. ©2004, Pacific Investment Management Company LLC.