MARCH 2006

The Informational Efficiency of Stock Prices: A Review

JAMES L. DAVIS

Introduction

ONE OF THE MOST ENDURING QUESTIONS IN FINANCE is whether securities prices are set in an informationally efficient manner. Dozens of academic studies over the years have reached different conclusions on the issue. Since there is a lot of money at stake, the debate can become quite heated. If securities prices are informationally efficient, then the fees charged by active managers are not justified.

This paper is an attempt to pull together some of the pertinent evidence on the issue. The literature is so broad that it cannot be adequately surveyed in a single article. Excellent surveys already exist in certain areas, and those areas will not be re-examined here. The primary focus of this survey is the performance of professional money managers. This seems to be the most practical aspect of the abstract concept of informational efficiency. Most people are interested in the issue only to the extent that it affects what they do with their own money.

The next section describes the characteristics of an informationally efficient market. Following that is a discussion of the empirical evidence on money manager performance. A discussion of how much information is reflected in prices and a brief conclusion brings things to a close. Throughout the discussion, the emphasis is on the stock market. This is the most thoroughly studied of the securities markets, and it is also the market in which inefficiencies are most likely to be uncovered.

What Are the Characteristics of an Efficient Market?

If securities prices reflect all publicly available information in an unbiased way, what would be the resulting traits of securities markets? Three come to mind:

1. *Return predictability.* The difference between realized returns and expected returns should not be predictable. There are two parts to this—a time series part and a cross-sectional part. The time series part says that investors should not be able to predict which time periods will produce abnormally high or low stock market returns. The cross-sectional part says that investors should not be able to predict which stocks are likely to have abnormally high or low returns during a subsequent period. In both cases, realized returns are compared to an estimate of expected returns to determine what part of the return is abnormal.
2. *Financial market link to the real economy.* Cross-sectional differences in expected return should be related to risk factors that have a meaningful effect in the economy. If a company's operating results are highly sensitive to a particular macroeconomic variable, and if this variable affects large segments of the economy, then the company's stock should earn a risk premium to compensate investors for this non-diversifiable risk. Sensitivity to pervasive risks in the real economy should be reflected in stock prices.[1]
3. *Performance persistence.* Professional money managers should not consistently be able to earn high returns by analyzing financial information. If the information is already reflected in securities prices, gathering and analyzing the available information should not provide a reliable payoff for investors. In any given period, some investors will realize high returns just by chance. But in an efficient market, certain skilled investors should not be able to profit consistently from their analytical abilities.

There are complications associated with all three of these characteristics. When assessing return predictability, it is necessary to specify an expected return for each security and the overall market before one can test for predictability of the abnormal return. This requires a model for expected returns. If a researcher finds evidence of predictability, it could be because the researcher is using the wrong model of expected returns. Fama (1976, 1991) refers to this as the *joint hypothesis* issue; tests of predictability are implicitly testing for both return predictability *and* the validity of a particular model of expected returns. Fama (1991) provides an excellent discussion of the evidence on predictability, and that discussion will not be repeated here (also see Rubinstein, 2001). For investors, the bottom line is that certain variables do appear to have predictive power for returns. However, it is not clear whether these variables are

[1] Note that a risk factor will only be priced if the hedging demands of economic agents do not offset each other. For example, consumers of gasoline and home heating oil are hurt if oil prices rise, while employees of (and investors in) energy firms are hurt if oil prices fall. There would only be a risk premium for unanticipated oil price changes if the corresponding hedging demands do not offset each other.

predicting abnormal returns or rational variation in expected returns. People on both sides of the efficiency debate interpret the results to support their views. Consequently, the empirical studies on predictability do not provide conclusive evidence either for or against the efficient markets hypothesis.

The financial market link to the real economy has also been studied extensively, including a thorough and rigorous discussion by Cochrane (2005). The evidence is mixed, and more work is sure to be done in this area. For an example of a study that finds evidence of a link between the broad market and expected business conditions, see Campbell and Diebold (2005). Xing and Zhang (2005) find evidence that value stocks have higher average returns than growth stocks because their operating performance is more sensitive to economic activity.

The performance persistence issue is the main focus of this paper. The work done in this area provides the most direct evidence on the question of informational efficiency. If prices do not accurately reflect available information, astute investors should be able to earn positive abnormal returns by searching for improperly valued securities. The studies discussed below present evidence that this is difficult to do on a consistent basis.

Empirical Evidence

There are some empirical irregularities that researchers point to when they argue that markets are inefficient. For example, Froot and Dabora (1999) find that Royal Dutch and Shell Transport have frequently not been priced in line with their relative claims on cash flows. In the early 1900s, the two companies merged their interests with an agreement that entitled Royal Dutch and Shell to split the two entities' combined cash flows on a 60/40 basis. This agreement remained in force until the two firms formally merged in 2005. In theory, Royal Dutch should have been consistently priced at 60/40=1.5 times the value of Shell Transport. Froot and Dabora (1999) document frequent and sometimes large deviations from this 60/40 parity relation. This violation of a simple parity relation is sometimes mentioned as evidence of inefficient prices.

What is the practical value of this information for investors? If the parity relation is violated by a wide margin, investors could buy the cheap (relative to parity) stock and short the expensive one. If the pricing discrepancy is big enough to cover the costs of the necessary (round trip) arbitrage trades, and if the activities of the arbitrageurs force prices back to parity, investors could profit. Do exploitable errors happen frequently enough—and are they large enough—for investors to earn abnormal returns?

One way to answer this question is to look at the investment results of professional money managers. If anyone can identify and exploit pricing errors, it should be the professionals. While the documentation of pricing irregularities is interesting, it is the performance studies that address the most practical issue: Can professional money managers use information to earn superior returns?

Dozens of performance studies have been conducted over the years, with widely varying results. Much of this variation in results is due to significant differences in methodology and data quality. To put it bluntly, some of the existing studies just aren't very good. Questionable methods, coupled with biased datasets, make the results of some studies very misleading. In the following discussion, I focus on studies that I believe to be of high quality.

Empirical studies of performance persistence can be differentiated along a few dimensions:

- *Sample size.* How long is the sample period, and how many funds are included? A small number of funds or a short sample period will provide no clue as to whether the results hold generally. Studies with small samples should be interpreted with extreme caution.
- *Benchmarks.* What model is used to identify and measure abnormal performance? A benchmark must control for risk, or else it is impossible to tell whether a fund's high returns are due to manager skill or portfolio risk. Also, a benchmark should be able to measure returns in excess of (or below) a passive strategy that requires no skill.
- *Selection bias.* Most mutual fund databases suffer from survivorship bias; non-surviving funds are absent or under-represented. How is this bias corrected or minimized? Since non-surviving funds tend to be poor performers, excluding them from a dataset will overstate the average fund's performance.

The choice of an appropriate benchmark requires some additional discussion. Currently, the three-factor model of Fama and French (1993) is probably the most widely used benchmark for academic studies of investor performance.[2] Since many researchers believe the three Fama/French factors measure risk, the model is a natural choice. It is still a valid choice even for those who do not believe all three factors to be related to risk. An investor can form a passive portfolio that takes a certain tilt toward small cap and high book-to-market companies. If active managers are able to add value, they should earn average returns higher than a passive strategy with similar size and style tilts. The Fama/French model is a convenient way of controlling for passive exposure to small or large cap companies, as well as to value or growth companies. Managers that truly possess security selection abilities should be able to pick stocks that outperform passive style-motivated strategies.

---

[2] The Fama/French three-factor model takes the form: $R_{it} - R_{ft} = a_i + b_i (R_{mt}-R_{ft}) + s_i SMB_t + h_i HML_t + e_{it}$ In this model, R_{it} is the return to portfolio *i* for month *t*; R_{ft} is the T-Bill return for month *t*; and R_{mt} is the return to the CRSP value-weighted index for month *t*. SMB_t is the realization on a capitalization-based factor portfolio that buys small cap stocks and sells large cap stocks. Similarly, HML_t is the realization on a factor portfolio that buys high-BtM stocks and sells low-BtM stocks. The s_i and h_i coefficients measure the sensitivity of the portfolio's return to the small-minus-big and high-minus-low factors, respectively. Portfolios of value stocks will have a high value for *h*, while growth portfolios will have a negative *h*. Large cap portfolios will load negatively on *SMB* (s_i will be negative), and small cap portfolios will have a positive value for *s*. See Fama and French (1993) for a detailed discussion of this model.

Table 1 provides a summary of how various studies compare along the three dimensions of sample size, benchmarks, and selection bias. This is not an exhaustive list; these studies are chosen to be a representative sample of the work done to date. Table 2 provides a summary of the main results of these studies.

Table 1

Summary Characteristics for Empirical Performance Studies

Study	Sample Period	Total Number of Funds	Benchmarks	Survivorship Bias
Jensen (1969)	1955-1964 1945-1954	115 56	Capital Asset Pricing Model	Not addressed
Grinblatt and Titman (1992)	1975-1984	279	Custom benchmark that controls for size, dividend yield and prior return	Not directly addressed
Lakonishok, Shleifer and Vishny (1992)	1983-1989	769*	S&P 500 Index	Not directly addressed
Hendricks, Patel and Zeckhauser (1993)	1974-1988	165	Various	"Mitigated"
Elton, Gruber, Das and Hlavka (1993)	1965-1984	143	Three-factor model: S&P 500 Index, non-S&P 500 stocks, and bonds	Apparently eliminated
Goetzmann and Ibbotson (1994)	1976-1988	728	Capital Asset Pricing Model	Effects are indirectly estimated but not eliminated
Brown and Goetzmann (1995)	1976-1988	Not reported; at least 1,000	Various	Mostly eliminated
Malkiel (1995)	1971-1991	Not reported; at least 724	S&P 500 Index, Capital Asset Pricing Model	Apparently eliminated for some tests
Elton, Gruber and Blake (1996)	1977-1993	188	Four-factor model: Stock market, size, value-vs.-growth, bond index	Apparently eliminated
Carhart (1997)	1962-1993	1,892	Capital Asset Pricing Model and four-factor model: stock market, size, value-vs.-growth, prior return	Eliminated
Quigley and Sinquefield (2000)	1978-1997	752**	Capital Asset Pricing Model and three-factor model: Stock market, size, value-vs.-growth	Mostly eliminated
Davis (2001)	1962-1998	4,686	Three-factor model: Stock market, size, value-vs.-growth	Eliminated

* Lakonishok, Shleifer and Vishny examined equity pension funds instead of mutual funds.
** Quigley and Sinquefield examined equity unit trusts in the UK.

To a large extent, the mutual fund performance literature can be divided into two parts—the studies done before Carhart's (1997) path-breaking work, and those done afterward. As Table 1 shows, Carhart's database reached an entirely different level in terms of sample size, when compared to earlier studies. Both the number of funds and

Table 2

Summary Characteristics for Empirical Performance Studies

Study	Main Results
Jensen (1969)	Funds don't cover costs, on average. There is very little evidence that any individual fund has been able to earn significantly higher returns than those that could have been expected by random chance.
Grinblatt and Titman (1992)	Five-year returns exhibit persistence. Losers show stronger persistence.
Lakonishok, Shleifer and Vishny (1992)	Before fees, pension funds underperform the S&P 500 Index by 1.3% per year on an equally weighted basis, and by 2.6% per year on a value-weighted basis. There is very little evidence of persistence in annual returns, but some persistence is seen at two-year and three-year horizons.
Hendricks, Patel and Zeckhauser (1993)	Last year's winners continue to perform well for one to eight more quarters. Last year's losers continue to perform poorly. Persistence in poor performance is stronger than persistence in good performance.
Elton, Gruber, Das and Hlavka (1993)	Funds underperform passive benchmarks on average. Funds with high fees and turnover underperform funds with low fees and turnover.
Goetzmann and Ibbotson (1994)	Both winners and losers repeat over horizons from one month to three years.
Brown and Goetzmann (1995)	The level of observed persistence depends on the time period under study, and is mostly due to poor performers.
Malkiel (1995)	Most funds underperform both the S&P 500 Index and the CAPM. There is evidence of performance persistence, but it is due entirely to the 1970s.
Elton, Gruber and Blake (1996)	Both one-year and three-year returns exhibit persistence, for both winners and losers. Strong persistence among losers is driven largely by high expenses.
Carhart (1997)	There is some evidence of persistence among winners, but this appears to be due to momentum in their underlying stocks instead of skill. Losers exhibit strong persistence, and high expenses account for much of this persistent poor performance.
Quigley and Sinquefield (2000)	UK equity unit trusts underperform the Fama/French three-factor benchmark, on average. Unit trusts that are focused primarily on small companies have especially poor performance. Bad performance persists, but good performance does not.
Davis (2001)	Some short-run persistence is observed among the best-performing growth funds and the worst-performing small cap funds.

the length of the sample period distinguish Carhart's work from earlier attempts. He also sets his research apart by explicitly addressing survivorship bias and distinguishing skill from simple momentum effects.

Several of the earlier studies found evidence of persistence among winners; funds with abnormally high returns in one period tend to have high returns in subsequent periods. This result is frequently cited as evidence that some managers are more skillful than others. Carhart shows that this short-run persistence is driven primarily by momentum in the funds' underlying stocks. It is not the case that managers effectively exploit momentum to enhance fund returns. Rather, Carhart's results show that some managers just happen to own stocks that develop upward momentum. When the momentum

in those stocks dissipates, so does the abnormal performance of the manager. The practical implication of Carhart's results is that investors cannot identify superior managers by looking at past returns.

Some studies provide evidence that not all managers possess the same abilities. For example, Chevalier and Ellison (1999) compare the characteristics of fund managers to fund performance to see if certain manager characteristics are associated with good returns. The variable that had the most predictive ability was the average SAT score of the manager's undergraduate institution; managers from schools with high SAT scores have high average returns. This could be interpreted as evidence of better skill or better information. However, the cross-sectional variation in fund returns is so large that even a manager from a school with an average SAT score[3] of 1420 cannot be expected to outperform a risk-adjusted benchmark at conventional significance levels. Chevalier and Ellison's manager characteristics model can explain only about 5% of the total variation in mutual fund returns (see their Table IV).

Wermers (2000) finds that mutual funds underperform the CRSP value-weighted market index by 1% per year for the 1975-1994 period, even though the stocks they own *outperform* the CRSP index by 1.3% per year. Why do mutual fund returns lag the returns of their underlying stocks by 2.3%? Wermers says that 1.6% of this 2.3% difference is due to fund expenses and transactions costs, and the remainder is due to fund holdings of securities other than common stock. In an effort to understand why the fund managers' stock picks have higher returns than the CRSP index, Wermers uses a characteristics-based model and finds that only part of the higher return can be explained by the characteristics of the stocks (size, book-to-market, and momentum). He says the part that can't be explained by stock characteristics is due to manager skill.

There may be some value added by the security selection efforts of some managers, but this added value does not cover the costs generated by the managers. The 1.6% in expenses and transactions costs mentioned above is larger than any reasonable estimate of value added through security selection. A regression of net fund returns on the four factors used by Carhart yields an annual abnormal return of -1.16%. Managers do not cover their costs on a consistent basis.

Some investors argue that small cap stocks are priced in a less efficient manner than large cap stocks, so that small cap pricing errors can more readily be exploited. This implies that there should be evidence of positive persistence among the small cap portfolio managers who are the best at exploiting valuation errors. However, the results of Davis (2001) and Quigley and Sinquefield (2000) reveal no evidence of reliable positive persistence among small cap managers. Small cap stocks do not appear to offer greater promise of positive abnormal returns.

---

[3] This test score is based on the old scoring system that was in place before the writing component was added to the SAT I exam in 2005. Under the old system, the maximum possible score was 1600.

The main message of the performance studies seems to be that it is very difficult to consistently earn positive abnormal returns by gathering and analyzing information. The average mutual fund underperforms a style-adjusted passive benchmark, and there is very little evidence of persistence in good performance when returns are adjusted for size and style tilts and simple momentum effects. Some managers may be more skilled than others, but the payoff to this skill does not appear to be large enough to cover the costs of active management.

How Much Information Should Be Reflected in Prices?

Grossman and Stiglitz (1980) present an insightful model of how information is reflected in securities prices. Information is costly to gather and analyze. Recognizing this, investors will only expend resources on information if they expect the resulting payoff to cover their costs. If the cost of information is high, it will take a large expected payoff to induce investors to gather and analyze information. In this model, the degree of informational efficiency of prices depends on the cost of information. An implication of this model is that investors should cover their costs, on average. They would collect and analyze information up to the point where the marginal benefit of this activity equals the marginal cost, but no more. Viewed in the light of this model's predictions, the evidence that money managers do not cover their costs is a bit of a puzzle. It appears that these managers engage in more information collecting than they should. They continue to gather information past the point where the expected marginal payoff equals the marginal cost. These results suggest that securities markets are actually more efficient than the Grossman and Stiglitz model would predict.

Society benefits from efficient securities markets, since an efficient market allocates investment capital to productive uses. The observation that fund managers do not cover their costs suggests that fund shareholders are subsidizing a public benefit. The activities of managers make markets more efficient, but the costs of these activities are not fully compensated through higher investment returns. Society is better off, but fund shareholders are paying the bill for this benefit. This appears to be an unintended benevolence.

Conclusions

Do pricing errors occur in financial markets? By now it is clear that they do. It is also clear that investors have a difficult time consistently exploiting these errors. The returns earned by professional investors look like what one would expect from a market that rationally reflects costly information in securities prices. Fama (1991, p. 1575) reminds us that "the extreme version of the market efficiency hypothesis is surely false." However, in light of the performance of money managers over the past several decades, it appears that "rational pricing with costly information" is a reasonable approximation of reality.

The helpful comments of Ken French, David New, Truman Clark, and Weston Wellington are gratefully acknowledged.

References

Brown, Stephen J., and William N. Goetzmann. 1995. Performance persistence. *Journal of Finance* 50:679-698.

Campbell, Sean D., and Francis X. Diebold. 2005. Stock returns and expected business conditions: Half a century of direct evidence. PIER Working Paper No. 05-025, University of Pennsylvania; CFS Working Paper No. 2005/22, Universität Frankfurt am Main.

Carhart, Mark M. 1997. On persistence in mutual fund performance. *Journal of Finance* 52:57-82.

Chevalier, Judith, and Glenn Ellison. 1999. Are some mutual fund managers better than others? Cross-sectional patterns in behavior and performance. *Journal of Finance* 54:875-899.

Cochrane, John H. 2005. Financial markets and the real economy. NBER Working Paper 11193, National Bureau of Economic Research.

Davis, James L. 2001. Mutual fund performance and manager style. *Financial Analysts Journal* 57:19-27.

Elton, Edwin J., Martin J. Gruber, and Christopher R. Blake. 1996. The persistence of risk-adjusted mutual fund performance. *Journal of Business* 69:133-157.

Elton, Edwin J., Martin J. Gruber, Sanjiv Das, and Matthew Hlavka. 1993. Efficiency with costly information: A reinterpretation of evidence from managed portfolios. *Review of Financial Studies* 6:1-22.

Fama, Eugene F. 1976. *Foundations of Finance*. New York: Basic Books.

Fama, Eugene F. 1991. Efficient capital markets: II. *Journal of Finance* 46:1575-1617.

Fama, Eugene F., and Kenneth R. French. 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33:3-56.

Froot, Kenneth A., and Emil M. Dabora. 1999. How are stock prices affected by the location of trade? *Journal of Financial Economics* 53:189-216.

Goetzmann, William N., and Roger G. Ibbotson. 1994. Do winners repeat? *Journal of Portfolio Management* 20:9-18.

Grinblatt, Mark, and Sheridan Titman. 1992. The persistence of mutual fund performance. *Journal of Finance* 47:1977-1984.

Grossman, Sanford J., and Joseph E. Stiglitz. 1980. On the impossibility of informationally efficient markets. *American Economic Review* 70:393-408.

Hendricks, Darryll, Jayendu Patel, and Richard Zeckhauser. 1993. Hot hands in mutual funds: Short-run persistence of relative performance, 1974-88. *Journal of Finance* 48:93-130.

Jensen, Michael C. 1969. Risk, the pricing of capital assets, and the evaluation of investment portfolios. *Journal of Business* 42:167-247.

Lakonishok, Josef, Andrei Shleifer, and Robert W. Vishny. 1992. The structure and performance of the money management industry. Working paper, University of Illinois at Urbana-Champaign, Harvard University, and the University of Chicago.

Malkiel, Burton G. 1995. Returns from investing in equity mutual funds 1971 to 1991. *Journal of Finance* 50:549-572.

Quigley, Garrett, and Rex A. Sinquefield. 2000. Performance of UK equity unit trusts. *Journal of Asset Management* 1:72-92.

Rubinstein, Mark. 2001. Rational markets: Yes or no? The affirmative case. *Financial Analysts Journal* 57:15-29.

Wermers, Russ. 2000. Mutual fund performance: An empirical decomposition into stock-picking talent, style, transactions costs, and expenses. *Journal of Finance* 55:1655-1695.

Xing, Yuhang, and Lu Zhang. 2005. Value versus growth: Movements in economic fundamentals. Working paper, Rice University and the University of Rochester.