A Critique of the Efficient Market Hypothesis*
*** Preliminary and Incomplete ***

Jonathan B. Berk
University of California, Berkeley, and NBER

First Draft: October, 2007
This Draft: January 22, 2008, 4:23pm

ABSTRACT

The Efficient Market Hypothesis (EMH) has outlived it usefulness in financial economics. I show that most forms of the hypothesis are untestable. The forms that are testable are easily rejected in the data. The idea that returns are unpredictable is misguided and impedes research in financial economics. Returns are only unpredictable under the risk neutral (martingale) probability measure. Under the actual probability measure, returns are predictable, and indeed, it is the *predictable* part of return that interests financial economists. It is time for the profession to recognize this reality dispense with the EMH.

JEL classification: G14.

*I am grateful for

> A finance professor and a student are walking down a street and the student sees a $100 bill lieing on the pavement. As the student is reaches down to pick up the money the professor says "Don't bother; if that really was a $100 bill it would already be picked up." — an unattributable old joke.

I usually tell this joke in my introductory finance classes and in my experience it is good for a laugh. However, once the laughter has died down, I add my own addendum: I ask the students if anybody has actually found $100 lying in the street. Invariably, in a class of 100 or so students, one or two students have, and so they raise their hands. And that is the real point of the joke. It is not that $100 bills are impossible to find in the street. Rather, because $100 is a significant amount of money it is unlikely that somebody would accidentally drop a bill of that size and even if they did, one would have to be extremely lucky to be the first one to spot it. This intution, that because financial markets are very competitive, money making opportunities should be rare, is intuition that underlies what Eugene F. Fama termed the *Efficient Market Hypothesis* (EMH).

The Efficient Market Hypothesis has arguably become the most influential concept in financial economics. There is no doubt that the insight, or perhaps better stated as the intuition, underlying the EMH was path breaking: Competition between financial market participants should move prices so that they are not predictable. By forcing economists to explicitly recognize the effect of competition in financial markets, Fama brought a discipline to finance that has shaped research ever since.

Yet, as important as this breakthough was, I will argue in this paper that the EMH has outlived its usefulness. I will show that the hypothesis is either untestable or rejected in the data (based on which definition of the EMH is adopted). The version of the hypothesis that postulates that returns should be unpredictable is wrong — returns should be and are predictable. Although it is true that most of the variation in returns is unpredictable, it is the *predictable* part of return that researchers and practitioners alike are interested in. This predictability could potentially derive from rational or behavioral model; the predictability itself is not evidence either for or against the rationality of markets.

Ironically, the EMH is a victim of its own success; the subsequent research in financial economics largely spurred by the EMH has advanced the field to the point that now the EMH has become a barrier to advancement. Continuing the above metaphor, in the last half century a huge amount of research resources have been devoted to showing how hard $100 bills are to find, while very few resources been devoted to studying the people lucky enough to find $100 bills. A priori, the question of how hard profit making opportunities are to find in financial markets would appear just as interesting as question of what they look like and how individuals find them. Yet, academic research on the latter topic is practically

non-existent.[1]

Historically, belief in the validity of the EMH is the price every financial economist has had to pay to publish research in finance.[2] Not surprisingly, therefore, papers that start with the assumption that profit making opportunities exist and are exploitable are difficult to publish, explaining the dirth of papers studying what they look like and how people find these opportunities. A second, and no less important, reason research on profit making opportunities in financial markets is scant is that these opportunities can only be available to a few market participants (otherwise market participants would compete the opportunities away). Such monopolistic power can derive from two sources: (1) informational advantages or (2) information processing advantages. Because of the existence and enforcement of strict insider trading rules, it is hard to argue that there are large differences in participants information sets. Nevertheless, the profession has almost exclusively concentrated on the first reason almost certainly because the second is ruled out by the EMH. The most general statement of the EMH as defined in Fama (1970) is the rational expectations hypothesis, enunciated first by Muth (1961). If all market participants have access to the same information set, then the rational expectations assumption requires that they all agree on the distribution of market returns. Thus the EMH effectively rules out differences in expected returns that result from differences in ability to process information.

By ignoring differences in ability, the profession has needlessly restricted the range of explanations for the behavior we observe. For example, researchers have puzzled why trading volume is so high and prices are so variable. By explicitly relaxing the EMH assumption by differences in ability of market participants to process information (which necessarily implies that some participants will not correctly compute expected returns), progress would likely be made towards resolving these empirical puzzles.

This paper is organized as follows. In the next section I will briefly review the literature that defined the EMH. The following section then reviews the evidence and concludes that some versions of the EMH are untestable and the versions that are testable are easily rejected in the data. As Fama himself points out (Fama 1970, p. 391) stating the EMH in terms of a testable hypothesis has proved difficult. I argue in Section 2 that, at least as far as predictability in returns is concerned, this task is not just difficult, it is impossible. Section 3 then makes a case that, 40 years after Fama's seminal paper, it is time to dispensing with the EMH because it stifles potentially interesting research in financial economics.

[1] Reference some of the few articles.

[2] Even contributions in the now voluminous behavioral finance literature rarely take aim at the EMH directly. More often these papers limit themselves to documenting apparently surprising empirical regularities or behavioral biases.

1 The EMH Defined

As Fama concedes in Fama (1970), the notion of market efficiency put forward in Fama (1965) was not well defined. As he states in Fama (1970, p.384):

> "The definitional statement that in an efficient market prices 'fully reflect' available information is so general that it has no empirically testable implications. To make the model testable, the process of price formation must be specified in more detail. In essence we must define somewhat more exactly what is meant by the term 'fully reflect.'"

Fama then goes on to define the EMH as the statement that if

$$\tilde{\epsilon}_{j,t+1} \equiv \tilde{r}_{j,t+1} - E[r_{j,t+1}|\phi_t] \tag{1}$$

then

$$E[\tilde{\epsilon}_{j,t+1}|\phi_t] = 0 \tag{2}$$

where ϕ_t is the information set at time t, $E[r_{j,t+1}|\phi_t]$ is the expected return of security j conditional on the information set at time t, $\tilde{r}_{j,t+1}$ is the realized return of security j at time $t+1$. As Fama points out, specified this way, the EMH is untestable because $E[r_{j,t+1}|\phi_t]$ is unobservable. To make the hypothesis testable further assumptions must be made. In Fama (1970) the working assumption is that expected returns are constant, which implies that the unconditional autocovariance of returns is zero, $\text{cov}(\tilde{r}_{j,t+1}, \tilde{r}_{j,t}) = 0$. The paper then tests whether this condition, which Fama term's a "fair game", can be rejected by the data. He concludes that it cannot. Subsequent to that paper the test of whether, in Fama's words, returns were a "fair game" became known as the test of "weak form" efficiency.

By 1988 it had become clear that the conclusion in Fama (1970) was premature and that indeed the hypothesis that returns are serially uncorrelated can be resoundingly rejected in the data.[3] However, rather than take this evidence as evidence against the EMH, the profession's reaction was that it merely showed the working assumption that expected returns are constant can be rejected. In the words of Lo and MacKinlay:

> "[Our] results do not necessarily imply that the stock market is inefficient or that prices are rational assessments of 'fundamental' values. As LeRoy (1973) and Lucas (1978) have shown, rational expectations equilibrium prices need not even form a martingal sequence, of which the random walk is a special case. Therefore,

[3] see, for example, DeBondt and Thaler (1985), De Bondt and Thaler (1987), Lo and MacKinlay (1988) and Conrad, Jennifer and Kaul, Gautam (1988)

without a more explicit economic model of the price-generating mechanism, a rejection of the random walk hypothesis has few implications for the efficiency of market price formation."

As the above quote makes clear, the EMH is not a well defined concept. The term is used by different people to mean different things. Many researchers define the EMH, as Ross (2004) does, as the statement that a positive pricing operator exists, given a particular information set. But the existence of this operator is equivalent to the absence of arbitrage opportunities — it does not restrict individual expectations to be rational, which we will see is the essence of Fama's original idea. This is a rather sad state of affairs because it is difficult to test concepts that are not well defined. To simplify matters in this paper, I will use to the definitions put forward by Fama himself in Fama (1991). He has three definitions in that paper, although I will lump two of the definitions together because it is unclear to me that an important distinction exists between them.

Fama's first definition, loosely stated, is that returns should be unpredictable given an information set. As we shall see in the next section, writing this intuition in terms as a testable hypothesis has proved difficult. The second definition, again loosely stated, is that, given an information set, money making trading strategies do not exist. Although tests of this latter definition are far less common in the literature, the notion that nobody can make money in financial markets is deeply ingrained in the profession. In the next section examines whether either definition is a testable hypothesis.

2 Is the EMH Testable?

One of the reasons the for the longevity of the EMH is surely that it is supposedly a testable hypothesis. In this section I will show that whether the EMH is testable depends on which definition is used. Ironically, the form of the EMH that is most often "tested" is actually untestable.

2.1 Tests for Return Predictability

Once it became clear that there is strong evidence against the hypothesis that returns are a "fair game" Fama's definition of weak form efficiency changed. Pointing out that it never was the statement that the autocovariance in returns was zero, in Fama (1991, p. 1575) Fama returns to his original definition: "security prices reflect all available information." However, stated this way, the hypothesis is untestable. Using the arguments in LeRoy (1976), let's see

why. Returning to (1) and taking expectations gives:

$$E[\tilde{\epsilon}_{j,t+1}|\phi_t] \equiv E[\tilde{r}_{j,t+1}|\phi_t] - E[r_{j,t+1}|\phi_t] = 0. \tag{3}$$

That is, (1) follows directly from (2) by the properties of the expectation operator, it requires no auxiliary assumptions, certainly not the EMH. To make sense, then, of Fama (1970) one has to take the position that what is being tested is whether investor expectations are correct, that is whether[4]

$$E[\tilde{r}_{j,t+1}|\phi_t] = E[r_{j,t+1}|\phi_t]. \tag{4}$$

The problem with testing (4) is that investor expectations are unobservable and because only people can have expectations, it makes no sense to talk about the "market's expectations." So an alternative approach is to ask the question how should (or do) investors *form* their expectations? The answer to this question is what we call an "asset pricing model" — a model for how investors should (or do) compute the expected return of an asset.

Consider how we test asset pricing models. Specifically, let $\bar{R}_{j,t+1}(\phi_t)$ be the expected return of asset j specified by the asset pricing model under consideration conditional on all the information available at time t. If the asset pricing model holds exactly then by the properties of the expectation operator, actual returns must have the following property:

$$\tilde{r}_{j,t+1} = \bar{R}_{j,t+1}(\phi_t) + \tilde{\epsilon}_{j,t+1} \tag{5}$$

with

$$E[\tilde{\epsilon}_{j,t+1}|\phi_t] = 0. \tag{6}$$

Put another way, if there is any predictability in $\tilde{\epsilon}_{j,t+1}$, $\bar{R}_{j,t+1}(\phi_t)$ must be mispecified, which implies the asset pricing model does not hold exactly. Thus a test of whether (6) holds is solely a test of whether the asset pricing model holds. At no point do we need the EMH to infer (6); thus a test of whether this condition holds is just a statement of whether the asset pricing model correctly calculates the expected return. It tells us nothing about whether the EMH holds — that is whether investors themselves correctly calculate expected returns.[5] For example, assume that returns are perfectly predictable in that they equal the negative of last period's return. Then $\bar{R}_{j,t+1}(\phi_t) = -\tilde{r}_{j,t}$ and $\epsilon_{j,t+1} = E[\tilde{\epsilon}_{j,t+1}|\phi_t] = 0$, so (6) holds but

[4]The insight that fully rational investors must have correct expectations is today known as the rational expectations hypothesis. It was first introduced by Muth (1961) which, inexplicably, is not cited in Fama (1965) or Fama (1970).

[5]Fama (1991, p. 1575) (as well as the rest of the profession) appear to recognize this problem, labeling it the "joint hypothesis" problem: "[Market efficiency] must be tested jointly with some model of equilibrium, an asset-pricing model." What is not clear is what is meant by "jointly."

nobody would call such a market efficient.

This point is worth belaboring. The insight here is that when the asset pricing model holds exactly, $\bar{R}_{j,t+1}(\phi_t)$ is how investors *should* compute the expected return, but, in principal, there is no requirement that they actually *do* compute expected returns that way. For instance, one could have an equilibrium in which investors do not have rational expectations. A correctly specified asset pricing model would still provide $\bar{R}_{j,t+1}(\phi_t)$ exactly, and hence (6) holds, but because investors do not have rational expectations in this equilibrium, the EMH fails. So merely observing that (6) holds does not imply that the EMH holds.[6]

In summary, it is possible for (1) investors not to have rational expectations even though an asset pricing model correctly specifies the expected return and (2) investors to have rational expectations even though an asset pricing model that correctly specifies expected returns remains to be discovered. Consequently, the only way to test this form of the EMH is to directly observe investor expections. But investor expectations are not observable so the EMH untestable.

2.2 Tests of Whether Investors Can Make Money

The second definition of the EMH is the notion of whether investors can make money in markets. At first glance it might appear that testing whether investors can make money in markets is doomed by the same arguments as testing for predictability in returns. This is largely true for most of the traditional tests of this form of the EMH, but as well see below, not all.

Many tests of of this form of the EMH have concentrate on showing that particular trading strategies can or cannot make money. But because the return to any trading strategy might in fact be a reward for bearing risk, without a model of risk and return (an asset pricing model) it is difficult to say definitively whether an investor can make money from a particular trading strategy.[7] However, an even more fundamental problem exists. Fama (1965, p. 34) writes:

> "The main conclusion will be that the data seem to present consistent and strong support for the [random-walk] model. This implies, of course, that chart reading, though perhaps an interesting pastime, is of no real value to the stock market investor. This is an extreme statement and the chart reader is certainly free to

[6] A similar argument can be made in terms of the pricing operator — a positive pricing operator merely guarantees that arbitrage opportunities do not exist. It does not imply that investors have rational expectations.

[7] Recently, Ross (2004) has argued that it might be possible to identify money making trading strategies without knowing the relationship between risk and return using variance bound tests.

> take exception. We suggest, however, that this and other studies in support of the random-walk model is now so voluminous, the counterarguments of the chart reader will be completely lacking in force if the are not equally well supported by empirical work."

This view, that I believe sums up the view of most of the profession, assumes that if a trading strategy that made money exists, it would be possible to identify it. But the unfortunate fact is that anybody who found such a trading strategy would never reveal this information because by revealing it, the opportunity would disappear. That is, competition between investors drives out profitable trading strategies so the only way a profitable trading strategy can exist is if it remains a secret and thus unobservable. Thus so called "strong form" efficiency, that no profit making trading strategies exist, cannot be rejected by finding profitable trading strategies. Even if markets are not strong form efficient, it will be impossible to find such strategies.

One might argue that although it is impossible to observe a money making trading strategy while the strategy makes money, one should be able to observe whether the strategy made money *ex post*: Once enough people discover the strategy and drive profits to zero, the existence of the strategy will become public knowledge. One could then see whether indeed the strategy made money using historical prices. The problem with this approach is that mining data for profit making strategies is always possible. Hence it is not clear how one would differentiate a money making strategy that truly did exist at a profit opportunity in the past from a strategy that historically yielded above normal profits by pure chance. In both cases, the data would show that it was possible to use strategies to make money in the past, but going forward, neither strategy would be expected to keep make money.

The only way to differentiate true money making strategies from artifacts of the data is by demonstrating that investors knew about the strategy *ex-ante* and were able to take advantage of this knowledge by actually making money from the strategy. That is, instead of trying to observe the strategies themselves, one can test this version of the EMH by simply testing whether the people employing the strategies are making money. Indeed, this approach to testing this form of the EMH was pioneered in Fama (1965). That paper examines the return of active mutual fund managers and reaches the conclusion that because active manager returns are no higher than passive manager's, they cannot be employing money making strategies. In addition, Fama recognizes that some managers did do well in certain years, but attributes this performance to luck because performance is inconsistent: Funds that do well in any given year are no more likely to do well in subsequent years. Ever since that study, the idea that money making strategies are so rare that even professionals cannot consistently find them has become deeply ingrained in the profession and is widely

cited as the strongest piece of evidence in favor of the EMH.

It idea that because active portfolio managers as a group earn essentially the same return as passive strategies implies that active managers lack skill and thus money making strategies do not exist is perhaps one of the biggest misconceptions in financial economics. As Berk and Green (2004) show, the average return to investors in a mutual fund cannot be used as a measure of managerial skill. If investors could earn abnormal returns simply by picking particular mutual funds, there would be excess demand to invest in these funds and so the market would not clear. An equilibrium can only be sustained if investors cannot earn abnormal returns investing in funds, that is, at all times the expected excess return of mutual funds must be zero. That means that, on average, active and passive managers should earn the same return and there should be no predictability in performance of any fund, regardless of the skill level of managers. So the fact that active managers do not consistently make higher returns than passive managers is not evidence that active managers are unskilled, and, more importantly, that money making strategies do not exist.

One might conjecture that the return a mutual fund manager makes before fees is a measure of the manager's skill level. But this inference ignores an insight in Berk and Green (2004) — that the percentage fee charged is irrelevant. In that model a manager can either charge a small percentage fee and manage a large portfolio or a large percentage fee and manage a large portfolio. What determines how much the manager makes is his ability not the percentage fee he charges. The form of the contract is irrelevant so long as the contract he uses extracts all the economic rents he generates. In particular he could charge very small percentage fee. In this case is pre fee and post fee return would be almost the same, and the difference would have nothing to say about his ability to find money making strategies.

If the return of mutual fund managers (either pre or post fee) cannot be used as a measure of their skill level, what can? As Berk and Green (2004) argue, the correct measure of the manager's skill level is the amount of money the manager himself makes, that is, the total amount of money charged in fees. Only people with a skill in short supply earn economic rents. In this case the skill is the ability to find money making strategies and so the economic rents they make (as proxied by the amount of money they make) is an indication of how successful they are. Ironically, the evidence on this is unequivocal. Mutual fund managers are amongst the highest paid members of society.[8] Based on the size of the economic rents

[8]The May 2006 Occupational Employment and Wage Estimate released by the Department of Labor lists "Financial Managers" as the 31st highest paid occupation in the U.S. (out of a total of 821 occupations) with an average annual compensation of $101,450 (the mean wage over all occupations was $39,190). Mutual fund managers are most likely amongst the higher earners in this class. For example, the 2005 Investment Management Compensation Survey by Russell Reynolds Associates, New York, N.Y., and CFA Institute, Charlottesville, Va. lists the average annual compensation of mutual fund managers as $390,000, suggesting that if the occupation was more narrowly defined as just encompassing mutual fund managers, it would rank

mutual fund managers are able to extract, one would have to conclude that money making strategies must exist. Thus, rather than provide strong evidence in favor of the EMH, the mutual fund industry in fact provides strong evidence against the EMH.[9]

3 Implications

The unfortunate conclusion is that the depending on what version of the EMH you use, it is either untestable or rejected in the data: Returns are not a random walk and the performance of mutual fund managers indicates that profit making trading strategies exist and some people are able to take advantage of them. Some might argue, as Fama (1991, p.1576) has, that the EMH still has an important role in finance because it has changed the profession's views about the cross-sectional and time series behavior of returns. In doing so it has changed the practice of finance, and the empirical literature it stimulated "is among the most successful in empirical economics, with good prospects to remain so in the future." Fama (1991, p.1576)

Although the EMH certainly has generated a huge empirical literature, it is hard to understand how a concept that is untestable could possibly be useful empirically. Surely the importance of the concept is *theoretical*. I think it is undeniable that in this respect that EMH, or more generally, the rational expectations hypothesis, has played a central and critically important role in the development of modern financial economics. But even here, I think the EMH might have outlived its usefulness.

The EMH was developed at a time when people thought that returns should be predictable based on the passed performance of the company. Their mistaken intuition was much like the mistaken intuition for why people believed that skilled actively managers should consistently beat the market. Good stocks should remain good: hence stocks with high past returns should continue to produce high returns. Fama developed the EMH to counter this appealing but incorrect intuition by pointing out that competition between investors determined stock returns, not the quality of the company. Putting this intuition another way, the only way markets can clear is if all stocks with the same riskiness earn the same expected return. Unfortunately, that is not how the EMH states the intuition. The EMH states the intuition in terms of predictability. The essential take away from the EMH is that, except for a few minor details, stock returns should not be predictable. Unfortunately that take away is destructive because the opposite is true — under the neo-classical

even higher.

[9]Berk and Green (2004) estimate that 80% of managers have skill and that on average these strategies can outperform by as much as 6.5% per annum.

paradigm, stock returns *should be* predictable.

Under the neo-classical paradigm, if one assumes that there is some persistence in the riskiness of companies, then because the expected return of a company is a function only of the company's riskiness, there must be persistence in expected returns. And if expected returns are persistent, then realized returns must be persistent as well. That is, tomorrow's return must be predictable based on today's information (which includes past returns). One might counter by pointing out that the EMH refers only to the part of return orthogonal to the expected return, but as we have already seen, the unpredictability of that part of return is tautological — it derives solely from the properties of the expectation operator. The central idea behind the EMH, that returns in some sense should be unpredictable, is wrong. Religious adherence to this incorrect concept is a barrier to productive research.

Given that for the last forty years the EMH has been taken as a cornerstone of modern finance, it might be difficult to accept that it has in fact inhibited advancement in the field. Consequently, let me provide two specific examples of what I mean.

3.1 Momentum

The empirical observation that winners (stocks with high past returns) earn higher returns than losers (stocks with low past returns) is known as momentum and is widely regarded as an anomaly, that is, evidence against the hypothesis that asset prices are determined by rational investors. In fact momentum, by itself, has *nothing* to say about rationality. Regardless of what process determines expected returns the only condition required to generate momentum profits is persistence. But persistence in returns is unrelated to whether returns reflect risk rather than a behavioral effect. So observing momentum says nothing about the rationality of asset prices. Let me be more explicit about this argument.

Returning again to (1) and rearranging terms:

$$\tilde{r}_{j,t+1} = E[r_{j,t+1}|\phi_t] + \epsilon_{j,t+1}. \tag{7}$$

As we have already pointed out, by the property of the expectation operator we must have

$$E[\epsilon_{j,t+1}|\phi_t] = 0. \tag{8}$$

Now if the average return of all stocks at time t is labelled $\bar{r}_t$, then the unconditional

autocovariance is

$$\begin{aligned} E[(\tilde{r}_{j,t+1} - \bar{r}_{t+1})(\tilde{r}_{j,t} - \bar{r}_t)] &= E[E[r_{j,t+1} - \bar{r}_{t+1}|\phi_t]E[r_{j,t} - \bar{r}_t|\phi_{t-1}]] + E[E[r_{j,t+1} - \bar{r}_{t+1}|\phi_t]\epsilon_{j,t}]] \\ &\quad + E[\epsilon_{j,t+1}E[r_{j,t} - \bar{r}_t|\phi_{t-1}]|\phi_t] + E[\epsilon_{j,t+1}\epsilon_{j,t}|\phi_t] \\ &= E[E[r_{j,t+1} - \bar{r}_{t+1}|\phi_t]E[r_{j,t} - \bar{r}_t|\phi_{t-1}]] + E[E[r_{j,t+1} - \bar{r}_{t+1}|\phi_t]\epsilon_{j,t}]]. \end{aligned}$$

Momentum profits are positive whenever the above expression is positive. The sign of this expression depends on the dynamics of expected returns. For example, if stocks have constant, but different expected returns, then the second term is zero and the first term is positive, implying momentum profits. Constant expected returns can derive from either a rational economy (stock's have different risks) or a irrational economy (high profitability stocks earn high returns). Stocks need not have constant expected returns, however. In Berk, Green, and Naik (1999) at short horizons (time periods) the second term is negative and dominant, so that contrarian profits exist. At longer horizons the second term is not important, but there is persistence in expected returns across stocks. Thus the first term is positive and hence momentum profits exist. Finally, consider a behavior model in which risk has nothing to do with expected returns, indeed, assume that expected returns are random and i.i.d. Then both terms are zero and there is no momentum in the economy.[10] I know of no researcher who has argued that the absence of momentum profits is evidence of *ir*rationality in asset prices, but it seems to me that it is hard to conceive of a realistic model linking risk and return that does not admit profits based on portfolios formed on past returns. The fact that the absence of momentum profits is universally accepted as evidence in favor of the rational paradigm it testament to the damage the EMH has done to the profession's understanding of the behavior of asset returns.

3.2 Ability

Another reason why the EMH is now counter productive is that it rules out the possibility that market participants can differ in ability. Although in many situations in economics, such an assumption makes sense, it is hard to justify in a model of financial markets that potentially includes millions of participants. However, because of the dominance of the EMH in financial economics, all research that has focused on differences in the market participants

[10] Although the existence of momentum profits alone provides no insight into the rationality of markets, researchers have argued that other characteristics of momentum strategies provide insight on this question. A recent paper Conrad, Kaul, and Lei (2007) argue that it current evidence is inconsistent with the hypothesis that momentum is driven by firm specific characteristics.

has almost exclusively restricted attention to differences in information.

One might suppose that the distinction between differences in information and ability is largely semantic — that any model in which participants have differences in ability is isomorphic to a model in which all participants have the same ability but have different information sets. However, I do not share this view, or perhaps more accurately, I do not know of a convincing argument that justifies this view. Traditionally, models with asymmetric information have assumed that participants all have rational expectations. A model in which participants have symmetric information but have differences in ability (in say, computing expected return from the information) cannot have all participants have rational expectations. Because some participants have better abilities than others, participants must have different expectations and only one set of expectations can be correct.

By clinging to the EMH, research on the effect of heterogenous ability on asset prices has been stifled. Given the existence of insider trading laws, it is quite likely that differences in ability are far greater than information differences in the economy. Indeed, two of the most important puzzles in asset pricing, the volatility of asset prices and the volume of trade in financial markets, might be resolvable in a model that explicitly accounts for differences in abilities.

4 Conclusion

At the time the EMH was developed the profession's understanding of the relation between risk and return was in its infancy. Subsequent research on the EMH contributed enormously to the development of the now well accepted paradigm that links returns to risk. The core of this paradigm is the idea, which follows from the absence of arbitrage alone, that any asset can be priced using a positive pricing kernel. Given the existence of this paradigm, the EMH does not add anything substantive or testable to our understanding of the return generating process in the economy.

The Efficient Market Hypothesis is not a well defined concept, and is either untestable or easily rejected by the data (depending on how it is defined). The widely held interpretation that the hypothesis implies that returns should be unpredictable is highly misleading. Consequently, evidence of predictability is regarded my many financial economists as evidence against market rationality. This misconception presents a barrier to productive research.

Perhaps the most important effect the EMH has had on the development of financial economics is the important role it played in forcing researchers to think carefully about the relation between risk and return. The result of this process is one of the great accomplishments in asset pricing theory — the recognition that a probability measure, the martingale

measure, exists under which all assets have the same expected return (the risk free rate) and all returns are unpredictable. It is universally recognized that the actual probability measure differs from the martingale measure and hence assets do not all have the same expected return. It is somewhat of a puzzle that researchers do not apply the same logic to the predictability of returns — because the true measure differs from the martingale measure, (excess) returns need not be unpredictable.

If returns are related to risk, and risk is persistent, then actual returns will be predictable. Surely this is one of the key results that follows from the paradigm linking risk and return that was developed in the last 50 years. A common argument is that because the unpredictable component of realized returns is so large, it is "as if" returns are unpredictable. But this argument misses two key points. Firstly, although the unpredictable components of returns dwalf the predictable components for individual assets, this fact is not necessarily true for large portfolios. Secondly, and more importantly, the unpredictable component of returns is not the interesting component. The interesting economics are in the predictable components, indeed, it is the predictability of these components that informs us about the underlying economics determining returns. Denying this predictability needlessly restricts our ability to explain what we observe from first principles. It is time for the profession to explicitly recognize this fact by dispensing with the EMH.

References

Berk, J. B., and R. C. Green (2004): "Mutual Fund Flows and Performance in Rational Markets," *Journal of Political Economy*, 112(6), 1269–1295.

Berk, J. B., R. C. Green, and V. Naik (1999): "Optimal Investment, Growth Options, and Security Returns," *The Journal of Finance*, 54(5), 1553–1607.

Conrad, J., G. Kaul, and Q. Lei (2007): "Momentum is Not an Anomaly," Discussion paper, University of Michigan.

Conrad, Jennifer, and Kaul, Gautam (1988): "Time-Variation in Expected Returns," *The Journal of Business*, 61(4), 409–425.

De Bondt, W. F. M., and R. H. Thaler (1987): "Further Evidence on Investor Overreaction and Stock Market Seasonalit y," *Journal of Finance*, 42(3), 557–81, available at http://ideas.repec.org/a/bla/jfinan/v42y1987i3p557-81.html.

DeBondt, W. F. M., and R. M. Thaler (1985): "Does the Stock Market Overreact?," *Journal of Finance*, 40, 793–805.

Fama, E. F. (1965): "The Behavior of Stock-Market Prices," *Journal of Business*, 38(1), 34–105.

——— (1970): "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance*, 25(2), 383–417.

——— (1991): "Efficient Capital Markets: II," *Journal of Finance*, 46, 1575–1617.

LeRoy, S. F. (1973): "Risk Aversion and the Martingale Property of Stock Prices," *International Economic Review*, 14(2), 436–446.

——— (1976): "Efficient Capital Markets: Comment," *The Journal of Finance*, 31(1), 139–141.

Lo, A. W., and A. C. MacKinlay (1988): "Stock Market Prices do not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies*, 1, 41–66.

Lucas, R. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

Muth, J. F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29(3), 315–335.

Ross, S. A. (2004): *Neoclassical Finance*. Princeton University Press.