

A Herd of Cats

Laurence B. Siegel

Updated November 2005 • Links updated September 2017¹

Forthcoming in *Investing for Institutions: Thoughts on the Bottom Line*,
by Barclay L. Douglas

Introduction

Before getting into the science of investing, it's helpful to remember that the goal of that science is to *predict animal behavior*. That's what active investment management seeks to do. Since you can earn the return on any given benchmark without any thought or effort, just by indexing, any effort that you choose to expend must be directed to *beating* the benchmark or passive investment – to earning alpha. That requires “knowing,” in a probabilistic sense, what people are going to do before they do it. It requires making forecasts of human behavior. And people are animals.

Now, the reason for the science: behavioral psychologists can surely predict the behavior of a herd of cats with more accuracy than non-scientists can, or than the cats themselves can. But the forecasts are inherently probabilistic: one would do well to get them a little more right than wrong. And that's how it is with investment management too: an approach to forecasting that is on average a little more right than wrong is about all you can hope for. In fact, such an approach would be a big winner, and would make the person who devises it very wealthy.

Rather than fully developing this theme at the outset, I'll return to it as I visit the various issues that I'm choosing to highlight in this review of investment practice. But it pays to note, early and emphatically, that the goals of the science of investing are modest (and we should not pretend that investing is more scientific than it is). Likewise, expected alphas are modest, because forecasting animal behavior is hard. A popular writer on science, trying to sum up the behavioral research on animals that was taking place in university laboratories, once said something like this: “Animals, when observed under carefully controlled conditions, do as they damned please.” So do people.

¹ Laurence B. Siegel is the Gary P. Brinson director of research at the CFA Institute Research Foundation and an independent consultant. He may be reached at lbsiegel@uchicago.edu.

On benchmarks and benchmarking

Full disclosure

Much of my work in recent years has been a collaboration with Barton Waring, lately of Barclays Global Investors (now BlackRock). He will recognize many of the thoughts in this section – and even some of the words – as his, or more properly as ours jointly. I hope that he and his colleagues will forgive my presumption in publishing what follows under my own name.²

Active management is a zero-sum game

Investment management is a peculiar business: the average practitioner of it *doesn't have any skill*. Due to what the Nobel Prize-winning economist William Sharpe calls the “arithmetic of active management,” active management (if measured against a properly-constructed benchmark) is a *zero-sum game*: the winnings of those who beat their benchmark are supplied by the losses of those who are beaten by it. And that's before costs; after fees and other investor costs are subtracted, active management is a negative-sum game.

It's not exactly correct, of course, to say that active investment managers don't have skill – in fact, almost the opposite is true on an individual basis. If you had dinner with an “average” active portfolio manager, you would find him or her to be a skillful human being, and you would learn something. In a recent book, the investment manager Robert Hagin reconciles these seemingly contradictory claims about manager skill as follows:

[P]rofessional investors are so talented, so numerous, and so dedicated to their work that as a group they make it difficult for any of their brethren to distinguish themselves, particularly in the long run. And the harder investors try to produce above-average investment returns, the more they trade; the more they trade, the more likely they are to end up with below-average investment returns.³

We are unaccustomed to thinking this way about human accomplishment. Not only does an average doctor have considerable skill, even a below-average doctor – a “C” student – would be a welcome sight if I happened to get in a car accident! The active management of investment portfolios is the only field that I know of – except league sports – in which

² See, especially, Waring, M. Barton, and Laurence B. Siegel, “The Dimensions of Active Management,” *Journal of Portfolio Management*, Spring 2003, <https://larrysiegeldotorg.files.wordpress.com/2014/07/barclays-dimensions-jpm.pdf>; and Siegel, Laurence B., and M. Barton Waring, “TIPS, the Dual Duration, and the Pension Plan,” *Financial Analysts Journal*, September/October 2004 (vol. 60, no. 5), <http://www.cfapubs.org/doi/pdf/10.2469/faj.v60.n5.2656>.

³ Hagin, Robert, *Investment Management: Portfolio Diversification, Risk, and Timing – Fact and Fiction*, John Wiley, Hoboken, N.J., 2004.

the considerable abilities of the participants cancel each other out in such a purely mathematical and deterministic way.⁴

This is not to say that some active portfolios will not beat others. Of course they will. The trick is to distinguish skill from luck when they do so, and to measure to what extent that outperformance arises from *beta* (market bets, whether intended or unintended) or *alpha* (the part of return not explained by beta).

To do this we need to spend some effort understanding benchmarks, and we need to define “pure” alpha as distinct from “naïve” alpha. The goal in all of this is to help *forecast* which managers will earn pure alpha in the future, so we can decide which managers to invest with – although this is harder than it sounds – even pure alpha can be generated by luck, so there’s more to forecasting future alpha than determining whether the past alpha earned by a manager was statistically significant.⁵

Get your absolute return right here! The parable of Will Rogers

We’d all like to be absolute-return investors. The first absolute-return investor was Will Rogers, who, when asked the secret of his success, said, “Investing is easy. Just buy some good stock and hold it ’til it goes up, then sell it.” He was then asked, “What do you do if it don’t go up?” “Then I don’t buy it,” he replied.

A couple of generations ago, *long-only* managers told their clients that they were absolute-return investors. They typically did not worry much about the prospects or composition of market “averages” like the Dow Jones – they bought stocks that they thought represented good value and/or the prospect of continued dividend increases. (Thus began, a long time ago, the value and growth styles of investing.)

In 1964, William Sharpe and his Capital Asset Pricing Model (CAPM) put an end to that sophistry.⁶ By showing how to apportion the return on any portfolio between the return from market exposure (beta) and the return from manager skill, or luck masquerading as skill (alpha), Sharpe demonstrated clearly that long-only equity managers were *relative*-return investors. They could add or subtract return beyond that of the (beta-adjusted) market benchmark, but they could not ignore the return of the market benchmark.

⁴ If active managers had no skill at all, however, there would be no one to set the prices of securities, and these prices would therefore be unrelated to their value (the market would be very inefficient). It is because of the skill of active managers that the market is at least somewhat efficient. See Ibbotson, Roger G., and Gary P. Brinson, *Global Investing*, McGraw-Hill, New York, 1993, pp. 38-41; and Grossman, Sanford J., and Joseph P. Stiglitz, “On the Impossibility of Informationally Efficient Markets,” *American Economic Review*, June 1980 (Vol. 70, No. 3).

⁵ There’s no simple recipe for picking winning managers, and there can’t be a recipe since everyone would follow it and it would lose its effectiveness. I recommend trying to figure out if a manager has knowledge that is not already widely disseminated (it is a good sign if the manager can teach me something really new and substantial). Naturally, it also pays to do business with people of good sense and character.

⁶ The CAPM was simultaneously discovered by John Lintner, Jack Treynor, and Jan Mossin, but Sharpe was by far its most visible advocate.

Stated another way, anyone who buys stocks has an opportunity cost of capital. The passive, holdable alternative to buying the active portfolio of stocks is buying the benchmark, through a low-cost index fund. The return on that index fund (after adjustment for beta) is thus the opportunity cost of capital for the active manager. If the active manager doesn't earn more than the opportunity cost of capital, he or she has destroyed value for the customer. Because investment concepts are sometimes confusing (and managers sometimes blow smoke to encourage the confusion), it always pays to look at investment questions from a corporate finance perspective to clarify one's view.

Today, we're faced by a new crop of would-be absolute-return investors: hedge funds. The absolute-return claims of these funds should be evaluated just as critically as the old absolute-return claims of long-only investors. It's harder to evaluate hedge funds' claims because the market factors to which they're exposed are more numerous and exotic than the market factors for long-only investors. I'll discuss these in detail later, in the section on alternative investing. For now, it suffices to say that hedge funds, like long-only managers, are adding alpha only if they produce return above and beyond that which is available from following some sort of simple, formulaic trading rule. In hedge funds as in long-only management, you should try not to pay alpha fees for beta performance!

The role of benchmarks⁷

To identify manager skill – that is, to distinguish between alpha and beta – you need good benchmarks. In ordinary English, a benchmark is a standard of acceptable performance; thus a benchmark in investment management is a paper portfolio, constructed to help determine whether an actual portfolio is being managed well. Any judgment-free, rule-based portfolio might suffice (thus we have equal-weighted benchmarks, for example).

If, however, one accepts the CAPM as a reasonable first cut on the way that assets are priced, then the market capitalization-weighted portfolio of all of the securities in an asset class is a special benchmark indeed. For such a benchmark is *a priori* mean-variance efficient; to beat this benchmark, you not only need skill, you need to have more skill than the average participant in the market. In other words, a cap-weighted benchmark *is* the market; the security positions held by all participants in the market, including active managers and index funds, sum to the positions or weights in the benchmark; and, when compared (after proper adjustment for risk) to such a benchmark, active management is, mathematically, a zero-sum game.⁸

⁷ This topic is developed further in my book, *Benchmarks and Investment Management*, CFA Institute Research Foundation, Charlottesville, Va., 2003. It can be downloaded free of charge at <http://www.cfapubs.org/doi/pdf/10.2470/rf.v2003.n1.3922>.

⁸ It is a zero-sum game before manager fees and other costs are taken into account; after these costs have been deducted from portfolio returns, active management is a negative-sum game.

Cap-weighted benchmarks are thus well adapted to the three main uses to which they've been put: (1) as a basis for performance measurement; (2) as portfolios (index funds); and (3) as the starting point for actively managed funds. This last use arises from the observation that if the benchmark is mean-variance efficient in the absence of special information, then when one collects such special information it should be used to *modify* (raise or lower) the benchmark weights, not to build an unrelated portfolio consisting of 30 or so of one's favorite stocks. Enhanced index funds, risk-controlled quantitative active funds, and market-neutral equity hedge funds are engineered to make use of this observation. Such funds make more effective use of security-level insights than does traditional active portfolio management.

Recently, benchmarks have been roundly criticized. One common critique is that index funds, or active funds managed with the benchmark as the starting point, have too large a weight in the securities that have gone up the most (whether they deserve their high prices or not) and that consequently have large market capitalizations. This critique resonates with investors who bought large-cap stocks at the height of the 1999-2000 "bubble."⁹

But all this critique says is that you could have beaten the cap-weighted benchmark by buying value stocks, or small caps, or anything but technology. It is an assertion by the critic that he or she knew how to add alpha at that point in time. If that's the case – and it's reasonable, at any time, to have a view on how to add alpha – then the active portfolio emerging from this insight should be compared to the cap-weighted benchmark to see if alpha was in fact added. The "bubble" did not negate the usefulness of benchmarks; it just made it relatively easy to add alpha for a while, by overweighting stocks that were out of favor.

At any rate, no one has been able to devise a benchmark-construction rule that gives a portfolio demonstrably more efficient, *a priori*, than the cap-weighted portfolio. Cap-weighted benchmarks are here to stay, as are active managers who would (and should) try to beat them.

And when active managers place their bets, they should do so mindful of the fact that they're going to be compared to the benchmark, not just by being rewarded for earning "pure" active return, but by being *punished* for taking "pure" active *risk*. They should not act as Will Rogers pretended to do, just buying stocks and hoping they go up.

Defining pure alpha

I referred earlier to "pure" alpha and "pure" active risk. By pure alpha I simply mean the regression alpha, which adjusts for the amount of beta risk taken. It is also best to include style factors (value-growth and large-small) in the regression equation. This procedure

⁹ Following the late Robert Bartley, of *The Wall Street Journal*, I have put quotes around the word "bubble" out of respect for the generally effective way that markets have allocated risk capital.

removes the part of the manager's return that could have been achieved by holding low-cost index funds and style index funds. Pure active risk is the standard deviation of the time series of pure alphas. It's not enough to just subtract the manager's return from the return of the benchmark, and to call the result the alpha; that method produces a naïve "alpha" that does not account for the amount of market beta risk or style exposure taken by the manager.

Phooey on style boxes

Investors would also do well to realize that the consultant's emphasis on manager style, while useful, has grown into a pseudoscience in which managers are classified into style "boxes" – large growth, small value, and so forth. In fact, only style index funds (and enhanced style index funds) line up with styles in this way. Active managers, in real life, lie on a continuum between (or beyond) the various styles. The degree of exposure that a manager has to a style is that manager's "style beta," and most managers will have exposures to more than one style.¹⁰ Returns-based style analysis is the most common tool used for identifying these exposures, although holdings-based style analysis may be helpful too. Classifying managers into a few clearly defined styles is no more sensible than classifying people as thin, fat, short, or tall. Most people are somewhere in between.

By regarding style as a continuum or scalar quantity, investors can build portfolios of active managers that take advantage of each active manager's expected alpha while, at the total portfolio level, controlling or avoiding "misfit risk" (the risk of having an unintended style bias in the overall portfolio of managers). To accomplish this, one can run a "manager structure optimizer," described by Waring, Whitney, Pirone and Castille (2000) and elaborated in Waring and Siegel (2003). Optimization across managers is beyond the scope of this article, but it is a useful technique that you may wish to look into further.¹¹

Alpha-beta separation and portable alpha

If you follow this logic where it naturally leads, you can think of an investment fund as consisting of *two* portfolios: a portfolio of betas (asset classes, or market and style exposures), and a portfolio of alpha sources. Traditionally, most alphas have been "bundled" with beta, either in the form of traditional, long-only, actively managed funds (where the alpha-beta blend is obvious) or in hedge funds (where it may not be as obvious but it's there). But today, portable-alpha programs are becoming more popular or at least more widely discussed. In such programs, the investor seeks alpha in any asset class or strategy that seems promising, then independently takes beta bets through futures or swap markets so that the betas in the overall portfolio sum to the desired mix.

¹⁰ Incidentally, this scalar view of style also allows investors to hire managers who do not fit well into a style category, but who instead time between styles, or who have a style that is not well explained by the traditional growth, value, large, and small-cap factors.

¹¹ Waring, M. Barton, Duane Whitney, John Pirone, and Charles Castille, "Optimizing Manager Structure and Budgeting Manager Risk," *Journal of Portfolio Management*, Spring 2000.

Portable-alpha strategies are entirely theoretically correct: there is no reason why the alpha decision and the beta decision should have anything to do with one another. But the investor must be very careful about taking the leverage risk and counterparty risk needed to put portable alpha into practice. Moreover, some advocates have taken the logic of portable alpha to a silly extreme, arguing that you should hold *only* hold the portfolio of alphas, taking no beta risk. Such a strategy is an increasingly popular topic of discussion among institutional investors, but it seems like a terrible idea to me.

Why is it a terrible idea to buy the alphas and not the betas? Betas have a natural return, or expected payoff from the way that risky assets are priced in equilibrium. The equity risk premium is the best-known and most important example of such a natural return; later, I'll discuss ways that the equity risk premium can be estimated. The expected equity risk premium may or may not be lower today than it was on average in the past, but either way, it's still there: when a CEO gets up in the morning and decides to go to work, instead of liquidating the company and investing the proceeds in Treasury bills, *that's* the equity risk premium. He thinks he can beat the Treasury bill return by operating the company.

Alpha sources, however, do not provide any natural return at all, since they depend entirely on having greater skill than one's competitors – on taking away return from someone else. Thus, in the absence of a defined liability – that is, where the object is to make as much money as possible, subject to concerns about excessive risk – I'd rather rely on beta than on alpha. And in liability-driven situations, it's better to find the portfolio of betas that offers the best fit to the betas inherent in the liability (which may be real interest rate exposure, inflation risk exposure, or some sort of beta other than equity beta), and use that portfolio as a starting point, adding incremental beta or alpha sources as appropriate in the quest for higher returns.

Liabilities as a “benchmark”: How should investors determine their policy portfolios?

The first task that confronts any investor is figuring out what the normal or “policy” mix of his or her investments will be. The policy mix should be and usually is expressed as a mix of *betas*, of exposures to general market factors or benchmarks. One can then, if one chooses, try to beat this mix of benchmarks through active management of either the asset mix, the portfolios themselves, or both.¹²

Starting in the late 1970's or early 1980's, and informed by modern portfolio theory and the discovery by Roger Ibbotson and Rex Sinquefeld that equities had produced a large

¹² Thus “tactical asset allocation” is just a kind of active management. Specifically, it's active management of asset class weights (beta exposures) instead of active management of the securities in a given asset class. Any excess return earned thereby should be credited to the manager as alpha, not beta.

and fairly stable return premium over fixed-income assets, traditional practice has been to identify the opportunity set of asset classes, estimate the expected return and risk of each, and estimate the correlation matrix – then optimize. That framework is all well and good, and should be preserved.

But what about the investor’s liability? Shouldn’t the mix of assets have something to do with the reason for which the pool of assets was gathered together in the first place? Shouldn’t the investments be managed to pay the liability? Shouldn’t the risk of an asset, when estimated for use in an optimization problem, reflect the way the asset co-moves with the liability – the extent to which the asset is helpful in paying that liability?

The answer to all these questions, of course, is yes. The liability is the ultimate benchmark. There are many valid reasons to take risk relative to the liability benchmark, say, by holding additional equities. But many investors hold a policy portfolio that is unrelated to the liability and don’t seem to be aware that this mismatch exposes them to risk. It doesn’t seem to have even occurred to them. It’s no surprise that pension funds are in trouble.

Defined-benefit pension plans

Defined-benefit (DB) pension plans have a clearly defined liability but are among the worst offenders in neglecting to match assets and liabilities in any meaningful way. Even though the liabilities for retired and soon-to-retire employees in a DB plan are very much like a bond, and the liabilities for employees expected to retire in the more distant future are also mostly bond-like, most pension plans held 60% or more in equities throughout the “bubble.” They moved, as a herd, into equities as the two-decade-long bull market unfolded, and their equity exposure was at an all-time high when the market peaked in the spring of 2000.

No wonder they have had so much trouble managing the risk to the sponsor’s balance sheet! They’ve been focusing on the wrong risk. The relevant risk is that of a beta mismatch between the assets and the liabilities, not the risk of being different from one’s peers, of straying from the herd. Defined-benefit pension plans should have held more in long-duration bonds and TIPS and should start doing so now if they haven’t already.

In a series of articles published in the *Journal of Portfolio Management*, Barton Waring describes in detail a framework for managing DB plans using asset-liability or “surplus” optimization, with the pension liability incorporated into the optimization problem as an asset held short.¹³ I won’t go into any more detail (you can read his articles), but the approach he prescribes is easily implementable, it’s good for the shareholder, and it’s

¹³ Waring, M. Barton (2004a), “Liability-Relative Investing,” *Journal of Portfolio Management*, Summer 2004; and Waring, M. Barton (2004b), “Liability-Relative Investing II,” *Journal of Portfolio Management*, Fall 2004.

good for the employee.¹⁴ DB pension plans shouldn't necessarily avoid all equity risk; as I noted earlier, equity beta risk has the great advantage that the expected payoff from taking it is large (although volatile) and does not depend on special skill. But if DB plans do take equity risk, it should be with full awareness that they are adding to total plan risk, forcing the shareholders to in effect guarantee the equity investments in the pension plan.¹⁵

In the presence of a liability, then, the investor may regard his or her problem as managing *three* portfolios: (1) the portfolio of betas that best defeases the liability (by matching the beta exposures present in the liability); (2) the portfolio of asset betas that are not present in the liability (you take these risks to earn extra return, beyond the defeasance of the liability); and (3) the portfolio of alphas. Lee Thomas of PIMCO is among those who have emphasized this three-way parsing of investment risk.

Will sound asset-liability management cause the market to crash?

Barton Waring is hardly alone in recommending that pension managers keep at least one eye on the liability. Much current literature addresses this issue, some of it in language stronger than Waring's; Kevin Kneafsey (2003), for example, advocates managing the pension plan to defease the liabilities and to take no additional beta risk.¹⁶ (Kneafsey does suggest trying to add alpha.) The U.K. company Boots moved, for a time, to a 100% fixed income allocation in its pension fund, and the U.S. government-sponsored Pension Benefits Guaranty Corporation is reducing its policy allocation to equities to about 15% from a much larger number. If these trends continue, should we worry about deteriorating stock prices?

In the short run, maybe some concern is appropriate. But, if investors are rational, pension plans will *not* cause the market to decline by selling equities to better align their assets with their liabilities. Here's the reason: as pension plans are currently managed, a share issued by a corporation that sponsors a DB plan is a *double* bet on the stock market: first, on the company, and second, on the investment by the pension plan in the S&P 500. If the sponsor unwinds this double bet, holding stock in the company becomes a pure bet on the company; the leveraged bet on the S&P 500 goes away. Thus the risk (beta) of the

¹⁴ If the pension plan is financially sound, the shareholder benefits by not having to commit as much of the corporate balance sheet to guaranteeing the plan. The employee gains from not having to worry about getting the promised benefits (technically speaking, he or she applies a lower discount rate to the expected benefits). See Bader, Lawrence N., "Pension Deficits: An Unnecessary Evil," *Financial Analysts Journal*, May/June 2004 (vol. 60, no. 3), pp. 15-21.

¹⁵ And, in addition, forcing other plan sponsors (through the mechanism of the Pension Benefit Guaranty Corporation) to guarantee these investments, if the shareholders become insolvent. Finally, the employees (pension beneficiaries) face investment risk taken by the sponsor, since the PBGC does not pay the pension benefits of bankrupt sponsors in full. See Bader (2004) and Waring (2004a, 2004b).

¹⁶ Kneafsey, Kevin, "Solving the Investor's Problem," *Investment Insights*, Barclays Global Investors, San Francisco, August 2003 (vol. 6, issue 5). In what seems to me like an afterthought, Kneafsey does allow for taking equity beta risk, but only as a kind of alpha.

sponsor's stock *decreases*. To keep their risk exposure constant – to stay fully invested – non-pension investors will have to *increase* their equity allocations. Holdings of the S&P 500 by corporate pension plans will simply be replaced by non-pension holdings of the same assets, and the price won't change.¹⁷

Now, I've argued that markets are not all that rational – that human behavior causes price to depart from value, sometimes for long periods. Investors don't always perceive opportunities that are right in front of them (e.g. TIPS in 2000), while falling all over each other to invest in opportunities that don't exist (tech stocks in 2000).

So a decrease in equity allocation by pensions could do some temporary damage to the stock market, or at least put a damper on prices that would otherwise rise more quickly. But a dramatic decrease in the risk of companies with DB plans (and thus of a broad market benchmark, such as the S&P 500) is not hard for analysts and asset allocators to identify, and it decreases the cost of capital or discount rate for the companies' cash flows, so investors with any analytical ability will find them to be a bargain and buy more of them. Every finance MBA learns these principles early in his or her education. It's not exactly rocket science, although – admittedly – sometimes it's a struggle to get practitioners to apply the basic principles they learned in school.

And I would argue that a corporate sector with a financially sound DB pension program is worth more to shareholders than one with a shaky program or one with only defined contribution plans (which will not accumulate enough money to allow most employees ever to retire comfortably). Managing DB plan assets to defease their liabilities should, in the long run, cause the stock market to go *up*.

Endowments, foundations, and individual investors

The principles I've discussed here come into sharpest focus when applied to defined-benefit pension plans, which have clearly defined liabilities, but they apply to any pool of assets. Endowments and foundations, for example, may be regarded as having a liability consisting of their planned spending or payout. (Private foundations are required to pay out 5% of their asset value annually in grants and administrative expenses, while university endowments face no such formal requirement but behave similarly.)

Defined contribution pension plans, and other investments by individuals, may be viewed as having a liability shaped by the individual's income needs in retirement. For most of these classes of investors, the liability looks much like an inflation-linked bond, so that U.S. Treasury-issued TIPS (which, unfortunately, have yields too low to satisfy most investors' return requirements) may be combined with equities and other risky assets to help defease their "liabilities."

¹⁷ This idea has been circulating in the financial community and is not original to me.

Individual investors, in fact, probably stand to reap the largest benefit from an asset-liability view of the problem they are facing – if only because this group has been particularly poorly served by the financial system. The assets that are relevant to a full asset-liability study for the individual include:

- Human capital (discounted present value of future labor income)
- Real estate and/or businesses
- Discounted present value of expected DB pension income
- Discounted present value of expected Social Security income
- Tax-deferred investment accounts, including defined-contribution pension accounts
- Taxable investment accounts
- Cash and other miscellaneous assets

Liabilities include living expenses at an acceptable standard for the remainder of one's life; taxes; and mortgages and other debts. Anything remaining is "owner's equity," although one might classify a reserve for emergencies as a liability rather than equity.

Typically the only assets over which the individual has control, in the sense of deciding the asset-class mix (and the active strategy, if any, for beating the benchmark represented by that mix) are the tax-deferred and taxable investment accounts. Based on this analysis one can arrive at an asset mix for these accounts that best diversifies the risks inherent in the other assets and that enhances the return of the overall "portfolio."

In addition, by purchasing single-payment fixed annuities,¹⁸ the individual investor can hedge mortality risk (the risk of outliving one's money), as defined-benefit pension plans do by their very structure.

Asset allocation, security selection, and behavior

Much ink has been spilled on the question of whether asset allocation or security selection is more fruitful.¹⁹ I tend to side with those who say asset allocation, since

¹⁸ Or certain other types of annuities. Annuities with the income payment indexed to inflation, to the extent that any such annuities are issued at competitive prices, should be of particular interest to individual investors.

¹⁹ The idea that asset allocation is responsible for 90% or more of investment performance is usually attributed (mistakenly) to Brinson, Gary P., Brian D. Singer, and Gilbert L. Beebower, "Determinants of Portfolio Performance II: An Update," *Financial Analysts Journal*, May/June 1991, which updates an earlier study by Brinson, L. Randolph Hood, Jr., and Beebower. Actually the Brinson studies only indicated that some 90% of the variation in pension plan returns (across time – not between funds) was attributable to the policy mix of asset classes. The 90%, in other words, is the average, across funds, of the R^2 's from the regression of each fund's return on its policy benchmark. This result simply shows that market factors explained most of the return, as one might expect when most plans have similar policy mixes and don't take much active risk.

security selection (active management) is a zero-sum game; the expenditure of effort in active management is thus subject to an especially harsh version of diminishing returns.²⁰ But whatever the exact proportions are, *asset allocation is easy and security selection is hard*, so most of the effort in the investment business goes into security selection, which is *inherently* behavioral. It's all about figuring out how people are going to react to information. (While some analysts would rather portray their craft as that of figuring out the difference between the price and the value of an asset, that calculation is all for naught if people do not act to reduce or eliminate the difference and provide to the analyst the profit he or she requires.)

So there are two parts to the investment process: policy and active risk, beta and alpha, engineering and behavioral science. Matching up betas on the asset side with betas on the liability side – and taking the additional step, if one chooses to do so, of figuring out what additional beta risks you want to take on the asset side in the hope of earning additional return – is engineering. All of the remaining effort that you care to expend consists of trying to earn alpha; that's behavioral science. Practitioners could be forgiven for calling the latter activity an art.

What is the expected return on the stock market? (How much can we know about the future?)

It would be very helpful to know what the average return on the stock market is going to be in the long run. (Other asset classes are important too, but the stock market is such a large fraction of wealth, and thus of the set of investment opportunities, that I devote this section to it.)

Many prominent authors, including the late Chris Hensel, Don Ezra, William Jahnke, and Ronald Surz, have weighed in on this controversy; for a full discussion of these issues, see Ibbotson, Roger G., and Paul D. Kaplan, “Does Asset Allocation Policy Explain 40, 90, or 100 Percent of Performance?” *Financial Analysts Journal*, January/February 2000, http://www.mangustarisk.com/doc/pdf/Does_Asset_Allocation_Explains_40_90_100_Performance.pdf.

The key to understanding these issues, as Ibbotson and Kaplan point out, is to specify very carefully the question being asked; about 40% of the *return variation between pension funds* is attributable to differences in asset allocation, with the remainder (60%) attributable to differences in active return (that is, to timing and security selection). Since active management is a zero-sum game, almost exactly 100% of the *level of returns* is explained by the asset-class mix. And the commonly cited 90% figure is as explained above. These two questions are very different and, naturally, have very different answers.

²⁰ Mark Kritzman and Sébastien Page have written an excellent article (“The Hierarchy of Investment Choice,” *Journal of Portfolio Management*, Summer 2003) in which security selection, not asset allocation, is shown to have greater impact. But you must read the article carefully. They say at the outset that it is “normative, not positive” – in other words, that it says how you should behave, not how people actually do behave. Moreover, they answer the slightly odd “normative” question: Conditional on having equal skill in all investment activities, should you apply that skill to asset allocation or security selection (presuming you can do only one and not both)? In other words, they're asking where the greater *opportunity* is, not where the better *results* have actually been achieved. Given Kritzman and Page's careful specification of the question, they are correct in saying that security selection is the more fertile field in which to apply skill.

Why is the long-run expected return on equities so important? It helps you plan ahead and decide how much to save and invest. It also helps you decide how much risk to take, since you should take more risk if the reward for risk-taking is high than if it is meager. Speaking a little more technically, the first decision in managing any pool of assets is setting the policy benchmark – in other words, determining which point on the efficient frontier to hold as a starting point (before adding active management or alpha decisions). This whole process is referred to as asset allocation. The asset allocation process depends on estimates of the return and risk of each asset class, and well as the correlation between each asset class and every other.

Let's set aside the question of risk and correlation, supposing that for these parameters some kind of analysis of the past gives reasonably good estimates for to the future. We'll focus here on expected returns.

Roger Ibbotson (my former boss) and Rex Sinquefield, drawing on the work of Lawrence Fisher and James Lorie, among others, showed that, over 1926 to the present, equities have earned a compound annual return some 7% higher than the inflation rate, or 5% higher than fixed income assets. (The difference between these two numbers is the real interest rate, which thus averaged 2% over the period measured.)

When graphed on a log scale over the nearly eight decades of the Ibbotson-Sinquefield study, the cumulative index line for stocks looks almost like a straight line (see Exhibit 1). (I've extended the time period in Exhibit 1 back to 1871, to show that the Ibbotson-Sinquefield observation is, in round terms, confirmed by earlier results. Also, having anchored the starting date to 1926 for so large a portion of my career – I was an employee of Ibbotson Associates from 1979 to 1994 – it's fun to do something different.) The illusion of no risk for long-term equity investors is so powerful that only a Scrooge would point out that the sharp jag downward in 1929-1932 represents an 89% decline in the stock market!

Ibbotson and Sinquefield's measurement of the historical returns on these asset classes is, of course, accurate. (While the returns they measured are before fees and other costs, and are for indices that were not "holdable" through index funds over much of the time period, these are not valid critiques of their methods. That's how you measure the return on any asset class.) Where Ibbotson and Sinquefield, early in their body of work, took a leap into the darkness was to assume that the return-generating process (for equities in excess of bonds, that is, for the equity risk premium) is stable over time. Their method also asserts that because the process itself is invisible, the historical return is the only helpful estimate that we have of it.

This set of assumptions leads to the famous Ibbotson and Sinquefield forecasts, which are probabilistic – that is, they forecast not just the expected return but the range of likely deviations around it, as shown in Exhibit 2. In other words, their approach recognizes that the forecasts are subject to uncertainty; they may not turn out to be accurate. These forecasts formed the conceptual basis for just about all of the asset allocation work in the

investment management industry from the late 1970's until the dividend discount model (DDM) counter-revolution of the last few years.

Are stocks risky in the long run?

Let's examine a little more closely my earlier observation that Exhibit 1 gives a powerful illusion of no risk in the stock market.

Now that we've seen where the beginning and end points in Exhibit 1 are, we can draw a straight line through them (or a best-fit regression line through the full data set) and "see" that whenever there is a deviation from the straight path, the market eventually snaps back to it and crosses it. Thus there's no risk to the truly long-term investor: returns are self-evidently mean-reverting, so that if you wait long enough, you'll earn the long-run average return!

Well, maybe not. (A little logic shows that returns must be in some sense mean-reverting – if extraordinarily good returns cause stocks to become overpriced, they are more likely to be followed by poor ones, and vice versa. But wait a minute.) While there's risk in the deviations around the line, as 1929-1932 and more recent episodes demonstrate, the biggest risk comes from the fact that we didn't know in advance what the slope of the line would be. In other words, you don't know what the mean you're reverting to is. And you never will.

Ibbotson and Sinquefeld themselves never said that stocks were riskless in the long run. To the contrary, their method emphasized the risk of stocks by drawing wide confidence bands around the forecast means. That stocks are riskless (or have low risk) if you wait long enough is just a misunderstanding of Ibbotson and Sinquefeld. Those who adhere to that misunderstanding sometimes use Ibbotson and Sinquefeld's data to support their cause, but they shouldn't.

Survival bias

A number of investigators, including Roger Ibbotson himself in his collaboration with Gary Brinson, *Global Investing*, pointed out that historically-based forecasts of long-run rates of return may be biased because one is observing only markets that were lucky enough to have survived. This principle is best illustrated relative to a hypothetical international portfolio of equity index funds, purchased at the beginning of the last century, when there were no "developed markets" and Europe, North and South America, and other parts of the world were bursting with "emerging markets" (see Exhibit 3).

As it turned out, an investor who held funds in the United States, the United Kingdom, and a few other small countries would have enjoyed uninterrupted equity markets up to the present, but investors who bought Germany, Japan, Russia, Austria-Hungary, China, and so forth would not have. All of these countries now have stock markets, but at some

point the investor would have lost everything and, to remain investors, would have had to inject new capital earned in the labor market.²¹

This observation implies that the return achieved after-the-fact by investors in the U.S. or U.K. equity index is an overestimate of the return they expected before-the-fact. In other words, the U.S. or U.K. result is one of the better outcomes that an investor in 1900 might have hoped for. Most investors fared much worse. That such a wide range of outcomes was not only possible but likely must have been known by investors in 1900, a date when the U.S. Civil War and various upheavals in Europe were within living memory. Thus, the expected return on equities is almost surely lower than the historical number. Fama and French (2000) appear to have confirmed this conjecture by using the dividend discount model to estimate the returns that U.S. investors expected or required at each point in historical time; they come up with an equity risk premium of about 3%.²²

Procyclicality

Another critique of the Ibbotson-Sinquefeld method is that it provides forecasts that are procyclical. As the stock market soared in the late 1990's, causing the historical average equity risk premium to increase, rate-of-return *forecasts* also became larger (because the forecasts embodied the future-equals-past assumptions). This defies the ordinary logic of investing, in which expected returns decline as the price of an asset rises. Thus forecasts that were reasonable (and which were vindicated by later results) in 1974 seemed extravagantly optimistic at the price levels that prevailed in 1999.

This critique is correct, and should warn investors to think for themselves (and to be cognizant of the assumptions underlying a forecasting method) rather than accepting a single forecasting method for all time. The Ibbotson-Sinquefeld method not only assumes that the future will resemble the past but that the market is fairly priced. If the market is substantially over- or underpriced, you have to use a different method, one that includes the current price as an input. The DDM fits this criterion.

To his credit, Ibbotson's consulting firm, Ibbotson Associates, now uses multiple methods including the future-equals-past method, a version of the DDM, and a method (that of Ibbotson and Chen) that combines aspects of both.²³

²¹ I develop this theme further in "Are Stocks Risky? Two Lessons," *Journal of Portfolio Management*, Spring 1997. William Goetzmann, with various co-authors, has produced a large body of literature on survival bias (in country markets, mutual funds, hedge funds, etc.); see his remarkable and fun Web page at <http://viking.som.yale.edu>.

²² In addition to the well-known published literature on the equity risk premium, I recommend Grinold, Richard C., and Kenneth F. Kroner, "The Equity Risk Premium," *Investment Insights*, Barclays Global Investors, San Francisco, July 2002 (Volume 5, issue 3); and Leibowitz, Martin L., et al., *AIMR Equity Risk Premium Forum*, November 8, 2001, accessed in September 2017 at http://www-stat.wharton.upenn.edu/~steele/Courses/956/EquityPremium/EquityRiskPremiumForum_TIAA_AIMR_2001.pdf

²³ Ibbotson, Roger G., and Peng Chen, "Long-Run Stock Returns: Participating in the Real Economy," *Financial Analysts Journal*, January/February 2003 (vol. 59, no. 1), pp. 88-98.

The DDM counter-revolution

The dividend discount model, of course, had long existed at the time Ibbotson and Sinquefeld did their seminal work and also gave forecasts of the return on equities. The DDM states that the return-generating process is *visible*. Reduced to its simplest form, it says that investing is like a savings account. That portion of earnings that is not spent (on capital improvements, paid out as dividends, and so forth) is added to the stock's fundamental value. As the fundamental value of a stock rises, so does its market value. Roger Ibbotson, in a personal communication, even once told me that he thought that the DDM gave the theoretically best forecasts, since it is forward looking and, if the dividend growth rate estimate is correctly formed, avoids the survival and procyclicality biases, and other measurement errors, of the future-equals-past method.

The DDM, however, wasn't widely used in forecasting of long-run rates of return on market benchmarks because forecasts of the dividend growth rate were (and are) notoriously inaccurate. It seemed as difficult to forecast the dividend growth rate – needed for a DDM estimate of the expected return on the stock market – as to forecast the stock market total return itself; so why bother?

In 1984, however, Jeffrey Diermeier (then of First Chicago Investment Advisors) pointed out, in the course of writing an article he co-authored with Roger Ibbotson and me, that corporate profits cannot reasonably be expected to grow indefinitely as a share of gross domestic product (GDP).²⁴ Thus the long-run growth rate of real (inflation-adjusted) GDP – which, in the United States, has been about 3% – is also likely to be a good estimate of the real growth rate of aggregate corporate earnings and, by extension, of dividends. One can use this information to rough in a forecast of the stock market:

$$E(R_S) = \text{dividend yield} + \text{real GDP growth} + \text{inflation} = 2\% + 3\% + 3.5\% = 8.5\%$$

The numerical estimate represents conditions in late 2005. We'd add an extra one-half point to reflect the likelihood that the dividend yield, long depressed by tax policy, will rise over time. This is the expected return on all corporations, including those privately held, and is not calculated on a per-share basis.

The estimate then needs to be refined to take account of share buybacks and new issues (that is, to convert the estimate to a per-share basis); and to recognize any difference in expected growth rates between corporations in general and those in a particular capitalization-weighted benchmark (say, the S&P 500). Any expected change in valuation (the price/earnings or price/dividend ratio) also needs to be included; we've assumed no such change. While these adjustments can be complex, 9% seems like a good current estimate of the expected equity return. This number is only slightly lower than that produced by the future-equals-past method.

²⁴ Diermeier, Jeffrey J., Roger G. Ibbotson, and Laurence B. Siegel, "The Supply of Capital Market Returns," *Financial Analysts Journal*, March/April 1984 (vol. 40, no. 2), <https://larrysiegeldotorg.files.wordpress.com/2017/09/supply-faj84.pdf>.

The use of a DDM with a link to GDP to estimate equity returns did not become widely accepted after my 1984 article with Diermeier and Ibbotson. One reason is that the forecasts it gave were too low; the higher future-equals-past forecasts were vindicated by subsequent results. The DDM forecasts were too low partly because the method, as conceived at that time, included no allowance for changes in valuation (the P/E ratio or dividend yield of the market); but valuations increased sharply over the 1984-2000 period, as bond yields and inflation rates fell. Today, however, some analysts expect a *decrease* in valuation and include that forecast in the model; we do not do so.

Today, DDM-based equity forecasts with a link to GDP are commonplace, and among investment professionals they have largely supplanted the future-equals-past method. (Some financial planners and other participants in the retail investment market still place strong emphasis on future-equals-past.) Robert Arnott, Clifford Asness, Peter Bernstein, Richard Grinold, and Kenneth Kroner, among others, have refined the DDM (and argued for its logical superiority) to the point where one can refer unironically to a “DDM counterrevolution.”²⁵ The state of the art of estimating expected equity returns has come full circle.

Alternative investing: Understanding its risks and costs

Alternative investing has been defined elsewhere as investing in anything other than publicly traded stocks, bonds, or cash. Thus private equities (including venture capital and buyout funds), real estate, commodities, timber and farm land, and oil and gas deals are all alternative investments.

While this definition is a start, the rise of hedge funds as a widely used investment vehicle gives us pause, since the underlying holdings of many hedge funds are long and short positions in stocks, bonds, and cash. Thus a more useful definition might be “anything that is not publicly traded stocks, bonds, or cash, or that involves short selling.” The use of leverage and derivatives is also characteristic of hedge funds but is not required for an investment to be considered a hedge fund. Alternative investments might also be defined as those in which the general partner receives 20% or so of the profits attributable to the limited partner’s capital, but that’s too cynical. I think.

Because of the relatively recent and sudden popularity of hedge funds, this section focuses on these funds. The other, older types of alternatives are still important but hedge funds need and deserve analytical attention.

²⁵ Grinold and Kroner (2002) was already cited. See also Arnott, Robert D., and Peter L. Bernstein, “What Risk Premium Is ‘Normal’?” *Financial Analysts Journal*, March/April 2002 (vol. 58, no. 2), pp. 64-85; and Asness, Clifford S., “Bubble Logic: Or, How to Learn to Stop Worrying and Love the Bull,” AQR Capital Management working paper, New York, August 2000. Accessed in September 2004 at <http://ssrn.com/abstract=240371> (a download fee applies for the full article).

Risk characteristics of hedge funds: Ordinary beta, exotic beta, true alpha

We have spent much effort separating alpha and beta, and exhorting investors to do so at every opportunity. Let's apply this analysis to hedge funds.

Hedge funds, like other funds, are exposed to a mix of market (beta) risks and security-specific (alpha) risks. (A few quantitative, market-neutral hedge funds do hedge out all or almost all market risks, but those are very much the exception.) Many hedge funds have exposure to ordinary equity beta, but what makes market risk particularly difficult to identify and isolate in hedge funds is the presence of "exotic beta."

What do I mean by exotic beta? Basically anything other than exposure to the S&P 500 or one of its international equivalents. Examples of exotic, but still market-related or beta, risks taken by hedge funds include:

- Interest rate risks (duration, yield curve risk, credit risk)
- Value minus growth stocks
- Credit spreads
- Optionality (for example, being "short volatility," as if selling put options)
- Buy acquired, short acquirer (merger arbitrage)
- Buy convertible bonds, short stock of issuer (an example of capital structure arbitrage)

In what sense are these beta risks? Cliff Asness has pointed out that if you can write down a recipe – a security selection or trading rule – then the resulting risk exposure should be characterized as beta. This is a bit of an extreme view. I'm more comfortable with the idea that it's beta only if you can replicate the risk position with low-cost index funds, futures, options, or exchange-traded funds (ETF's).

But Asness' view makes the worthwhile point that beta is more common, and thus alpha scarcer, that it appears on the surface. While investors are accustomed to thinking of beta as exposure to the stock market, as it usually is in traditional, long-only investing, they should get used to thinking of beta more broadly as they move into hedge funds and other structures where beta exposure is harder to detect. And investors should try to avoid paying "alpha fees" – fees that can be justified only if true alpha is being created – for "beta performance" (that is explained by either ordinary or exotic beta).

It's difficult for investors to perform this analysis. Returns-based style analysis of hedge funds, with the exotic betas as the "styles," would be the preferred method if the requisite data existed and were of reliable quality. To some extent, investors will have to look at the hedge funds' holdings (to the extent they can get this information), and interview the managers, to estimate the exposures to exotic beta.

Cliff Asness has also noted that alpha can *become* beta over time. The first person to notice that value stocks had experienced higher long-term returns than growth stocks should be insulted to hear the value-minus-growth spread called beta. It was alpha, and

whoever discovered this deserved to earn alpha-related fees. But now that everybody knows about the value effect, and now that there are value and growth style index funds, value-growth has become beta. An interesting thought.

Investors should also be on the lookout for exotic beta in *traditional* investments. Style and size exposures are the most common of these betas. Investors should try to avoid paying alpha fees for performance that can be attributed to style exposure and that could have been captured through a style index fund. As I suggested in the section on style boxes, returns-based style analysis can be used to thus separate alpha and beta.

To summarize, you should analyze a hedge fund as you would any active manager. It's all about separating alpha from multiple betas, and controlling costs by paying, when possible, only for the true or pure alpha. You can get beta almost for free, so you should.

Are we paying alternative investment firms too much?

Of course we are.

That doesn't mean that the delivery of consistent alpha should be cheap. Anyone who can forecast human behavior well enough to add alpha (relative to a properly constituted benchmark) is worth plenty.

But the absolute dollar size of hedge fund fees has become ridiculous. The two-and-20 structure (2% of assets plus 20% of profits, per year) that was necessary to get a "crazed genius" (Cliff Asness' wry phrase) to share his investment wizardry with a few wealthy friends simply isn't appropriate in an institutional setting where billions of dollars are being allocated. This structure has made it so attractive to start a hedge fund that there are now more than 8000 of them – most of which will go out of business.

The high fee structure of hedge funds (and private equity) does attract some smart and highly motivated people. But it also attracts some "dumb" ones, some of whom will succeed, due to luck (or due to the selling of beta as apparent alpha). In the long run the smart fund managers will drive the others out of business, but the winnowing process is slow, because it takes a long time to determine whether the alpha is statistically meaningful. Meanwhile, an awfully large number of slow-witted, lucky centimillionaires will have been created.

So far, hedge fund firms that have tried to compete on fees have gotten nowhere. Hedge fund fees will be eroded by competitive forces, but because of this opacity, the erosion will take place more slowly than one might expect.

At any rate, hedge fund managers can't be worth ten times the fees collected by traditional managers. They're the same guys.

Conclusion

Like the other authors in this book, I've been asked by Barclay Douglas to identify beliefs I hold with certainty, those I hold with near certainty, those held with better than 50% conviction, and those widely held by others that I think are completely wrong. I'm going to keep my answers short, because I've been long-winded on many topics, some of which bear on this question.

Beliefs held with certainty

I'm afraid I can't be much help on beliefs held with certainty. Bertrand Russell said, "Never be absolutely certain of anything." That degree of skepticism may not be appropriate for one's most cherished moral or spiritual beliefs, but I think it's about right for beliefs held about investing money or about any other enterprise that involves forecasting animal behavior.

Beliefs held with near certainty

Beliefs I hold with near certainty are also in short supply, but I have one: markets "work," so price will bear some relation to value. If the price and value of an asset or class of assets are widely separated, they may become even more widely separated before you run out of patience or capital, but eventually they will move closer together. They really will, every time. So fundamental analysis is not useless, the market is not completely efficient but it tends in that direction, and there is alpha for the taking by those who are more astute than the average of the manager population.

Beliefs held with better than 50% conviction

- Most individual investors, and many institutions, would do best to buy and hold a portfolio of low-cost index funds, diversified across asset classes. You can't control returns (much less alpha) but you can control costs, which are a one-for-one decrement to returns. If you think you have skill at identifying active managers who will beat their benchmarks on a properly risk- and cost-adjusted basis, by all means do so, but know that your confidence level is probably too high.
- There are worse things than relying on committees and building consensus. Most of the "cowboys" who would be heroes are actually riding wrong-end-to on their horse. Committees can stop them. A few real heroes will be unrecognized this way, but much damage will have been prevented. In other words, two heads are better than one (this sequence does not expand indefinitely – an investment committee of more than six is probably too big).

Commonly held beliefs that are totally wrong

- Hedge funds and other alternative investments are the only investments worth having...NOT. While hedge funds may have been winners when traditional long-only managers were losing money in the 2000-2002 bear market, they have the same job as every other active manager – beating some sort of passive benchmark (the benchmark could be cash if the manager is truly hedged so that he or she has a zero beta). This is a zero sum game and half of those who try will fail – before fees! Thus both “pure” alpha sources and traditional long-only active managers are as likely to deliver a negative alpha – again before fees – as a positive one. Meanwhile, *beta* bets – bets on the equity risk premium or on some other market-related risk premium – will be rewarded in the long run; those are *not* a zero-sum game.
- Good investment managers eat their own cooking? Baloney! Maybe they should if they are a global multi-asset-class manager or outsourced CIO, but there aren't many of these. I believe in prudence and diversification, and so should an investment manager when acting on his or her own behalf as an individual investor. When we institutional investors hire these managers to act as our agents, we diversify widely, investing a maximum of, say, 3% of our total fund's wealth in a given alpha source, and sometimes much less. It's hypocritical – if not downright irrational – for us to expect a manager to invest most or all of his or her own assets in the (almost always specialized) fund they're managing.

As a concluding thought, I'd like to draw on the wisdom of the world's greatest investment philosopher: Buffett.

Not Warren. Jimmy. (The two Buffetts are distant cousins.)

There's a community of Jimmy Buffett fans on the Internet called the Church of Buffett, Orthodox, who believe that Buffett's commercial work amounts to “apostasy” and that Buffett's “spiritual core” resides in his earliest, most personal songs. “If I ever had to defend myself to the Church of Buffett,” he said in the August 17, 1998 edition of *Time* Magazine, “I would only say that the bitterest artists I know are those who had the chance to jump through the hoop and chose not to take it. They stayed on as coffeehouse singers. But I jumped through, not knowing what was on the other side.”

What was on the other side may not be great art, but it's more than good enough, and it has delighted a great many people. Following one's dream can be a satisfying choice, but it involves the risk of withdrawing into a self-constructed world instead of engaging the realm of humanity. The bitterness that comes when the dream doesn't work out is a high price to pay for personal purity. For goodness' sake, do something that matters to other people, and that improves their lives.