TO: Financial Commission

FROM: Russell J. Morreale, Staff Liaison

SUBJECT: An Informational Report of Active Investment Management

RECOMMENDATION:

Accept and discuss a presentation by Chair Donald Korn

BACKGROUND

In recent year the City embarked on a new investment model as a basis for increasing portfolio diversification while maintaining liquidity, capital preservation and prudently enhancing yield as, in the order listed, the primary portfolio investment objectives.

DISCUSSION

At recent Commission meetings the discussion of active investment management has arisen as a point of discussion and interest. Tonight's presentation by Chair Korn serves to introduce this topic as a basis for information sharing and general discussion.

To start discussions, Chair Korn has generously provided several pertinent articles of interest and background regarding this area of conversation. They are as listed below:

Attachments:

A. The Arithmetic of Active Management
B. Investment Management Fees Are (Much) Higher Than You Think
C. How consistently do active managers win? (available in hard copy only)
D. The Paradox of Skill
E. What Practitioners need to know
F. Financial Analysts Journal Commentaries

Attachment A

The Arithmetic of Active Management

From the Board

The Arithmetic of Active Management

by William F. Sharpe, Timken Professor Emeritus of Finance, Stanford University, and Chairman, William F. Sharpe Associates

"Today's fad is index funds that track the Standard & Poor's 500. True, the average soundly beat most stock funds over the past decade. But is this an eternal truth or a transitory one?"

"In small stocks, especially, you're probably better off with an active manager than buying the market."

"The case for passive management rests only on complex and unrealistic theories of equilibrium in capital markets."

"Any graduate of the ____ Business School should be able to beat an index fund over the course of a market cycle."

Statements such as these are made with alarming frequency by investment professionals.[1] In some cases, subtle and sophisticated reasoning may be involved. More often (alas), the conclusions can only be justified by assuming that the laws of arithmetic have been suspended for the convenience of those who choose to pursue careers as active managers.

If "active" and "passive" management styles are defined in sensible ways, it *must* be the case that

(1) before costs, the return on the average actively managed dollar will equal the return on the average passively managed dollar and
(2) after costs, the return on the average actively managed dollar will be less than the return on the average passively managed dollar.

These assertions will hold for *any* time period. Moreover, they depend *only* on the laws of addition, subtraction, multiplication and division. Nothing else is required.

Of course, certain definitions of the key terms are necessary. First a *market* must be selected—the stocks in the S&P 500, for example, or a set of "small" stocks. Then each investor who holds securities from the market must be classified as either *active* or *passive*.

- A *passive investor* always holds every security from the market, with each represented in the same manner as in the market. Thus if security X represents 3 per cent of the value of the securities in the market, a passive investor's portfolio will have 3 per cent of its value invested in X. Equivalently, a passive manager will hold the same percentage of the total outstanding amount of each security in the market.[2]
- An *active investor* is one who is not passive. His or her portfolio will differ from that of the passive managers at some or all times. Because active managers usually act on perceptions of mispricing, and because such perceptions change relatively frequently, such managers tend to trade fairly frequently—hence the term "active."

Over any specified time period, the *market return* will be a weighted average of the returns on the securities within the market, using beginning market values as weights.[3] Each passive manager will obtain precisely the market return, before costs.[4] From this, it follows (as the night from the day) that the return on the average actively managed dollar *must* equal the market return. Why? Because the market return must equal a weighted average of the returns on the passive and active segments of the market. If the first two returns are the same, the third must be also.

This proves assertion number 1. Note that only simple principles of arithmetic were used in the process. To be sure, we have seriously belabored the obvious, but the ubiquity of statements such as those quoted earlier suggests that such labor is not in vain.

To prove assertion number 2, we need only rely on the fact that the costs of actively managing a given number of dollars will exceed those of passive management. Active managers must pay for more research and must pay more for trading. Security analysts (e.g., the graduates of prestigious business schools) must eat, and so must brokers, traders, specialists and other market-makers.

Because active and passive returns are equal before cost, and because active managers bear greater costs, it follows that the after-cost return from active management *must* be lower than that from passive management.

This proves assertion number 2. Once again, the proof is embarrassingly simple and uses only the most rudimentary notions of simple arithmetic.

Enough (lower) mathematics. Let's turn to the practical issues.

Why do sensible investment professionals continue to make statements that seemingly fly in the face of the simple and obvious relations we have described? How can presented evidence show active managers beating "the market" or "the index" or "passive managers"? Three reasons stand out.[5]

1. Footnotes appear at end of article.

- First, the passive managers in question may not be truly passive (i.e., conform to our definition of the term). Some index fund managers "sample" the market of choice, rather than hold all the securities in market proportions. Some may even charge high enough fees to bring their total costs to equal or exceed those of active managers.
- Second, active managers may not fully represent the "non-passive" component of the market in question. For example, the set of active managers may exclude some active holders of securities within the market (e.g., individual investors). Many empirical analyses consider only "professional" or "institutional" active managers. It is, of course, possible for the average professionally or institutionally actively managed dollar to outperform the average passively managed dollar, after costs. For this to take place, however, the non-institutional, individual investors must be foolish enough to pay the added costs of the institutions' active management via inferior performance. Another example arises when the active managers hold securities from outside the market in question. For example, returns on equity mutual funds with cash holdings are often compared with returns on an all-equity index or index fund. In such comparisons, the funds are generally beaten badly by the index in up markets, but sometimes exceed index performance in down markets. Yet another example arises when the set of active mangers excludes those who have gone out of business during the period in question. Because such managers are likely to have experienced especially poor returns, the resulting "survivorship bias" will tend to produce results that are better than those obtained by the average actively managed dollar.
- Third, and possibly most important in practice, the summary statistics for active managers may not truly represent the performance of the average actively managed *dollar*. To compute the latter, each manager's return should be weighted by the dollars he or she has under management at the beginning of the period. Some comparisons use a simple average of the performance of all managers (large and small); others use the performance of the median active manager. While the results of this kind of comparison are, in principle, unpredictable, certain empirical regularities persist. Perhaps most important, equity fund managers with smaller amounts of money tend to favor stocks with smaller outstanding values. Thus, *de facto*, an equally weighted average of active manager returns has a bias toward smaller-capitalization stocks vis-a-vis the market as a whole. As a result, the "average active manager" tends to be beaten badly in periods when small-capitalization stocks underperform large-capitalization stocks, but may exceed the market's performance in periods when small-capitalization stocks do well. In both cases, of course, the average actively managed *dollar* will underperform the market, net of costs.

To repeat: Properly measured, the average actively managed dollar must underperform the average passively managed dollar, net of costs. Empirical analyses that appear to refute this principle are guilty of improper measurement.

This need not be taken as a counsel of despair. It is perfectly possible for *some* active managers to beat their passive brethren, even after costs. Such managers must, of course, manage a minority share of the actively managed dollars within the market in question. It is also possible for an investor (such as a pension fund) to choose a set of active managers that, collectively, provides a total return better than that of a passive alternative, even after costs. Not all the managers in the set have to beat their passive counterparts, only those managing a majority of the investor's actively managed funds.

An important corollary is the importance of appropriate *performance measurement*. "Peer group" comparisons are dangerous. Because the capitalization-weighted average performance of active managers will be inferior to that of a passive alternative, the former constitutes a poor measure for decision-making purposes. And because most peer-group averages are not capitalization-weighted, they are subject to additional biases. Moreover, investing equal amounts with many managers is not a practical alternative. Nor, *a fortiori*, is investing with the "median" manager (whose identity is not even known in advance).

The best way to measure a manager's performance is to compare his or her return with that of a *comparable passive alternative*. The latter—often termed a "benchmark" or "normal portfolio"—should be a *feasible* alternative identified *in advance* of the period over which performance is measured. Only when this type of measurement is in place can an active manager (or one who hires active managers) know whether he or she is in the minority of those who have beaten viable passive alternatives.

Footnotes

1. The first two quotations can be found in the September 3, 1990 issue of *Forbes*.
2. When computing such amounts, "cross-holdings" within the market should be netted out.
3. Events such as mergers, new listings and reinvestment of dividends that take place during the period require more complex calculations but do not affect the basic principles stated here. To keep things simple, we ignore them.
4. We assume here that passive managers purchase their securities before the beginning of the period

in question and do not sell them until after the period ends. When passive managers do buy or sell, they may have to trade with active managers; at such times, the active managers may gain from the passive managers, because of the active managers' willingness to provide desired liquidity (at a price).

5. There are others, such as differential treatment of dividend reinvestment, mergers and acquisitions, but they are typically of less importance.

Ten Commandments of Financial Statement Analysis

by William H. Beaver, Joan E. Horngren Professor of Accounting, Graduate School of Business, Stanford University

Individuals must pass a proficiency test before obtaining a driver's license. By contrast, investors need not pass any proficiency test before trying to use financial statements as part of their investment analysis. Investors are not required to have taken a course in accounting or financial statement analysis. They are not required even to have read or understood books written on the subject. Yet analyzing financial statements requires at least as much knowledge and skill as driving an automobile. Perhaps each financial statement should contain a warning to potential users, similar to those found on many products. The warning would include at least the following 10 commandments.

1. Thou shalt not use financial statements in isolation, but only in the broader context of other available information. The additional information includes data on economy-wide conditions and industry-wide conditions.

2. Thou shalt not use financial statements as the only source of firm-specific information. There are many other sources of information about the company. Consider, for example, the popular financial press and periodicals, as well as analysts' reports.

3. Thou shalt not avoid reading footnotes, which are an integral part of financial statements. Financial statements cannot be reasonably analyzed without reading and understanding the footnotes. By analogy, a temperature of 10 degrees is meaningless in isolation, unless one knows whether it is being measured on the Celsius or Fahrenheit scale. In a given country, a uniform temperature scale may be assumed. The same is not true of the accounting methods used under generally accepted accounting principles. GAAP, for example, permits a variety of inventory and depreciation methods. A description of a company's accounting policies is included as a part of the footnotes.

4. Thou shalt not focus on a single number. The investor should read and understand *all* the material presented in the financial statements. Financial statements are not designed to be reduced to a single number. Net income is not intended to be *the* number that summarizes *all* the information relevant to making an investment decision. A user must analyze growth and leverage, among other factors, as well as profitability.

5. Thou shalt not overlook the *implications* of what is read. It is not sufficient simply to know that a company is a high-growth firm or a highly leveraged firm; one must also know that such characteristics typically imply higher risk, as well.

6. Thou shalt not ignore events subsequent to the financial statements. Financial statements are not forecasts of the future. The annual financial statements report the financial condition of the company as of year-end. They do not purport to capture the effects of events that occur after year-end. They thus become increasingly out-of-date as the year progresses. The rate of deterioration in timeliness is related to many factors, including the growth rate of the firm.

7. Thou shalt not overlook the limitations of financial statements. Financial statements report on only a specified set of events, not all events or all possible financial effects of a single event. Financial statements do not generally represent estimates of the market values of the reported assets and liabilities, nor do they reflect changes in the market values of those assets and liabilities.

8. Thou shalt not use financial statements without adequate knowledge. Investors should be sufficiently competent to read, understand and analyze financial statements. Otherwise, the investor cannot be called a user of financial statements in any meaningful sense.

9. Thou shalt not shun professional help. If unwilling or unable to attain adequate knowledge, the investor should defer to someone who does have such ability, such as a financial analyst. If unwilling or unable to obtain help, the investor should hand over a portion of the investment process (hence a portion of the investment decision itself) to a professional manager.

10. Thou shalt not take unnecessary risks. If unwilling or unable to obtain professional help, the investor should undertake investments where investment risk is minimal, or where analysis of financial statements is not an issue. Investment in U.S. Treasury bills is one example.

Of course, there may be more than 10 commandments for financial statement analysis, but these capture the primary issues.

Concluded on page 18.

Attachment B

Investment Management Fees Are (Much) Higher Than You Think

GUEST EDITORIAL

Charles D. Ellis, CFA
Chairman of the Whitehead Institute
Cambridge, Massachusetts

Investment Management Fees Are (Much) Higher Than You Think

Although some critics grouse about them, most investors have long thought that investment management fees can best be described in one word: low. Indeed, fees are seen as so low that they are almost inconsequential when choosing an investment manager. This view, however, is a delusion. Seen for what they really are, fees for active management are high—much higher than even the critics have recognized.

When stated as a percentage of *assets*, average fees do look low—a little over 1% of assets for individuals and a little less than one-half of 1% for institutional investors. But the investors already own those assets, so investment management fees should really be based on what investors are getting in the *returns* that managers produce. Calculated correctly, as a percentage of returns, fees no longer look low. Do the math. If returns average, say, 8% a year, then those same fees are not 1% or one-half of 1%. They are much higher—typically over 12% for individuals and 6% for institutions.

But even this recalculation substantially understates the *real* cost of active "beat the market" investment management. Here's why: Index funds reliably produce a "commodity product" that ensures the market rate of return with no more than market risk. Index funds are now available at fees that are very small: 5 bps (0.05%) or less for institutions and 20 bps or less for individuals. Therefore, investors should consider fees charged by active managers not as a percentage of total returns but as *incremental* fees versus risk-adjusted *incremental* returns above the market index.

Thus (correctly) stated, management fees for active management are remarkably *high*. Incremental fees are somewhere between 50% of incremental returns and, because a majority of active managers fall short of their chosen benchmarks, *infinity*. And when market returns are low, as in recent years, management fees eat up even more of an investor's return. Are any other services of any kind priced at such a high proportion of client-delivered value? Can active investment managers continue to thrive on the assumption that clients won't figure out the reality that, compared with the readily available passive alternative, fees for active management are astonishingly high?

Fees for active management have a long and interesting history. Once upon a time, investment management was considered a "loss leader." When pension funds first mushroomed as "fringe benefits" during the post–World War II wage-and-price freeze, most major banks agreed to manage pension fund assets as a "customer accommodation" for little or no money—that is, no explicit fee. With fixed-rate brokerage commissions, the banks exchanged commissions for cash balances in agreed proportions. The brokers got "reciprocal" commission business, and the banks got "free" balances they could lend out at prevailing interest rates. In the 1960s, a few institutional brokerage firms, including DLJ, Mitchell Hutchins, and Baker Weeks, had investment management units that charged full fees (usually 1%) but then offset those nominal fees entirely with brokerage commissions.

When the Morgan Bank took the lead in charging fees by announcing institutional fees of one-quarter of 1% in the late 1960s, conventional Wall Street wisdom held that the move would cost the bank a ton of business. Actually, it lost only one account. Thus began nearly a half century of persistent fee increases, facilitated by client perceptions that fees were comfortably exceeded by incremental returns—if the right manager was chosen. Even today, despite extensive evidence to the contrary, both individual and institutional investors typically expect their chosen managers to produce significantly higher-than-market returns. That's why fees have seemed "low."

A relatively minor anomaly is getting more attention: While asset-based fees have increased substantially over the past 50 years—more than fourfold for both institutional and individual

Guest Editorial is an occasional feature of the Financial Analysts Journal.
This piece reflects the views of the author and does not represent the official views of the FAJ *or CFA Institute.*

Attachment C

How consistently do active managers win?

Available in Hard Copy Only

Attachment D

The Paradox of Skill

The Paradox of Skill

Why Greater Skill Leads to More Luck

Michael Mauboussin

ChangeThis 100.03

Variation in batting averages must decrease as improving play eliminates the rough edges that great players could exploit, and average performance moves toward the limits of human possibility.

—STEPHEN JAY GOULD

Okay, you have gotten the memo on improving skill: 10,000 hours, hard work, deliberate practice, grit, and attentive teacher. We've all heard it. You also recognize that in many of life's activities, the results you achieve combine skill and luck. No debate there. Now, what if I told you that in many cases improving skill leads to results that rely more on luck? That's right. Greater skill doesn't decrease the dependence on luck, it increases it. If you have an interest in sports, business, or investing, this lesson is for you.

Stephen Jay Gould was a renowned evolutionary biologist at Harvard University who loved to write about baseball. One of his best essays was about why no player in Major League Baseball had maintained a batting average of more than .400 for a full season since Ted Williams hit .406 in 1941. Gould considered several conventional explanations, including more night games, demanding travel, improved fielding, and more extensive use of relief pitching. None checked out.

Maybe Williams was some sort of freak player, Gould thought, better than all of those who came before him as well as all of those who followed. That's implausible, he concluded, because in every sport where performance is measured versus a clock, including swimming and running, athletes have improved. Baseball players, too, are better than they were in the past: faster, stronger, more fit, and better trained.

So how do we solve the mystery of the vanishing .400 hitter? The best approach is to set up a simple model that explains how greater skill can lead to a greater reliance on luck. We'll then apply our model to other realms to see if it explains what we see there. In each case, we'll see that luck has more sway even as participants hone their skill. It's the paradox of skill.

> *"What if I told you that in many cases improving skill leads to results that rely more on luck? That's right. Greater skill doesn't decrease the dependence on luck, it increases it.*

The Jars of Success

Imagine two jars, one representing skill and the other luck, that are each filled with cards with numbers printed on them that comprise a bell curve. Bell curves are defined by a mean, or average, and a standard deviation. From the top of the bell, the curve slopes down the sides symmetrically with an equal number of observations on each side. Standard deviation is a measure of how far the sides of the bell curve are from the average. A skinny bell curve has a small standard deviation and a fat bell curve has a large standard deviation.

So most cards in each jar have values at or near the mean, and a few cards are marked with numbers that have values far from the mean. To determine an outcome, you draw one number from the skill jar, one from the luck jar, and add them. Relating this to batting averages, you could say that a player has a certain amount of hitting skill—the number he drew from that jar—and some luck. A great player can have an unlucky season that results in a batting average below his true skill, or a below-average player can enjoy substantial luck and hit at an average that overstates his skill. Hitting .406 as Williams did requires tremendous skill and terrific luck. He drew numbers from both jars that were far above average.

Let's put some numbers to the averages in each jar. Let's start with the luck jar. While for a season some players will have good luck and others bad luck, we can safely assume that luck is zero on average. That says that the average of the skill jar will approximate the batting average for all of the players combined, which has vacillated around .260-.270 in the last 75 years or so. The reason that average skill hasn't gone up, even though the hitters today are better than in the past, is that batting average represents a duel between pitcher and hitter. If pitchers and hitters improve roughly in lockstep, the overall skill can improve sharply even as the batting average remains steady. The arms war (pun intended) between pitchers and hitters creates the illusion of stability even as the players improve.

Here was Gould's crucial insight: the standard deviation of skill has gone down over time. Imagine the bell curve going from being fat to skinny. The extreme values are closer to the average. So even if the luck distribution doesn't change a bit, you should expect to see the standard deviation of batting averages decline over time. And that is precisely what Gould showed. The standard deviation of batting averages was .0326 in the 1940s, when Williams achieved the feat, and was .0274 in the first decade of the 2000s. In statistical terms, hitting .380 in 2011 is the equivalent to the .406 that Ted Williams hit 70 years earlier.

Why did the range of skill from the best to the worst narrow so much? Two factors can explain a great deal of the phenomenon. When professional baseball began, it drew only white players from the Northeastern part of the U.S. But over time, the league began recruiting players of all races, from all parts of the U.S., and eventually from all around the world. This greatly expanded the pool of talent. Hungry players from the Dominican Republic, Venezuela, and Japan brought a new level of skill to the game. In addition, training has improved greatly since the 1940s, which has certainly had an effect on this convergence of skills. Combine more access to talented players with sharpened training techniques and you get a higher, and more uniform, level of skill throughout the league.

That the bell curve in the skill jar gets skinnier over time while the bell curve in the luck jar remains the same means that as skill improves for the population, luck becomes more important in determining results. On average, players have greater skill today than they did in years past but their outcomes are more tied to luck. This extends to other realms as well.

A good theory makes predictions that we can test. The paradox of skill says that in fields where there is no offsetting interaction (for example, pitcher versus hitter) and no luck, we should see absolute results improve and relative results cluster. This is precisely what we see in events such as swimming and track and field.

Naturally, human physiology limits absolute performance—a man can run only so fast and a woman can swim only so swiftly. But we see improvement and convergence broadly. For example, the winning time for the men's Olympic marathon dropped by more than 23 minutes from 1932 to 2012. As revealing, the difference between the time for the winner and the man who came in 20th shrunk from 39 minutes to 7 ½ minutes over the same period. Luck and interaction can partially obscure the paradox of skill, but the core elements are there in case after case.

> *“A good theory makes predictions that we can test. The paradox of skill says that in fields where there is no offsetting interaction (for example, pitcher versus hitter) and no luck, we should see absolute results improve and relative results cluster.*

Business: If You're Not Getting Better, You're Getting Worse

Now let's take a look at the business world. It's important to start with the acknowledgement that luck plays a large role in the results for business. Just as in baseball, where the difference between a hit and an out might be six inches of flight trajectory, business has a lot of randomness.

There are a few sources of that randomness. For one, you never know what your competitors are going to do. Sometimes companies compete in an orderly fashion and the outcome is good for the industry. Other times competitors may develop a strategy to drop prices, or add capacity, that forces a reaction. So even if you know what your plans are, you don't know those of your competitors. Game theory is a branch of economics that studies how players act and react to one another, and as you add players to the competition, the unpredictability rises quickly.

Customers are another source of randomness in business. Naturally, companies spend lots of time and effort anticipating what their customers want and need, but the success rate of new products shows that there's no easy way to do so. And even if a company can decipher its competitors and customers, it has to deal with changes based on technology. Consider the media business: how many executives in the newspaper, radio, and television industries properly anticipated the changes of the last couple of decades? Who knows where things are going from here? Business has its own version of the luck jar, and there's a wide range of numbers.

What about skill? The paradox of skill teaches a couple of lessons that executives sometimes ignore. Phil Rosenzweig, a professor at IMD Business School, provides a concrete example. In the mid-1990s, a large U.S. retailer set out the goal of improving its inventory turnover ratio, a crucial measure of capital efficiency. And its effort proved to be a rousing success, as its turns went from 3.4 times in 1994 to 4.6 times in 2002. Indeed, you might envision the board promising and delivering management bonuses based on such a nice improvement.

Here's the problem: the retailer's number one competitor also happened to be focused on inventory turnover and was able to take its ratio from 5.1 times to 8.1 times during the same period. So even as the first retailer strengthened its absolute performance, its relative position weakened. This is one of the lessons of the paradox of skill. Getting better in an absolute sense doesn't matter if it's offset by the competition. Hitters today are much better than they were in the past, but so are the pitchers. The improvement is obscured by the interaction. Likewise, the first retailer was better in 2002 than it was in 1994 but it actually lost ground relative to its prime competitor.

Research has pointed out the variance of quality in consumer goods has narrowed over time, another finding that's consistent with the paradox of skill. In years past, companies offered products across a wide spectrum of quality, and prices by and large reflected that quality gap. For instance, some automobiles were cheap and shoddy, and others were expensive but well made.

Over time, the gap in quality has narrowed. As a consequence, customers now rely less on price-quality trade-offs and more on other variables, including convenience, after-sale service, and store location. This can enhance the role of luck in securing the sale. In business as in baseball, the skill distribution has likely tightened allowing luck to play a growing role in outcomes.

> *“Getting better in an absolute sense doesn't matter if it's offset by the competition. … The improvement is obscured by the interaction.*

Investing: A Random Walk Because of Skill

Perhaps nowhere is the paradox of skill more evident than in the world of investing. Luck is such a big deal that one of the industry's all-time best-selling books is called A Random Walk Down Wall Street. But it is a random walk only because investors are so collectively skillful. Companies, analysts, the government, and the media disseminate gobs of information that investors quickly incorporate into prices. Advances in technology mean that there is massive computing power available to crunch numbers. And the spoils of success are sufficiently high that many of the best and brightest students are drawn to the investment world.

The challenge is that most investment firms have access to the same information, whiz-bang computers, and sharp graduate students. Those things don't set you apart. Since stock prices generally reflect all of the information that's out there, it's only new information that moves prices. And because by definition you can't predict new information, stock prices tend to follow a random walk. The random walk story is not exactly true, but its emphasis on how hard it is to beat the market is well placed. Similar to baseball and business, as skill increases luck becomes more important.

> *“ Even if it is not perfectly linear, improvement over time occurs. Efficiency grinds upward.*

But in one important respect, investing is quite different than those fields. Take sports as a starting point. As time goes on, the ability of the athletes marches inexorably toward the limit of human performance. That last bit of performance improvement is hard fought because there's only so much a body can do. It's also why some athletes turn to chemicals to enhance their performance. But even if it is not perfectly linear, improvement over time occurs. Efficiency grinds upward.

Now compare sports to investing. In investing, efficiency means that value and price are one and the same. The price of a stock accurately reflects the present value of all of the cash flows in the future and news is rapidly and accurately assimilated. Indeed, economists have done lots of experiments to show that a group of investors will settle on an efficient price under normal conditions. The problem is the conditions are not always normal in investing. From time to time, investors follow one another in a herd, leading to prices that veer far from value.

The euphoric dot.com bubble that peaked in early 2000 and the acute fear that created a market low in early 2009 are but two recent examples. In these cases, it's still hard to beat the market but you'd be hard pressed to say that the market is efficient.

What To Do About the Paradox of Skill

We may never see another .400 hitter in professional baseball. That's alright. It reflects "the spread of excellence," using Stephen Jay Gould's phrase. You'll see skill increasing and luck becoming more important in shaping results in many places that you look. So, what should you do about it? Here are three suggestions:

- **Find realms where the variance of skill is still wide.** If you compete in a field where the range of skill is wide, the more skillful will succeed at the expense of the less skillful.

Investing is a good case in point. In developed markets, large and sophisticated institutional investors dominate the trading scene. The skillful players compete with one another and it's hard to gain an edge. In some developing markets, by contrast, large institutions compete with less sophisticated individuals. Research shows that, on average, the institutions earn excess returns at the expense of the individuals. But it's not always easy to know if you're the most skillful player. Warren Buffett, the famous investor and chairman and CEO of Berkshire Hathaway, makes the point in the context of poker: "If you've been playing poker for half an hour and you still don't know who the patsy is, you're the patsy."

➔ **Think relative, not absolute.** Essential to the paradox of skill is the idea that you can measure improvement in skill either on an absolute scale or relative to competitors. In activities where there is no direct interaction or luck—say, a 100-meter dash—absolute skill is all that matters. But when there is interaction and luck, you have to measure relative performance. Here's why this is so important, using business as an example. There are a slew of best-selling books that offer a simple formula for corporate performance improvement. These miss the mark because they fail to consider what competitors may do. Results are a combination of your actions with those of your rivals. If all companies are getting better in lockstep, no company is gaining an edge.

➔ **Focus on process, not outcome.** If you want to become world-class as a violinist or a chess player, areas where little luck is involved, you need roughly 10,000 hours of deliberate practice. What's crucial is that your results, as you improve, will be a reliable indicator of your skill. As a result, feedback in these domains can be clear and unequivocal. If you compete in a field where luck plays a role, you should focus more on the process of how you make decisions and rely less on the short-term outcomes. The reason is that luck breaks the direct link between skill and results—you can be skillful and have a poor outcome and unskillful and have a good outcome. Think of playing blackjack at a casino. Basic strategy says that you should stand—not ask for a hit—if you are dealt a 17. That's the proper process, and ensures that you'll do the best over the long haul. But if you ask for a hit and the dealer flips a 4, you'll have won the hand despite a poor process. The point is that the outcome didn't reveal the skill of the player, only the process did. So focus on process.

One final thought. Once you've embraced the paradox of skill, you'll see that it's appropriate to have an attitude of equanimity toward luck. If you've done everything you can to put yourself in a position to succeed, you should accept whatever results appear. Some days you'll be lucky, and the results will exceed your expectations. Some days the results will be disappointing because of bad luck. **The best plan will be to pick yourself up, dust yourself off, and get ready to do it again tomorrow.**

Attachment E

What Practitioners need to know

What Practitioners Need to Know. . .

by Bob Hagin
LTCB-MAS Investment Management

In addition to being relevant to practitioners, articles in Financial Analysts Journal *are required to meet high standards of academic excellence. As a byproduct, readers are frequently faced with terms such as "standard error of the estimate" or "t-test" that—even though possibly studied years ago—are certainly not everyday terms in an investment practitioner's world. We introduce here a new column designed to provide short "English-language" explanations of "What Practitioners Need to Know." For our first topic, we have selected t-tests.*

. . . About t-Tests

Suppose an investor has outperformed the average active manager by 2 per cent each year for the past four years. What can we infer about this investor's skill?

There is an old anecdote, "If you place an infinite number of monkeys in front of an infinite number of typewriters, one of them will type the full text of *King Lear*." Similarly, there is some probability that beating the average manager by 2 per cent per year for four years could happen by pure chance. Thus we know enough to be cautious about immediately inferring that this investor has skill (just as we would likely not bestow the name "Shakespeare" on the monkey that, by pure chance, typed *King Lear*).

What we need in a situation such as this is a way to estimate the likelihood that what we observe (in this case, an investor who has outperformed the average manager by an average of 2 per cent per year for four years) has *not* occurred by chance. Intuitively, we know that our confidence in saying that someone has used skill to outperform the other managers by 2 per cent per year over the past four years depends on the distribution of annual active returns.

If, on one hand, the returns of active managers vary so widely that, on average, two-thirds of all managers fall within a range of plus or minus 10 per cent of the average manager, we would not be very confident that someone who was above the mean by 2 per cent did so by skill. On the other hand, if the returns of active managers are so tightly clustered around the mean that, on average, two-thirds of all managers fall within a range of plus or minus 1 per cent of the average return, we would certainly have more confidence that a manager who beat the average by 2 per cent per year for four years did so because of skill. Statistical tests, such as the t-test, allow us to quantify these intuitions.

Normal Distributions

Many statistical tests assume that data are normally distributed. One of the most important features of a *normal distribution* is that it is completely described by its mean and standard deviation.

In brief, standard deviation measures variation around an average. Say the average of 30 monthly rates of return is 15 per cent, and the average of the returns' deviations from this average is 10 per cent. Conceptually, the returns between 5 per cent (15–10) and 25 per cent (15 + 10) fall within *one* standard deviation of the average return. Similarly, the returns between −5 per cent [15 − (2 × 10)] and 35 per cent [15 + (2 × 10)] fall within *two* standard deviations of the average.

In a normally distributed sample, approximately 68 per cent of the values are within one standard deviation of the mean, approximately 95 per cent of the values are within two standard deviations of the mean, and more than 99 per cent of the values are within three standard deviations of the mean.

How can we use this knowledge to gauge our confidence in an investor's skill? Consider a broad population of stocks with a normal distribution of returns. The return on a portfolio of

stocks selected from this population is not likely to have the same average return as the population itself. After all, the portfolio holds a smaller number of stocks, and the weights of the stocks in the portfolio may differ from their weights in the overall population.

We can, however, determine the range of probable mean returns for a sample of stocks drawn randomly from the overall population. We noted earlier that 95 per cent of normally distributed observations fall within a range of plus or minus two standard deviations from the observations' mean. There is thus only a 5 per cent chance that the mean of a sample portfolio drawn randomly from a normally distributed population will fall outside a band defined by plus two and minus two standard deviations from the population mean.

Such bands can be used to define *confidence intervals*—intervals within which we are likely (with 68 per cent probability, or 95 per cent probability, or 99 per cent probability) to find the means of random samples drawn from a normally distributed population. The calculation of confidence intervals allows us to quantify our confidence that what we observe are true differences and not merely observations that are likely to occur by mere chance.

If we simplify our problem to the case of an investor whose performance for one year was 2 per cent above the mean performance of a population of similar managers, we merely have to determine where this 2 per cent lies in relation to the returns within a given confidence interval. If 2 per cent lies within the confidence interval, we would conclude that this level of return could easily be the result of chance. If 2 per cent lies outside this confidence interval, we would conclude that this level of return is not likely to be the result of chance.

In our simplified case, in which we know the standard deviation of the population, we can *standardize* our sample mean (which in this case is only one measurement) by calculating a *z-score*. This can be done by dividing the difference between the sample mean and the population mean by the population's standard deviation:

$$z = \frac{\text{Sample Mean} - \text{Population Mean}}{\text{Population Standard Deviation}}.$$

Assume that the average active return from the population of investment managers is zero, the standard deviation is 6 per cent and our investor's "sample mean" is 2 per cent. We thus obtain a z-score of 0.67—(2-0)/6. This means that (given our simplifying assumptions) an active return of 2 per cent falls only 0.67 standard deviation above the mean—well short of the two standard deviations that contain 95 per cent of the observations. Thus our investor's 2 per cent active return is within the 95 per cent confidence interval and likely the result of chance, not skill.

Our problem becomes slightly more complicated when we (more realistically) measure our manager over more than one year. Technically, instead of determining the probability of one observation, we need to compare the means of two distributions (the mean of our population and the mean of our manager). To complicate matters further, in practice we usually need to make inferences from a relatively small sample size.

This was the problem that W. S. Gosset faced around the turn of the century. If we think of the dreary working conditions portrayed so well by Charles Dickens, it is easy to imagine Gosset's problem. Gosset—a chemist at the Guinness Breweries—was asked to make inferences about the quality of various brews. But Gosset had two problems.

First, quite understandably, Guinness was unwilling to supply Gosset with a large number of samples. But this limitation on sample size spurred Gosset to an important discovery. He found that, when working with small samples, errors were introduced unless the normal distribution was replaced with a distribution that had more variability (and a higher probability of large deviations).

Having discovered something of great importance to the scientific community, Gosset faced a second problem: Guinness prohibited him from using his name to publish the results of his on-the-job discovery. Undaunted, and believing in the importance of his discovery, Gosset published his findings anonymously under the pen name "Student." Statisticians have ever since been introduced to "Student's t" or, as it has come to be known, the "t-test." The t-test is especially important for financial researchers who—working with annual returns—share Gosset's problem of being forced to work with small samples.

The t-Test

The ingredients of a t-test are quite intuitive. First, we expect to have more confidence in statistics derived from large samples than in statistics derived from small samples. We might expect, however, that an increase in sample size of 10 will have a larger effect on our confidence when the sample goes from 10 to 20 than when it goes from 90 to 100. We will skip the mathematics, but it turns out that the confidence interval around the mean is governed by the square root of the sample size. In the case of the 95 per cent confidence interval defined by plus or minus two standard deviations from the mean:

$$95\% \text{ Confidence Interval} = \frac{2 \times \text{Standard Deviation of Sample}}{\text{Square Root of Sample Size}}.$$

The "2" in front of the "standard deviation of sample" refers to the two standard deviations associated with the 95 per cent confidence interval. Note that an increase in sample size from four to nine has the same relative effect as an increase from nine to 16.

Let's take an easy example. Say that the market-relative return of all managers has a mean of zero, our sample size is four, and its standard deviation is 6 per cent (which is typical). Plugging these values into the foregoing equation, we have:

$$95\% \text{ Confidence Interval} = 2\% \pm \frac{2 \times 6\%}{2}$$

$$= +8 \text{ to } -4.$$

Thus, if we calculated the active return above or below the mean return for a four-year random sample, and repeated the experiment 100 times, we would expect 95 per cent of these returns to be within the 95 per cent confidence interval. Note that, although the confidence interval will bob around for each sample, 95 per cent of the intervals so formed will capture the true mean of the population.

What would happen if the standard deviation or the size of the sample changed? If we cut standard deviation in half—to 3 per cent—the confidence interval would also be halved; it would then range from minus 1 to plus 5 per cent [$2\% \pm (2 \times 3\%/2)$]. Similarly, if we increased the sample size from four to 16 years, the confidence interval would be reduced to from minus 1 to plus 5 per cent [$2\% \pm (2 \times 6\%/4)$].

How confident can we be that an observed value is not a chance deviation from the mean? If we do not know the standard deviation of the underlying population, or if we have a small sample (less than 30), we can answer this question directly with a t-test.

The t-distribution is more "spread out" than the normal distribution. Its exact form is a function of the sample size (degrees of freedom)—the smaller the sample, the more spread out the distribution becomes.

The equation for t contains the three things that intuitively affect our confidence—the mean of the active return (the larger the active return, the more confidence), the standard deviation of the active return (the more variable the return, the more likely we are to conclude that an above-average return occurred by chance) and the number of years (the more years, the more confidence). Specifically:

$$t = \frac{\text{Active Return} - \text{Average Return}}{\dfrac{\text{Standard Deviation of Active Return}}{\text{Square Root of Number of Years}}}.$$

A useful rule of thumb is that a t-statistic must be at least 2.0 to be significant. To be more specific, the t-distribution (following Gosset's insight) depends on the sample size. By consulting a table of t-statistics for various sample sizes, we find that with a sample size of four we need a t-statistic of at least 2.35 to conclude that there is only a 5 per cent chance that beating the market by 2 per cent per year for four years was an accident. Because the t-distribution becomes less spread out as we increase sample size, the t-statistic required for statistical significance decreases as sample size increases. When the sample size reaches 30, the t-distribution is very close to the normal distribution.

The data in Table I show that when the annual standard deviation of the active investment return is 6 per cent, we cannot reject the hypothesis that our manager's average active return is zero until the manager outperforms the average manager by 2 per cent for 25 *years*. At this point, the required t and the calculated t are approximately equal.

Table I Effect of Sample Size on Required t-Statistic

Sample Size	*Required t*	*Calculated t*	*Active Return*	/	*(Std. Dev.*	/	$\sqrt{Yrs.}$)
4	2.35	0.67	2.0	/	(6.0	/	2)
9	1.86	1.00	2.0	/	(6.0	/	3)
16	1.75	1.33	2.0	/	(6.0	/	4)
25	1.71	1.67	2.0	/	(6.0	/	5)
36	1.69	2.00	2.0	/	(6.0	/	6)

The foregoing example illustrates two things. First, it is important *not* to rely on intuition in estimating confidence. In this example, as in many studies contained in *Financial Analysts Journal*, we need an objective estimate of the likelihood that the results are sensible. t-tests—one of the major ways to estimate that likelihood—tell us the probability of being able to rely on the data used to make the analysis. Second, the example was selected to illustrate the differences between the level of confidence that would be acceptable to many practitioners (beating the market by an average of 2 per cent for four years) and the objective standards to which *Financial Analysts Journal* holds its researchers.

Current Issues *concluded from page 12.*

returns. Note, however, that the hierarchy of risk and returns conformed to expectations until 1989. Both observations can be helpful in understanding the fixed income securities market.

Footnotes

1. See E. Altman, "The Anatomy of the High-Yield Bond Market," *Financial Analysts Journal*, July/August 1987 and *Default Risk, Mortality Rates, and the Performance of Corporate Bonds* (Charlottesville, VA: Research Foundation of the Institute of Chartered Financial Analysts, 1990); and M. Blume and D. Keim, "Lower-Grade Bonds: Their Risks and Returns," *Financial Analysts Journal*, July/August 1987.
2. E. Altman and S. Nammacher, *Investing in Junk Bonds: Inside the High Yield Debt Market* (New York: John Wiley, 1987).
3. R. Ambarish and M. Subrahmanyam, "Defaults and the Valuation of High Yield Bonds," in E. Altman, ed., *The High Yield Debt Market: Investment Performance and Economic Impact* (Homewood, IL: Dow Jones-Irwin, 1990) and R. Bookstaber and R. Clarke, "Problems in Evaluating the Performance of Portfolios with Options," *Financial Analysts Journal*, January/February 1985.
4. J. Fons, "Default Risks and Duration Analysis," in *The High Yield Debt Market, op. cit.*
5. We used the mortality methodology in E. Altman, "Measuring Corporate Bond Mortality and Performance" (Working paper, Salomon Brothers Center, New York University, February 1988 and June 1989) and *Journal of Finance*, September 1989. This was updated with data from Altman, "Default Risk," *op. cit.*
6. See Altman, "Measuring Corporate Bond Mortality," *op. cit.* and P. Asquisth, D. Mullins and E. Wolff, "Original Issue High Yield Bonds: Aging Analysis of Defaults, Exchanges and Calls," *Journal of Finance*, September 1989.
7. See Altman, "Measuring Corporate Bond Mortality," *op. cit.*

Attachment F

Financial Analysts Journal Commentaries

Re: "Murder on the Orient Express: The Mystery of Underperformance"
by Charles D. Ellis, CFA
Financial Analysts Journal, July/August 2012, Vol. 68, No. 4:13-19.

A Comment

Victor S. Sidhu, CFA, Sidhu Group, Santa Monica, CA; *Financial Analysts Journal*, November/December 2012, Vol. 68, No. 6: 11–12

Exactly 37 years ago, on a summer's day in August, I sat back to enjoy the July/August 1975 issue of the Financial Analysts Journal. What I expected to be a relaxing moment abruptly turned very disturbing. Shouting out at me from the pages of an article called "The Loser's Game" was one Charles Ellis. I did not like what he was saying. Not at all. As a professional money manager, I believed we investment gurus could beat the market. Ellis bluntly said, "That premise appears to be false." He backed it up. And he converted me.

Now, 37 years later, he's baaack. In the current article, I recognize all the players and schemes and self-denials Charley describes. In the roles of chief investment officer and investment committee member, I have tried constructively over the decades to make right the process, to avoid the "crime." My efforts have met with limited success. I lack Charley's diplomacy and wit, and so my attempts to educate on the truth have been greeted consistently with emotions ranging from polite derision to near-violent hostility.

To Charley, I say bravo for writing this article. To CFA Institute, I commend you for publishing it. To readers, I urge you to study it carefully and take heed. What Charley describes is quite simply billion-dollar fraud. A crime. We had better do something about it. Because if we don't . . . Bernie, move over and make some room, 'cause here come da judge.

A Comment

"Murder on the Orient Express: The Mystery of Underperformance": Peter Drucker *Redux"*

Jeffrey E. Horvitz, Moreland Management Company, Beverly Farms, MA; *Financial Analysts Journal*, November/December 2012, Vol. 68, No. 6:11.

With his usual clarity, Charles Ellis assigned culpability to investment managers, investment consultants, fund executives, and investment committees (July/August). The essential problem is that the fees are not warranted by the service being delivered, and all parties to active management have a hand in the dysfunctional outcome. Ellis's critique is similar to the famous point made by Peter Drucker: Efficiency is doing things right; effectiveness is doing the right thing. The money management business is certainly efficient but woefully ineffective.

Investment management is remarkably similar to roulette, in which the management fees are the house take and the average investor will predictably "lose" (i.e., underperform the

broad market) by the amount of the fees. Investment managers cannot all be above average, but neither are they all average—and cross-sectional dispersion among managers can be large. The more concentrated the manager, the more likely the manager will deviate from the average, simply because of small sample size. This outcome gives the illusion that skill is available.

Most investment managers insist that they add value. Clearly, this is delusional for the majority, but without this fantasy, going to work each day would be too depressing. This behavior is like hamsters running hard and fast on a wheel—going nowhere but doing it with gusto. The reality is that money management is perhaps the only major industry that produces, on a cumulative net basis, almost no financial return (except to the managers themselves).

Very smart people, though often young and inexperienced, go into investment consulting. Because the most important consulting advice is imparted early, the ongoing relationship is of diminishing value. Experienced consultants must develop skill in making the trivial seem important—for example, the style box. Great effort goes into finding managers who can be put into a particular cell in a style matrix (value/growth × large/mid/small cap) and are "consistent." When the style matrix is filled, the result is usually close to an index fund. This exercise is efficient but not usually effective. What consultants can do very well is keep their clients from doing very stupid things; what they can't do is provide advice that leads to truly superior investment performance (otherwise, they would become hedge fund managers).

Pension and endowment executives suffer from the agency problem: It simply isn't their money, but it is their job and salary. A conventional loss is more acceptable than an unconventional loss and is unlikely to result in the loss of one's job; deviating from both the norm and conventional wisdom puts one's personal income at great risk. Fund executives are very much the middle managers discussed by Peter Drucker who spend a lot of time doing things (like reports and committee meetings) that do little to change the bottom line.

Some investment committees are populated with noninvestment people who are there for political or social reasons, not for their investment acumen. Other investment committees comprise mostly investment professionals. Because they are "in the investment business," they are expected to be expert in matters of portfolio management. In fact, most of these people are the very investment managers who charge high fees while delivering little in return. Naturally, investment managers will favor managers whose style is similar to their own. They prefer to do what they know, even if it is not what is needed—just as, for example, cardiologists overwhelmingly recommend treatment with medication whereas cardiac surgeons overwhelmingly recommend surgery.

And just as banks are now too big to fail, the investment business may now be too big to succeed.