Inefficient markets and the new finance

Lynn A. Stout
Professor of Law, UCLA School of Law; Principal Investigator, UCLA-Sloan Foundation Research Program on Business Organizations[1]

Abstract

By the early 1980s, the Efficient Capital Market Hypothesis (ECMH) had become one of the most widely-accepted and influential ideas in finance. More recently the idea of market efficiency has fallen into disrepute as a result of market events and growing empirical evidence of inefficiencies. This article argues that the weaknesses of efficient market theory are, and were, apparent from inspection of its initial premises, including the presumptions of homogeneous investor expectations, effective arbitrage, and investor rationality. By the same token, a wide range of phenomena inconsistent with the ECMH can be explained using market models that modify these three assumptions. To illustrate, this article explores three important strands of today's finance literature: (1) work on asset pricing when investors have heterogeneous expectations; (2) scholarship on how and why arbitrage may move public information into prices more slowly and incompletely than earlier writings suggested; and (3) the exploding literature in behavioral finance. Taken together, these literatures provide the framework for building a new and more powerful working model of securities markets.

1 This article is extracted from The Mechanisms of Market Inefficiency: An Introduction to the New Finance, 28 J. Corp. L. 633 (2003). I am deeply indebted to my research assistant, L. Nicolle Hollingsworth, for her invaluable assistance producing this version.

Inefficient markets and the new finance

In 1978, Michael Jensen famously pronounced 'there is no other proposition in economics which has more solid empirical evidence supporting it than the Efficient Markets Hypothesis' (1978, p. 95).[2] Even judged by the evidence of the time this claim included a dash of hyperbole. Yet Jensen's oft-quoted assertion illustrates how efficient market theory had captured the imagination of a generation of economists and finance theorists.

These days we are not so sure. The idea that securities prices reflect informed estimates of value has always coexisted uneasily with a darker view that sees stock prices as disconnected from economic reality. Recent events have fed skepticism. The past decade has seen the formation and bursting of a remarkable price bubble in technology stocks that rivaled the 1637 Dutch tulip bulb craze analyzed by Garber (1989). Indeed, the entire stock market seemed caught in the turbulence. In the Spring of 2000, the Standard & Poors 500 Index topped 1500. Six months later it hovered near 775, a nearly 50% decline in value.

Faced with such market gyrations, it is difficult not to conclude that efficient market theory fails in some fundamental respect to capture the reality of securities markets. This article argues that we need not have waited decades for this insight. Nor need we have suffered the 1990s bubble to find enlightenment. The weaknesses of the efficient market theory were apparent to anyone who cared to look for them, within a few years after the theory was first developed and disseminated.[3]

Like all theories, the Efficient Capital Markets Hypothesis (ECMH) simplifies a complex reality. Almost as soon as 'modern' finance (meaning the version of efficient market theory developed between 1950 and 1970) became prominent, a new generation of scholars began to revise, extend, and elaborate on it. This article explores some significant ideas that have come out of this postmodern work. In particular, it highlights three developing strands of thought in finance: (1) the growing literature on asset pricing when investors have heterogeneous expectations; (2) recent scholarship on the limits of arbitrage; and (3) the emerging discipline of behavioral finance.

Even a brief survey of these literatures suggests a pressing need to reexamine the ECMH's theoretical foundations. It also offers two important lessons. The first lesson of the New Finance is that, well before the mounting empirical evidence against market efficiency achieved critical mass, there was good reason to question the relationship between stock prices and economic value. The second lesson is that even if the relationship between price and value is more indirect and complex than the orthodox ECMH predicts, it does not lie beyond our understanding. There is no need to throw up our hands and abandon our attempt to model markets. To the contrary, if we are willing to modify some of the assumptions that underlie orthodox finance and let the analysis get a bit more complicated, we may be able to develop a much better understanding of how, exactly, securities markets set prices. To describe the current state of finance in the terms of Thomas Kuhn's classic The Structure of Scientific Revolutions (1970), the old paradigm of an efficient market is crumbling. But the outlines of a new paradigm are visible in the resulting cloud of intellectual dust.

Inefficient markets and investor disagreement

To understand the limits of traditional efficient market theory, we must first inquire into what theorists mean by efficient. A market commonly is described as efficient when prices fully reflect available information. This shorthand description does not provide much guidance, however, without a better understanding of what 'fully reflect' means. At least two different understandings of 'fully reflect' appear in the literature (often in the same article, and sometimes used interchangeably). The first focuses on whether information can be used to extract trading profits. As Grossman and Stiglitz observed (1980), a market is 'informationally efficient' with respect to information if a trader who obtains that information cannot make money trading on it. The concept of informational efficiency can be understood as a prediction about how quickly prices respond to new information.

2 Ironically, Jensen's remark appeared in an article in which he questioned market efficiency and noted that "widely scattered and as yet inconclusive evidence is arising which seems to be inconsistent with the theory."

3 By 1981 both the theory underlying the ECMH and its empirical soundness had been called into question by the pioneering work of Miller (1977), Shiller (1981), and others.

Yet informational efficiency, alone, does not imply that market prices respond to new information correctly or even that prices respond at all. This argument is explored later. For present purposes it suffices to note that the connection often drawn between the rapid absorption of information into price, and market prices that mirror economic value, reflects a second common understanding of market efficiency sometimes called 'fundamental value efficiency' [Stout (1995)]. Markets are efficient in the fundamental value sense if prices respond to new information not only quickly but accurately, so that the price reflects the best possible estimate (in light of available information) of economic value in terms of risk and return. As Barberis and Thaler (2003) remind us, informational efficiency and fundamental value efficiency are distinct concepts. Nevertheless, the two can be and often are made to go hand-in-hand, with fundamental value efficiency flowing naturally from informational efficiency. To accomplish this, however, requires a theory of how market participants use information to estimate values and set prices. The most widely-employed of these theories is the Capital Asset Pricing Model (CAPM).

When the CAPM is combined with the ECMH, the two theories produce a joint prediction that informationally efficient prices are also fundamental value efficient. How does this prediction follow from the two theories? Commentators usually assume arbitrage provides the answer, asserting that when prices depart from reasonable estimates of value, sharp-eyed arbitrageurs rush into the market to buy or sell and move prices back into line. This belief in the healing powers of arbitrage is so deeply embedded it has become mental habit, rarely subject to conscious scrutiny.

More careful inquiry reveals this is not the correct explanation. Combining the ECMH with the CAPM produces a prediction of fundamental value efficiency through a different and more troubling analytical path — by tautology. One of the most remarkable (yet remarkably unremarked-upon) characteristics of the classic CAPM is that the pioneers who developed it, including Sharpe (1964) and Lintner (1965), began by expressly assuming that investors share homogenous expectations regarding securities' future risks and returns. In other words, they assumed all investors agree. Yet people do not agree on the intrinsic worth of particular securities. This is why most mutual funds are actively managed, and why many stocks experience 100 percent or more annual turnover.

Despite its empirical unsoundness, the simplifying assumption of homogeneous investor opinion is harmless when CAPM is used for its original purpose of describing the individual investor's attitude towards risk and return. There is danger, however, when CAPM is used to argue that efficient markets price securities accurately. The reasoning becomes circular. If investors make identical estimates of future risk and return, inevitably the market must mirror this 'best' estimate because only one estimate exists — the estimate of the homogenous investor.

This observation sets the stage for introducing one of the most potentially powerful strains of the New Finance — the literature on asset prices in a world where investors disagree. To understand this literature, it is useful to start with the conventional CAPM and ask, what happens if we assume investors disagree? An interesting quirk of the original CAPM quickly emerges. When we acknowledge disagreement while still assuming perfect markets, price-moving arbitrage of the sort assumed by many commentators becomes impossible. Indeed, market equilibrium becomes impossible.

To see why, imagine a simple market with only one security, stock issued by Widget Corporation at $100 per share. Assume there are only two investors: Bull, who thinks Widget is worth $101, and Bear, who thinks Widget worth $99. With no wealth limitations, transactions costs, or short sales restrictions, even this very modest disagreement makes an equilibrium price impossible. Bull will see the chance to buy 'undervalued' Widget stock as a money machine, and will buy and buy until the supply of Widget stock is exhausted. The supply will never be exhausted, however, because Bear simultaneously sees a chance to make money by selling Widget short, and borrows and borrows it (presumably from Bull) to sell it short (again, to

Bull). Bull and Bear thus place infinite bets against each other and no equilibrium emerges.

To achieve equilibrium under conditions of disagreement we need some market imperfection. This complicates the analysis, which may explain why most basic texts teach only the homogeneous-expectations CAPM. Yet those willing to expend the extra intellectual effort needed to explore the heterogeneous expectations literature will be amply rewarded. With just a few plausible modifications, we can build alternative asset pricing models that explain many puzzling characteristics of securities markets that cannot be explained by the CAPM.

In illustration, consider a relatively simple variant of the CAPM first developed in a famous paper by Edward Miller (1977). Miller explored how prices are set in securities markets given three realistic assumptions, investor disagreement, limits on short selling, and risk aversion. For convenience, Miller's model will be labeled the Heterogeneous Expectations Asset Pricing Model (HEAPM)[4]. The HEAPM predicts that even at a very high price, a firm can sell a few shares of stock to a few very optimistic investors. If the firm wants to sell more shares, however, it must lower the price. Lowering price increases demand both because less-optimistic individuals who do not already own the stock become interested and buy, and because existing, more-optimistic but risk-averse shareholders become willing to put more eggs in one basket. The result is a classic downward-sloping demand function for equity securities.

This prediction of downward-sloping demand creates a useful opportunity to test the HEAPM's predictive power vis-à-vis the CAPM. The notion that demand for stock is downward-sloping seems plausible, even obvious, to non-experts. Among finance economists, however, it is (or more accurately once was) a form of heresy. This is because the standard CAPM, when combined with the ECMH, predicts that all stocks should be accurately priced and no stock should be a bargain relative to any other. Thus all stocks should be perfect substitutes. The demand for any particular stock accordingly should be perfectly flat. Raising price above the level set by an efficient market causes investors to refuse to purchase any shares, while lowering price below market creates infinite demand.

The CAPM's prediction of perfectly flat demand ultimately is an artifact of its investor homogeneity assumption. As finance theorists have published more papers exploring asset pricing under conditions of disagreement, the theoretical fragility of flat demand has become apparent. Equally important, numerous empirical studies, including Bagwell (1992), Shleifer (1989), Loderer et. al. (1991), Kaul et al. (2000) and Diether et al (2002) have found that stock demand curves in fact slope downwards[5]. These findings demonstrate that in many respects the HEAPM does a better job than the CAPM of predicting how actual markets behave. They also highlight a basic lesson of the New Finance: when investors disagree, market prices can change for a variety of reasons unrelated to any change in underlying value.

As an example, consider how the HEAPM may help explain the 1990s tech stock bubble [Stout (2000); Ofek & Richardson (2003)]. This explanation focuses on the role 'uncertainty' plays in stock prices. Stock returns are subject not only to risk (predictable variation) but also to uncertainty (unpredictable variation). Uncertainty encourages disagreement, and securities markets are subject to greater uncertainty at some times than others. In periods of greater uncertainty, increased ignorance about the future allows greater disagreement about value. Where investor opinion formerly ranged from pessimistic to optimistic, it now can range from wildly pessimistic to wildly optimistic.

If short sales are restricted, however, only optimists' views influence market price. The HEAPM consequently predicts that stocks of speculative firms whose futures are uncertain trade at higher prices (relative to expected risk and returns) than stocks of stable firms whose futures are more foreseeable. Similarly, it teaches that increased uncertainty that increases the dispersion of investors' opinions can raise prices even when the average investor's opinion of value remains

4 Other papers modeling prices under conditions of disagreement include Harrison & Kreps (1978), Jarrow (1980), Lintner (1969), Mayshar (1983), Sharpe (1970, p. 104-13), Treynor (1998), and Williams (1977).

5 In contrast to the extensive evidence supporting downward-sloping demand, very little evidence supports a horizontal demand function. Most commentators who subscribe to horizontal demand rely on a three-decade-old study by Myron Scholes (1972). Subsequent studies, including Kraus & Stoll (1972) and Mikkelson & Partch (1985) have not replicated the findings Scholes relied on to find horizontal demand.

unchanged. This analysis can be applied to the tech stock bubble. The development of the Internet and other new technologies promised to increase corporate profits, explaining some of the 1990s run-up in stock prices. But the appearance of new technologies also made it harder for investors to predict any single firm's future. Increased investor disagreement and short sales restrictions increased stock prices, especially for uncertain start-ups and information technology firms.

Bubbles are only one example of market phenomena the HEAPM can help explain. Others include, the mystery of why many investors do not fully diversify [Stout (1997)], the price-shifting effects of share issues and repurchases [Stout (1990)], the large premiums paid to target shareholders in takeovers [Stout (1990)], and a variety of so-called market anomalies [Stout (1997)]. One of the most interesting of these anomalies is the 'high-beta' effect, a finding that firms with very high levels of market risk do not, as the CAPM predicts, enjoy proportionately higher earnings. As Miller (1977) pointed out nearly three decades ago, the high-beta effect makes sense under the HEAPM if greater risk is correlated with greater uncertainty.

Yet the HEAPM does more than shed light on hitherto-unsolved market mysteries. It also undermines the theoretical foundation of the claim that stock prices in an informationally efficient market will be fundamental value efficient as well. After all, in a world where investors disagree and short sales are restricted, market price is nothing but the opinion held by the most-pessimistic member of the optimistic subset of investors who choose to hold a security. Why assume that relatively-pessimistic optimist's opinion is correct?

Once we recognize that fundamental value efficiency depends not only on the ECMH but on the CAPM's assumption of investor homogeneity, then the argument that efficient markets value securities accurately is in shambles. Fundamental value efficiency is a theoretical rabbit pulled out of a hypothetical hat. When investors disagree, there is no reason to assume market price necessarily mirrors the best estimate of value in terms of future risk and return. Nor is there reason to be surprised by the growing evidence that it often does not.

Inefficient markets and the limits of arbitrage

As seen above, in a world of disagreement we must question whether securities markets can be efficient in a fundamental value sense. As noted earlier, however, fundamental value efficiency is not the only possible understanding of efficiency. Many theorists who speak of efficient markets seem to be relying on the alternate idea of informational efficiency, that prices respond so quickly to new information it is impossible for traders to make profits on the basis of the information. Indeed, it has become common for finance economists whose faith in fundamental-value efficiency has been beaten out of them by market events and the accumulating evidence on anomalies to retreat to this intellectual position [Brealey & Myers (2000, p. 377); Malkiel (1999, p. 270-74)].

It is important to understand that when economists define market efficiency in terms of the difficulty of making arbitrage profits, they have abandoned the intoxicating and powerful claim that efficient markets produce accurate prices. A market can respond near-instantaneously to new information without producing prices that mirror fundamental value. One could have, for example, a market where prices respond quickly to a popular astrologer's reports, or to an auger's reading of a dead goat's entrails. Once we abandon the CAPM it is easy to imagine markets that are informationally efficient but not fundamental value efficient.

But are markets even informationally efficient? Can arbitrageurs in fact quickly move prices?

In the years immediately following the development of the ECMH it was subjected to extensive empirical testing as researchers analyzed how quickly prices responded to public announcements of stock splits, corporate mergers, and the like. These first studies found that prices seemed to respond to new information almost immediately, within hours or minutes of an announcement. Hence Professor Jensen's claim that 'no

other proposition in economics' had more empirical support than efficient market theory.

But consider the kind of information researchers initially used to test market efficiency. Merger announcements and stock split reports are widely disseminated and easy to understand. What happens when new information becomes available but investors must invest substantial time, trouble, or money to get it? What happens when the information is technical and difficult to understand? Do prices still change within hours or minutes? As Brealey & Myers note (2000, p. 363-65), more recent studies suggest the answers to these questions may be no. Many types of information important to valuing securities seem to be incorporated into prices far more slowly and incompletely than the conventional account of market efficiency suggests.

A good example is the widely-studied phenomenon of 'post-earnings-announcement-drift.' An unanticipated announcement of increased corporate earnings tends to be followed by abnormal positive returns over the next several months, while firms that announce unexpectedly poor earnings see abnormal negative returns over an extended period. Researchers have puzzled over these results. Bernard & Thomas (1989) suggest drift is evidence that the initial price response to the new earnings information is incomplete, and that the full implications of the new earnings information are digested by the market far more slowly than previously suspected.

Several recent papers, including Barberis & Thaler (2003), Gemmill & Thomas (2002), and Mendenhall (2004), explore the limits to arbitrage that explain such a delayed market response. Firstly, arbitrage is not costless. Information is expensive to acquire, process, and verify, and trading also involves costs, especially when would-be arbitrageurs want to sell short. (Indeed, sometimes short selling is essentially impossible.) Secondly, arbitrageurs control only finite amounts of money, and can usually only afford to hold a small percentage of a firm's stock. This means the subjective opinions of other, less well-informed investors will continue (in accord with the HEAPM) to influence market price. Thirdly, arbitrageurs are also likely risk-averse, meaning they will hesitate to take very large positions requiring them to bear otherwise-diversifiable risk. Finally, and perhaps most important, arbitrageurs can only profit from superior information if the rest of the market eventually comes to agree with their value assessments. As they say on Wall Street, 'a bargain that stays a bargain is no bargain.' This is true for individual arbitrageurs, who must worry that while they wait for other investors to come to agree with their assessments, something bad may happen to either the company or the economy as a whole. The problem of the 'bargain that stays a bargain' is even more worrisome for portfolio managers whose careers may depend on their last year's, or even last quarter's, performance.

These realities of trading life paint a picture of market efficiency quite different from the portrait offered in introductory finance texts. Basic texts describe the process of information adjustment in efficient markets as binary, like flipping a switch. Information is either public or private, either 'in' or 'out' of market price. Yet the flow of information into prices may better resemble the flow of a fluid into a vessel. How full the vessel gets, and how quickly, depends on both the diameter of the channel the fluid flows through (how widely the information is disseminated) and the fluid's viscosity (how complex, technical, or difficult to understand the information may be). Information that is easy to understand and that is trumpeted in the media may be incorporated into prices almost instantaneously. Information that is public but difficult to obtain, or information that is complex or requires a specialist's knowledge to comprehend, may take weeks or months to be reflected in price. Indeed, it may never be fully reflected at all.

This view of market efficiency as an imperfect and incremental process offers a number of important lessons. For example, if large numbers of investors are unwilling to invest the effort necessary to understand accounting conventions, and if the opinions of these investors influence market prices, accounting conventions (for example, whether options are expensed) can have a significant impact on market prices.

More obviously, the possibility of informationally inefficient markets undermines the most central implication of the ECMH–that you cannot 'beat the market' by trading on public information. A delayed market response creates an opportunity for the superior investor to make a trading profit. It is important not to exaggerate this opportunity. As noted above, trading inevitably involves transaction costs, and any arbitrageur who tries to profit from an anticipated delay in the market's response to new information must worry that the market will never appreciate the information (a bargain that stays a bargain is no bargain), or will react so slowly that the risk of taking a long-term undiversified position outweighs the likely speculative return. Nevertheless, the incremental model of market inefficiency suggests some traders can reap trading profits from some public information.

What happens in practice? During the ECMH's heyday, most academics believed studies had proven it impossible to beat the market, and that investment professionals who claimed they could were lying to their customers or to themselves. This view was based on a series of 1960s studies finding that mutual and pension funds run by professional portfolio managers on average underperformed market indexes [Ippolito (1993)]. If these institutions could not beat the market, the argument went, neither could anyone else.

Closer inspection reveals a weakness in this analysis. Mutual and pension funds own about half of all equities and account for an even larger percentage of all trading. It would hardly be surprising that institutions on average fail to beat the market when they very nearly are the trading market. Nevertheless, Ippolito (1993, p.42) and Brealey & Myers (2000, p. 361) note that some recent studies have found that mutual funds outperform market indexes enough to offset their research and trading expenses. And whether or not mutual funds on average beat the market, some fund managers rather consistently do so. Warren Buffett is one example of this 'hot hands' phenomenon [Stout (1997)].

A recent study by Coval et al. (2002) finds a similar pattern among individual traders. Coval and his coauthors studied the transaction records of over 110,000 individual traders. On average these individuals underperformed the market net of expenses. The average, however, obscured a striking pattern. While most individual traders did not outperform the market, a small minority did so consistently.

It seems some traders do beat the market. This is not to say the job is easy. It is hard to profit from trading on difficult-to-process information if the minority of investors who pay attention to it react quickly, and the balance remain asleep at the switch and never react at all. Most of us probably would do better to buy an index fund or throw darts at the Wall Street Journal. But we also should think twice before assuming markets are even informationally efficient.

Inefficient markets and investor irrationality

So far we have explored two important branches of the New Finance: asset pricing models that incorporate investor disagreement and investigations of limits to arbitrage. Large and growing literatures exist on each subject. Yet for sheer numbers of recent papers, there are few developments in finance, however intriguing or useful, that rival a third challenge to the traditional ECHM: behavioral finance.

Conventional finance, like conventional economics, adopts the 'homo economicus' model of human behavior. It assumes humans are rational actors with stable preferences who care only about their own welfare. Real people, however, are not always rational. Sometimes our emotions mislead us; sometimes we make foolish mistakes. The fundamental insight of behavioral finance is that emotion and error influence investment choices just as they influence decisions to play lotteries, use drugs, or wear seatbelts.

The trick is to figure out how this influence operates. Behavioral finance theorists rely on the psychological literature, and especially on studies of human behavior in experimental games, to identify forms of cognitive bias that cause people to make mistakes. They then examine whether these

biases can explain or predict market anomalies that cannot be explained or predicted by traditional finance.

It is difficult to overstate just how rapidly the behavioral finance literature has grown over the past decade. Scores of papers and several books have been produced on the topic[6]. Indeed, the number of behavioral finance papers being produced now rivals scholarly production in traditional finance.

At the same time, many behavioral finance studies that have captured public attention are not the sort to convince a skeptic that the field has much to offer in terms of developing our structural understanding of securities markets. For example, one recent study by Dichev and Janes (2003) reports that stock prices are influenced by the lunar cycle, while another by Kamstra et al. (2003) concludes that seasonal affective disorder causes stock returns to rise and fall with the seasonal lengthening and shortening of daylight hours. A third 'behavioral finance' theory that has been discussed in the national press [Teitelbaum (2000)] explains the 1990s market bubble as a consequence of increased use of antidepressants like Prozac and Zoloft, with an attendant collective surge in investor optimism. (Sadly, this entertaining theory suffers from a number of flaws, most obviously its inability to explain the Crash of 2000 absent evidence that the investing public suddenly and collectively went 'off its meds.')

These kinds of studies have raised behavioral finance's popular profile. They have also contributed to a perception that behavioral finance is more suitable for dinner party conversation than serious research. This perception has been reinforced by the rather large number of cognitive biases argued to influence securities prices, some of them quirky, some of them short-lived, and some of them apparently contradictory. The result has been an impression that, while behavioral finance may be useful to arbitrageurs, it has little to offer theorists beyond the prediction that securities prices depart from informed estimates of value in arbitrary and capricious ways.

This prediction is hardly trivial. A more careful review of the behavioral finance literature reveals, however, that in addition to offering insight into the effects of seasonal and lunar cycles, it can explain market phenomena far more enduring and consequentially.

Behavioral finance may, for example, explain the 'equity premium puzzle,' the observation that stock market returns seem much greater relative to the risk-free rate of return than can be explained by the market's somewhat higher level of risk. Bernartzi and Thaler (1995) have suggested that the equity premium reflects a psychological phenomenon called 'loss aversion,' an odd human tendency to neurotically dislike decreases in wealth, so much so that otherwise risk-averse actors accept additional risk to gain a chance of avoiding losses. According to this hypothesis, people demand a very large premium to hold stocks because, relative to risk-free investments like Treasury bills, stocks carry a much higher probability that in any particular year the investor's portfolio will decline rather than rise in value. This prospect of intermittent loss is so psychologically painful that investors tend to shun stocks. Thus loss aversion leads investors to refuse to hold equity securities unless they receive a disproportionately large return by doing so.

Behavioral finance may also further our understanding of price bubbles. As discussed earlier, heterogeneous-expectations asset pricing models offer a partial explanation for bubbles by predicting that when short sales are limited, an exogenous increase in the dispersion of opinions raises prices, while resolution of uncertainty brings prices back to prior levels. This explanation, while logical, seems insufficient to explain extreme bubbles like the 1990s tech-stock bubble or the 1637 tulip frenzy. Behavioral finance sheds additional light.

Behavioral studies find that people suffer from a form of systematic cognitive bias that Hirshleifer (2001) describes as the 'availability effect,' Barberis et al. (1998) call the 'representativeness heuristic,' Shiller (2000, p. 60-61) labels 'adaptive expectations,' and I have described as 'trust' [Stout (2002)]. Whatever label one prefers, this bias is familiar to any casual

6 For useful surveys see Barberis & Thaler (2003) and Hirshleifer (2001).

observer of human behavior: people extrapolate too readily from past events.

This tendency to extrapolate reconciles two well-established anomalies that seem, at first, inconsistent with each other: the finding that stock prices under-react to sudden departures from past patterns, and the finding that long-term changes in returns tend to be followed by reversals in a fashion that suggests investors overreact to perceived trends. Fama (1998 pp. 6-7) has used the apparent tension between these findings to criticize behavioral finance, suggesting it fails to explain why 'the same investors under-react to some types of events and over-react to others.' If people extrapolate too readily from the past, however, it makes sense that the same investors who under-react to earnings surprises also overreact to long-term shifts that seem to evidence a trend.

This possibility suggests a joint explanation for extreme bubbles like the 1990s tech stock bubble. In brief, new information technologies modestly increased many corporations' profits and, more dramatically, introduced uncertainty that increased the dispersion of investor opinion. The result (as predicted by the HEAPM) was a rise in market prices, especially for information technology companies whose stocks were difficult to short. This price rise persisted long enough that investors perceived a trend and revised their estimates of stock value upwards. This revision increased prices further, confirming the trend, leading to another upward revision in valuations, leading to a further price increase, in a self-reinforcing process that eventually produced prices experts viewed as unreasonable.

At this point the bubble was ripe for bursting. Burst it did. As the future became less murky and uncertainty resolved, prices declined modestly. This created the perception of a new downward trend, causing investors to revise their value estimates downwards, depressing prices further, confirming the trend, and so on. Thus combining two strands of the New Finance — heterogeneous expectations models and behavioral finance — allows us to understand an otherwise baffling market event.

Conclusion

By the early 1980s, market efficiency had achieved the status of stylized fact among economists. The years since have seen the steady accretion of empirical evidence against efficiency, including but not limited to evidence of pricing anomalies, downward-sloping demand, excessive volatility, delayed information response, and consistently superior traders. Now, in the light of three decades of accumulated experience, a disinterested observer would likely conclude the efficient market paradigm is crumbling.

Yet if efficient market theory is failing, what can replace it? Skepticism about the validity of the ECMH should not lead to skepticism about our ability to understand and predict market behavior. Indeed, the framework for a new understanding may be in place. In particular, three strands of work in contemporary finance — the expanding literature on heterogeneous expectations asset pricing models, scholarly work on the limits of arbitrage, and the emerging field of behavioral finance — show great promise as new ways of thinking about markets. Taken separately, each of these three bodies of work can explain market puzzles that have proven intractable under orthodox efficient market theory. Taken together, they offer the intellectual framework for building a new theory that permits numerous testable predictions. As this article suggests, such testable predictions include, issuing more shares decreases price while share repurchases raise price, short sales restrictions raise price while developments that make shorting easier decrease price, increased uncertainty raises price, prices under-react to information presented to the public in a form that is difficult to access or understand, limited opportunities for profitable arbitrage exist that can be exploited by a minority of traders, prices under-react to information surprises but overreact to perceived trends, equities offer excess returns relative to risk-free investments even after adjusting for non-diversifiable risk, and, last but not least, market price need not bear any close relationship to the best estimate, in light of the available information, of a security's expected risk and return.

As this partial list indicates, the New Finance paints a picture of capital markets that is far more nuanced and complex than the simple portrait sketched by what is still sometimes called 'modern finance' (described by one former Wall Street rocket scientist as 'finance theory with vacuum tubes')[7]. What is more, this portrait bears a much closer resemblance to the way actual market participants — executives, investment bankers, stock analysts, portfolio managers — think the market behaves. The possibility of convergence between the views of academics and those of practitioners should, perhaps, be viewed as cause for optimism. It would be worrisome if so many apparently intelligent people (for clarity it should be noted, with due deference to academics, that the practitioners are being referred to here) could be so wrong for so long about phenomena they experience firsthand and have great incentive to observe accurately.

Similar optimism permits the suggestion that a new paradigm in finance is not only urgently needed, it is well on its way to emerging. There is room for doubt in this regard. Many of the ideas that provide the foundations for the New Finance have been around for nearly as long as the ECMH itself. Edward Miller, for example, published his heterogeneous expectations asset pricing paper in 1977. A quarter-century later, most basic finance texts still discuss in detail only the conventional ECMH and CAPM. If alternatives are mentioned at all, it is briefly, or in footnotes.

Why is the intellectual progress so slow? Thomas Kuhn (1970) offered part of the answer in his influential book The Structure of Scientific Revolutions. As Kuhn observed, the individuals who pioneer a new theory are often those who are most reluctant to revise it in the face of contradictory evidence. Intellectual progress must await the arrival of a new generation of scholars less committed to defending the dominant paradigm.

This may be only part of the story, however. The remainder can be found in the New Finance itself. The New Finance teaches that the degree of efficiency we observe in markets ultimately depends on a number of variables, including the extent to which investors have rational and homogeneous expectations and the ease with which information flows into prices. Just as investors often disagree in their opinions of value, theoreticians often disagree in their views of the merits of different models. Just as complex and technical information may move into prices relatively slowly, complex and nuanced theories may move into the academic culture more slowly than theories easily captured in a sound bite. Finally, just as investors may irrationally overweight evidence of past performance, scholars may overweight past judgments about a theory's validity. A better understanding of how securities markets work accordingly gives us insight into other markets, including the market for ideas.

7 Peter Huang coined this wonderful phrase.

References

- Bagwell, L. S., 1992, "Dutch auction repurchases: An analysis of shareholder heterogeneity," Journal of Finance, 47:1, 71-105
- Barberis, N. C., and R. H. Thaler, "A survey of behavioral finance," in G. M. Constantinides, M. Harris, and R. M. Stultz, eds., 2003, "Handbook of the Economics of Finance," Elsevier Science Publishers, Amsterdam, The Netherlands, 1054-1123
- Barberis, N. C., A. Shleifer, and R. W. Vishny, 1998, "A model of investor sentiment," Journal of Financial Economics, 49:3, 307-343
- Bernard, V. L., and J. K. Thomas, 1989, "Post-earnings-announcement drift: Delayed price response or risk premium?," Journal of Accounting Research, 27 (supplement), 1-36
- Benartzi, S., and R. H. Thaler, 1995, "Myopic loss aversion and the equity premium puzzle," Quarterly Journal of Economics, 110:1, 73-92
- Brealey, R. A., and S. C. Myers, 2000, "Principles of Corporate Finance," 6th ed., Irwin/McGraw-Hill, Boston, Massachusetts, United States of America
- Cooney, J. W., C. F. Loderer, and L. D. Van Drunen, 1991, "The price elasticity of demand for common stock," Journal of Finance, 46:2, 621-651
- Coval, J. D., D. A. Hirshleifer, and T. Shumway, 2002, "Can individual investors beat the market?," Harvard NOM Research Paper No. 02-45, available online at http://ssrn.com/abstract=364000
- Dichev, I. D., and T. D. Janes, 2003, "Lunar cycle effects in stock returns," Journal of Private Equity, 6:4, 8-29
- Diether, K. B., C. J. Malloy, and A. Scherbina, 2002, "Differences of opinion and the cross section of stock returns," Journal of Finance, 57:5, 2113-2141
- Garber, P. M., 1989, "Tulipmania," Journal of Political Economy, 97:3, 535-560
- Gemmill, G., and D. C. Thomas, 2002, "Noise trading, costly arbitrage, and asset prices: Evidence from closed-end funds," Journal of Finance, 57:6, 2571-2594
- Fama, E. F., 1970, "Efficient capital markets: A review of theory and empirical work," Journal of Finance, 25:2, 383-417
- Fama, E. F., 1998, "Market efficiency, long-term returns, and behavioral finance," Journal of Financial Economics, 49:3, 283-306
- Grossman, S. J., and J. E. Stiglitz, 1980, "On the impossibility of informationally efficient markets," American Economic Review, 70:3, 393-408
- Harrison, J. M., and D. M. Kreps, 1978, "Speculative investor behavior in a stock market with heterogeneous expectations," Quarterly Journal of Economics, 92:2, 323-336
- Hirshleifer, D. A., 2001, "Investor psychology and asset pricing," Journal of Finance, 56:4, 1533-1597
- Ippolito, R. A., 1993, "On studies of mutual fund performance, 1962-1991," Financial Analysts Journal, 49:1, 42-50
- Jarrow, R. A., 1980, "Heterogeneous expectations, restrictions on short sales, and equilibrium asset prices," Journal of Finance, 35:5, 1105-1113
- Kamstra, M. J., L. A. Kramer, and M. D. Levi,, 2003, "Winter blues: seasonal affective disorder (SAD) and stock market returns," American Economic Review, 93:1, 324-343
- Kaul, A., V. Mehrotra, and R. Morck, 2000, "Demand curves for stocks do slope down: New evidence from an index weights adjustment," Journal of Finance, 55:2, 893-912
- Kraus, A., and H. R. Stoll, 1972, "Price impacts of block trading on the New York Stock Exchange," Journal of Finance, 27:3, 569-588
- Kuhn, T. S., 1970, "The Structure of Scientific Revolutions," 2nd ed., University of Chicago Press, Chicago, Illinois, United States of America
- Jensen, M. C., 1978, "Some anomalous evidence regarding market efficiency," Journal of Financial Economics, 6:2/3, 95-101
- Lintner, J., 1965, "The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets," Review of Economics and Statistics, 47:1, 13-37
- Lintner, J., 1969, "The aggregation of investors' diverse judgments and preferences in purely competitive securities markets," Journal of Financial and Quantitative Analysis, 4:4, 347-400
- Malkiel, B. G., 1999, "A Random Walk Down Wall Street," 7th ed., W. W. Norton and Company, New York, New York, United States of America
- Mayshar, J., 1983, "On divergence of opinion and imperfections in capital markets," American Economic Review, 73:1, 114-128
- Mendenhall, R. R., 2004, "Arbitrage risk and post-earnings announcement drift," Journal of Business, 77:4, 875-894
- Mikkelson, W. H., and M. M. Partch, 1985, "Stock price effects and costs of secondary distributions," Journal of Financial Economics, 14:2, 165-194
- Miller, E. R., 1977, "Risk, uncertainty, and divergence of opinion," Journal of Finance, 32:4, 1151-1168
- Ofek, E., & Richardson, M. P., 2003, "DotCom mania: The rise and fall of internet stock prices," Journal of Finance, 58:3, 1113-1138
- Scholes, M. S., 1972, "The market for securities: Substitution versus price pressure and the effects of information," Journal of Business, 45:2, 179-211
- Sharpe, W. F., 1964, "Capital asset prices: A theory of market equilibrium under conditions of risk," Journal of Finance, 19:3, 425-442
- Sharpe, W. F., 1970, "Portfolio Theory and Capital Markets," McGraw-Hill, New York, New York, United States of America
- Shiller, R. J., 1981, "Do stock prices move too much to be justified by subsequent changes in dividends?," American Economic Review, 71:3, 421-436
- Shiller, R. J., 2000, "Irrational Exuberance," Princeton University Press, Princeton, New Jersey, United States of America
- Shleifer, A., 2000, "Inefficient Markets: An Introduction to Behavioral Finance," Oxford University Press, New York, New York, United States of America
- Shleifer, A., 1989, "Do demand curves for stocks slope down?," Journal of Finance, 41:3, 579-590
- Stout, L. A., 1990, "Are takeover premiums really premiums? Market price, fair value and corporate law," Yale Law Journal, 99:6, 1235-1296
- Stout, L. A., 1995, "Are stock markets costly casinos? Disagreement, market failure, and securities regulation," Virginia Law Review, 81:3, 611-712
- Stout, L. A., 1997, "How efficient markets undervalue stocks: CAPM and ECMH under conditions of uncertainty and disagreement," Cardozo Law Review, 19:1/2, 475-492
- Stout, L. A., 2000, "Stock prices and social wealth," Harvard Law and Economics Discussion Paper No. 301, Georgetown Law and Economics Research Paper No. 249991, available online at http://papers.ssrn.com/abstract=249991
- Stout, L. A., 2002, "The investor confidence game," 68:2 Brooklyn Law Review, 68:2, 407-437
- Treynor, J., 1998, "Bulls, bears, and market bubbles," Financial Analysts Journal, 54:2, 69-74
- Williams, J. T., 1977, "Capital asset prices with heterogeneous beliefs," Journal of Financial Economics, 5:2, 219-239