

Five Myths of Active Portfolio Management

Most active managers are skilled.

Jonathan B. Berk

Proponents of efficient markets argue that it is impossible to beat the market consistently. In support of their view, they point to the evidence that active managers as a group do not beat the market. Their conclusion is that these investment professionals do not have the skills necessary to pick stock or time the market. Yet if this argument is correct, why do we have active portfolio managers at all?

Even more puzzling is managers' level of compensation. One of the first principles any student of microeconomics learns is that in a competitive market (which the capital markets surely are), people can earn economic rents only if they have a skill that is in short supply. If active managers cannot pick stocks or time the market, what rare skill do they have that makes them among the highest paid members of society?

Even people who allow for the possibility that some managers have skill have been hard pressed to find evidence of this skill in the data. Beyond a year, there is little evidence of performance persistence—managers who do well in one year are no more likely to do well the next year (see Carhart [1997]). This fact is widely interpreted as evidence that the performance of the best managers is due entirely to luck rather than skill (and is thus not repeatable).

As Gruber [1996] notes, the behavior of investors is just as puzzling. Why do investors continue to invest with active managers in the face of this evidence? Yet investors chase returns; a good year induces an inflow of capital, and a bad year induces an outflow of capital. The flow of capital into and out of actively managed mutual funds is sensitive to past fund performance, despite no strong evidence of persistence.

JONATHAN B. BERK is the Harold Furst associate professor of management philosophy and values in the Haas School of Business at the University of California at Berkeley.
berk@haas.berkeley.edu

My objective is to explain these facts and show why they are consistent with a rational and competitive financial market and active managers who have skill and add considerable value.

Let's begin by considering five hypotheses that many investment professionals take on faith:

1. The return investors earn in an actively managed fund measures the skill level of the manager managing that fund.
2. Because the average return of all actively managed funds does not beat the market, the average manager is not skilled and therefore does not add value.
3. If managers are skilled, their returns should persist—they should be able to beat the market consistently.
4. In light of the evidence that there is little or no persistence in actively managed funds' returns, investors who pick funds on the basis of past returns are not behaving rationally.
5. Because most active managers' compensation does not depend on the return they generate, they do not have a performance-based compensation contract.

At first glance, all five hypotheses appear plausible. In fact, I show that in a model in which rational investors compete with each other for the services of value-adding managers, none of them is true—they are five widely held myths.

THOUGHT EXPERIMENT

In Berk and Green [2004], we derive a theory of active portfolio management in an economy in which investors and managers are fully rational. I'll summarize that theory and show why it implies that the five hypotheses are myths.

Imagine an economy with skilled investment managers with differential ability who can generate positive risk-adjusted excess returns. Managers and investors alike know who these managers are.

Assume that managerial ability to generate excess returns cannot be effectively deployed at an arbitrarily large scale. Eventually, the amount of money under management grows so much that each additional dollar contributed reduces the expected return of the portfolio as a whole. Certainly, this assumption is consistent with the observed decentralization of the professional money management industry.

How is the equilibrium determined in this economy? Who gets money to manage? Well, as investors know

who the skilled managers are, money will flow to the best manager first. Eventually, this manager will receive so much money that it will impact the manager's ability to generate superior returns, and expected return will be driven down to the second-best manager's expected return. At that point, investors will be indifferent between investing with either manager, so funds will flow to both managers until their expected returns are driven down to the third-best manager.

This process will continue until the expected return of investing with any manager is driven down to the expected return investors can expect to receive by investing in a passive strategy of similar riskiness (the benchmark expected return). At this point, investors are indifferent between investing with active managers or just indexing, and an equilibrium is achieved. Notice that in this equilibrium all managers, regardless of their skill level, have the same expected return.

Now consider a more realistic economy in which neither investors nor managers themselves can initially separate good managers from bad. In this case, participants will have to use the information available to them to deduce managerial ability as best they can. Armed with their inferences, investors will choose managers as they did in the previous economy. Investors will rush to invest with the managers perceived as the best, thereby driving down these managers' perceived expected returns. Capital will flow until investors are indifferent between investing with active managers or indexing, so the expected return of investing with any active manager must again equal the expected return of indexing.

The differentiating characteristic between this economy and the first one comes with the passage of time. In the earlier economy, investors know who the best managers are, so they learn nothing from past performance. When investors do not know who the best managers are, they use managers' performance to update their inferences on who the better managers are.

That is, as time passes, differential skill levels among managers become apparent from managers' investment performance. Consequently, investors react by increasing their investment with managers who have done better than expected; funds will flow to these managers as long as investors believe they are capable of producing better-than-expected returns in the future. Funds will stop flowing to these managers when the managers have so much money under management they are no longer expected to produce superior returns. At that point investors are indifferent between investing with these managers and indexing.

Similarly, investors who invested with poorly per-

forming managers will withdraw funds until the amount of capital under management is reduced to the point that investors believe these managers can at least match the benchmark expected return. Once capital flows stop, the expected return from investing with all managers, regardless of their skill level, is the same as the expected return from indexing.

This scenario demonstrates a general insight. In equilibrium, investors who choose to invest with active managers cannot expect to receive positive excess returns (after fees) on a risk-adjusted basis. If investors do expect to make positive excess returns with a particular manager, there would be an excess supply of capital to that manager. Every investor in the economy who held assets of equivalent risk would want to sell those assets and invest with the active manager instead.

Markets can clear only when the expected return to investors in active funds equals the expected return in alternative investment opportunities. So the risk-adjusted expected excess return to investing with a skilled active manager must be zero.

If investors do not benefit from the manager's skill, who does? The answer is the *managers themselves*. By managing a large fund and charging a fee that is proportional to the amount of assets under management, the manager captures all the economic rents generated using his ability. Highly skilled managers will manage larger funds, earn more in fees, and extract more rents.

THEORY

Our thought experiment communicates the main point in Berk and Green [2004]. When capital is supplied competitively by investors, but ability is scarce, only participants with the skill in short supply can earn economic rents. Investors who choose to invest with active managers cannot expect to receive positive excess returns on a risk-adjusted basis. In any economy with rational profit-maximizing investors who compete with each other, all expected risk-adjusted excess returns must be zero, and realized excess returns must be unpredictable.

Let's now consider the five hypotheses laid out in the introduction in light of this theory.

Myth #1: Return Measures Managerial Skill

In what many people term an "efficient market," investors cannot use public information to beat the market because all investors have access to this information.

By using the information, investors compete away any benefit, and thus the expected return of this strategy is simply the market return (or whatever return is commensurate with the risk that is undertaken).

The identical argument applies to portfolio managers. It is no more plausible for investors to expect to earn excess returns picking portfolio managers than it is for them to believe they can earn excess returns picking individual stocks.

If investors find a manager who can consistently beat the market, they will flock to invest with this manager. Eventually, the manager will have so much money under management she will not be able to deliver superior performance. Competition among investors drives the manager's return down to the return investors could earn by themselves.

The result is that all managers (who hold portfolios of the same riskiness) are expected to earn the same return, regardless of their skill level. The immediate implication is that the return managers earn does not measure their skill levels.

Myth #2: The Average Manager Lacks Skill

Because investors must infer the skill level of managers, in some cases they will underestimate it. These managers are likely to have higher realized returns than investors expect. Hence, high realized returns are associated with managers whose skill is underestimated by investors. Similarly, managers whose skill level is overestimated by investors are likely to have low excess returns.

Investors cannot know which managers' skill levels are overestimated and which ones are underestimated. On average, managers must make the market return; otherwise investors could beat the market by simply investing across all managers. This statement is true *regardless* of the average skill level of active managers.


Hence, the empirical finding that, on average, active managers cannot beat passive managers does not imply that the average active manager lacks skill. It just means that the capital markets are competitive.

Myth #3: Skill Implies Persistence

When a manager does well, investors rationally infer that it is likely they have underestimated the manager's skill level. Hence funds flow to the manager. As long as investors believe the manager is capable of delivering superior performance, funds will continue to flow in. The flow of funds will cease only once the manager has

EXHIBIT

Distribution of Management Skill



Distribution of alpha as inferred by Berk and Green [2004]. The vertical line marks the level of the management fee—1.5%. Approximately 80% of the area below the curve lies to the right of this line.

so much money to manage investors believe, given manager skill level, that they are no better off investing with the manager than indexing.

A similar argument applies if a manager underperforms. In this case investors infer that they overestimated the manager's ability and so withdraw funds. In some cases this might mean that the manager will have to shut down the fund, but it is more likely that as investors withdraw funds, manager expected return will increase. If it rises to the expected return of the benchmark, investors will no longer have reason to withdraw funds.

Notice that in both cases the manager is expected to earn the benchmark return. Regardless of prior performance or skill level, all managers' expected returns going forward are the same—the benchmark expected return. Naturally, realized returns will differ across managers, but these differences must be unpredictable, given the information set at that time, or investment would flow to the managers with predictably better returns. Because past returns are in the information set, there can be no persistence in portfolio manager returns.

One can think of this argument as a version of weak-form efficiency applied to portfolio managers rather than individual stocks. Just as risk-adjusted stock returns should be unpredictable, so should portfolio manager returns.

Myth #4: Investors Chase Past Returns

Although there should be no persistence in active manager returns, investors still use past returns as a guide to investing their money. In fact, it is precisely because of

this behavior that there is no persistence in returns. So the evidence that capital flows are driven by past performance is not evidence of investor irrationality.

Investors are not chasing *past* performance; they are chasing *future* performance, and in doing so compete away the opportunity of benefiting from the skill that produced the historic superior performance.

Myth #5: Managers Do Not Have Performance-Based Contracts

In the portfolio management industry, managers are almost always compensated by a fraction of the total assets under management rather than as a function of their performance in excess of their benchmark. Given the emphasis in most industries on performance-based contracts, it seems puzzling that the contract in the mutual fund industry is rarely performance based. Consequently, people have argued for contracts that reward managers on the basis of the excess return they deliver to their investors (see, for example, Ambachtsheer [1994]).

As I have already argued, the return is not a measure of the skill level of the manager, so any contract that pays a manager as a function of the manager's return would *not* reward the more highly skilled managers with higher compensation.

Better managers manage larger funds, *ceteris paribus*, so the size of the fund is a (noisy) measure of managerial skill. A contract that pays a manager on the size of the fund managed therefore rewards more highly skilled managers with more pay. This means that the standard compensation contract that rewards managers according to the size of the fund is already a performance-based contract.

When a manager does well (beats the benchmark), funds flow in, and manager compensation rises because the pay is a fraction of the assets under management. Similarly, when a manager does poorly (underperforms the benchmark), funds flow out, and compensation falls.

HOW SKILLED ARE MANAGERS?

One of the most important research questions in money management is whether most managers have skill and thus add value. Almost all researchers who have investigated this question have used returns to measure skill, and thus all reach the same erroneous conclusion—that managerial skill is rare to non-existent. As we have seen, the fact that the average manager fails to beat a benchmark is simply evidence that capital markets are compet-

itive and cannot be used to infer anything about the skill level of managers.

One might suppose that to measure skill, one simply needs to measure the manager's return before fees. The problem with this reasoning is that the fee the manager charges for services is only a small part of the costs of managing money. To achieve high returns, management must identify undervalued securities and trade to exploit this knowledge without moving the price adversely. In doing so, managers expend resources and pay bid-ask spreads that diminish the return available to pay out to investors. At some point, these costs increase disproportionately, eventually driving the managers' expected returns down to the benchmark. To measure a manager's skill level, one would need to measure return absent these costs, clearly something that we cannot do directly.

In Berk and Green [2004] we use an innovative approach to measure skill indirectly. We build a model of the U.S. mutual fund industry that assumes skill is normally distributed with unknown mean and variance. We then search for values of these (and the other) unobservable parameters so that the output of the model matches the important features of the data: 1) on average the (after-fee) return of active managers matches their benchmark; 2) the five-year survival rate of funds matches the actual survival rates; and 3) the performance/flow relation matches the observed relation for two-year-old funds.

The Exhibit is the output of this process—the inferred normal distribution of managerial skill in the economy. Skill is defined to be the alpha (α) a manager adds before incurring the costs associated with running a large portfolio. One can think of this measure as the value added by the manager if the size of the portfolio is restricted to maximize the return; this is the highest return a manager *could* generate. The vertical line is the management fee used in the parameterization—1.5%.

If we define a manager who adds value as any manager whose alpha exceeds the fee charged, then the fraction of managers who add value is the area to the right of the vertical line. This amounts to 80% of the distribution. That is, when skill is measured correctly, the data are consistent with a conclusion that the vast majority of active managers add value.

Even more surprising is the extent of what the average manager adds. The mean of the distribution in the Exhibit is 6.5%. Given a management fee of 1.5%, this means that the data are consistent with an alpha of 5% for the average manager.

Of course, investors themselves never see this. Competition among them increases the size of the fund and drives the alpha to zero. Instead the managers themselves capture this value through the fees they charge.

CONCLUSION

I have argued that much of what we observe about the behavior of actively managed mutual funds is consistent with a world populated by rational value-maximizing investors who compete with each other. An important insight is that returns cannot be used to measure managerial skill. Because researchers generally use return to measure skill, they have drawn the erroneous conclusion that active managers add little value.

Given the overall levels of manager compensation, one would expect that managers in aggregate should have significant levels of skill and thus add considerable value. I show that when skill is measured correctly, the data are indeed consistent with many skilled managers who add considerable value but capture this themselves in the fees they charge.

ENDNOTE

The author thanks Richard Green, Audrey Grubman, and Jack Treynor for comments and suggestions. This research was supported by a grant from the National Science Foundation.

REFERENCES

- Ambachtsheer, Keith P. "Active Management That Adds Value: Reality or Illusion?" *The Journal of Portfolio Management*, 21 (1994), pp. 89-92.
- Berk, J.B., and R.C. Green. "Mutual Fund Flows and Performance in Rational Markets." *Journal of Political Economy*, 112 (2004), pp. 1269-1295.
- Carhart, M. "On Persistence in Mutual Fund Performance." *Journal of Finance*, 52 (1997), pp. 57-82.
- Gruber, M.J. "Another Puzzle: The Growth in Actively Managed Mutual Funds." *Journal of Finance*, 51 (1996), pp. 783-810.

To order reprints of this article, please contact Ajani Malik at amalik@ijournals.com or 212-224-3205.