

COPYRIGHT NOTICE:

Lo & MacKinlay: *A Non-Random Walk down Wall Street* is published by Princeton University Press and copyrighted, © 1999, by Princeton University Press. All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher, except for reading and browsing via the World Wide Web. Users are not permitted to mount this file on any network servers.

For COURSE PACK and other PERMISSIONS, refer to entry on previous page. For more information, send e-mail to permissions@pupress.princeton.edu

1

Introduction¹

ONE OF THE EARLIEST and most enduring models of the behavior of security prices is the Random Walk Hypothesis, an idea that was conceived in the sixteenth century as a model of games of chance.² Closely tied to the birth of probability theory, the Random Walk Hypothesis has had an illustrious history, with remarkable intellectual forbears such as Bachelier, Einstein, Lévy, Kolmogorov, and Wiener.

More recently, and as with so many of the ideas of modern economics, the first serious application of the Random Walk Hypothesis to financial markets can be traced back to Paul Samuelson (1965), whose contribution is neatly summarized by the title of his article: “Proof that Properly Anticipated Prices Fluctuate Randomly.” In an informationally efficient market—not to be confused with an allocationally or Pareto-efficient market—price changes must be unforecastable if they are properly anticipated, i.e., if they fully incorporate the expectations and information of all market participants. Fama (1970) encapsulated this idea in his pithy dictum that “prices fully reflect all available information.”

Unlike the many applications of the Random Walk Hypothesis in the natural and physical sciences in which randomness is assumed almost by default, because of the absence of any natural alternatives, Samuelson argues that randomness is achieved through the active participation of many investors seeking greater wealth. Unable to curtail their greed, an army of investors aggressively pounce on even the smallest informational advantages at their disposal, and in doing so, they incorporate their information into market prices and quickly eliminate the profit opportunities that gave rise to their aggression. If this occurs instantaneously, which it must in an idealized world of “frictionless” markets and costless trading, then prices must always fully reflect all available information and no profits can be gar-

¹Parts of this introduction are adapted from Lo (1997a,b) and Lo and MacKinlay (1998).

²See, for example, Hald (1990, Chapter 4).

nered from information-based trading (because such profits have already been captured). This has a wonderfully counter-intuitive and seemingly contradictory flavor to it: the more efficient the market, the more random the sequence of price changes generated by such a market, and the most efficient market of all is one in which price changes are completely random and unpredictable.

For these reasons, the Random Walk Hypothesis and its close relative, the Efficient Markets Hypothesis, have become icons of modern financial economics that continue to fire the imagination of academics and investment professionals alike. The papers collected in this volume comprise our own foray into this rich literature, spanning a decade of research that we initiated in 1988 with our rejection of the Random Walk Hypothesis for US stock market prices, and then following a course that seemed, at times, to be self-propelled, the seeds of our next study planted by the results of the previous one.

If there is one central theme that organizes the papers contained in this volume, it is this: financial markets *are* predictable to some degree, but far from being a symptom of inefficiency or irrationality, predictability is the oil that lubricates the gears of capitalism. Indeed, quite by accident and rather indirectly, we have come face to face with an insight that Ronald Coase hit upon as an undergraduate over half a century ago: price discovery is neither instantaneous nor costless, and frictions play a major role in determining the nature of competition and the function of markets.

1.1 The Random Walk and Efficient Markets

One of the most common reactions to our early research was surprise and disbelief. Indeed, when we first presented our rejection of the Random Walk Hypothesis at an academic conference in 1986, our discussant—a distinguished economist and senior member of the profession—asserted with great confidence that we had made a programming error, for if our results were correct, this would imply tremendous profit opportunities in the stock market. Being too timid (and too junior) at the time, we responded weakly that our programming was quite solid thank you, and the ensuing debate quickly degenerated thereafter. Fortunately, others were able to replicate our findings exactly, and our wounded pride has healed quite nicely with the passage of time (though we still bristle at the thought of being prosecuted for programming errors without “probable cause”). Nevertheless, this experience has left an indelible impression on us, forcing us to confront the fact that the Random Walk Hypothesis was so fully ingrained into the canon of our profession that it was easier to attribute our empirical results to programming errors than to accept them at face value.

Is it possible for stock market prices to be predictable to some degree in an efficient market?

This question hints at the source of disbelief among our early critics: an implicit—and incorrect—link between the Random Walk Hypothesis and the Efficient Markets Hypothesis. It is not difficult to see how the two ideas might be confused. Under very special circumstances, e.g., risk neutrality, the two are equivalent. However, LeRoy (1973), Lucas (1978), and many others have shown in many ways and in many contexts that the Random Walk Hypothesis is neither a necessary nor a sufficient condition for rationally determined security prices. In other words, unforecastable prices need not imply a well-functioning financial market with rational investors, and forecastable prices need not imply the opposite.

These conclusions seem sharply at odds with Samuelson's "proof" that properly anticipated prices fluctuate randomly, an argument so compelling that it is reminiscent of the role that uncertainty plays in quantum mechanics. Just as Heisenberg's uncertainty principle places a limit on what we can know about an electron's position and momentum if quantum mechanics holds, Samuelson's version of the Efficient Markets Hypothesis places a limit on what we can know about future price changes if the forces of economic self-interest hold.

Nevertheless, one of the central insights of modern financial economics is the necessity of some trade-off between risk and expected return, and although Samuelson's version of the Efficient Markets Hypothesis places a restriction on expected returns, it does not account for risk in any way. In particular, if a security's expected price change is positive, it may be just the reward needed to attract investors to hold the asset and bear the associated risks. Indeed, if an investor is sufficiently risk averse, he might gladly *pay* to avoid holding a security that has unforecastable returns.

In such a world, the Random Walk Hypothesis—a purely statistical model of returns—need not be satisfied even if prices do fully reflect all available information. This was demonstrated conclusively by LeRoy (1973) and Lucas (1978), who construct explicit examples of informationally efficient markets in which the Efficient Markets Hypothesis holds but where prices do not follow random walks.

Grossman (1976) and Grossman and Stiglitz (1980) go even further. They argue that perfectly informationally efficient markets are an *impossibility*, for if markets are perfectly efficient, the return to gathering information is nil, in which case there would be little reason to trade and markets would eventually collapse. Alternatively, the degree of market *inefficiency* determines the effort investors are willing to expend to gather and trade on information, hence a non-degenerate market equilibrium will arise only when there are sufficient profit opportunities, i.e., inefficiencies, to compensate investors for the costs of trading and information-gathering. The profits

earned by these industrious investors may be viewed as economic rents that accrue to those willing to engage in such activities. Who are the providers of these rents? Black (1986) gives us a provocative answer: noise traders, individuals who trade on what they think is information but is in fact merely noise. More generally, at any time there are always investors who trade for reasons other than information—for example, those with unexpected liquidity needs—and these investors are willing to “pay up” for the privilege of executing their trades immediately.

These investors may well be losing money on average when they trade with information-motivated investors, but there is nothing irrational or inefficient about either group’s behavior. In fact, an investor may be trading for liquidity reasons one day and for information reasons the next, and losing or earning money depending on the circumstances surrounding the trade.

1.2 The Current State of Efficient Markets

There is an old joke, widely told among economists, about an economist strolling down the street with a companion when they come upon a \$100 bill lying on the ground. As the companion reaches down to pick it up, the economist says “Don’t bother—if it were a real \$100 bill, someone would have already picked it up.”

This humorous example of economic logic gone awry strikes dangerously close to home for students of the Efficient Markets Hypothesis, one of the most important controversial and well-studied propositions in all the social sciences. It is disarmingly simple to state, has far-reaching consequences for academic pursuits and business practice, and yet is surprisingly resilient to empirical proof or refutation. Even after three decades of research and literally thousands of journal articles, economists have not yet reached a consensus about whether markets—particularly financial markets—are efficient or not.

What can we conclude about the Efficient Markets Hypothesis? Amazingly, there is still no consensus among financial economists. Despite the many advances in the statistical analysis, databases, and theoretical models surrounding the Efficient Markets Hypothesis, the main effect that the large number of empirical studies have had on this debate is to harden the resolve of the proponents on each side.

One of the reasons for this state of affairs is the fact that the Efficient Markets Hypothesis, by itself, is not a well-defined and empirically refutable hypothesis. To make it operational, one must specify additional structure, e.g., investors’ preferences, information structure, business conditions, etc. But then a test of the Efficient Markets Hypothesis becomes a test of several auxiliary hypotheses as well, and a rejection of such a joint hypothesis tells

us little about which aspect of the joint hypothesis is inconsistent with the data. Are stock prices too volatile because markets are inefficient, or is it due to risk aversion, or dividend smoothing? All three inferences are consistent with the data. Moreover, new statistical tests designed to distinguish among them will no doubt require auxiliary hypotheses of their own which, in turn, may be questioned.

More importantly, tests of the Efficient Markets Hypothesis may not be the most informative means of gauging the efficiency of a given market. What is often of more consequence is the *relative* efficiency of a particular market, relative to other markets, e.g., futures vs. spot markets, auction vs. dealer markets, etc. The advantages of the concept of relative efficiency, as opposed to the all-or-nothing notion of absolute efficiency, are easy to spot by way of an analogy. Physical systems are often given an efficiency rating based on the relative proportion of energy or fuel converted to useful work. Therefore, a piston engine may be rated at 60% efficiency, meaning that on average 60% of the energy contained in the engine's fuel is used to turn the crankshaft, with the remaining 40% lost to other forms of work, e.g., heat, light, noise, etc.

Few engineers would ever consider performing a statistical test to determine whether or not a given engine is perfectly efficient—such an engine exists only in the idealized frictionless world of the imagination. But measuring relative efficiency—relative to a frictionless ideal—is commonplace. Indeed, we have come to expect such measurements for many household products: air conditioners, hot water heaters, refrigerators, etc. Therefore, from a practical point of view, and in light of Grossman and Stiglitz (1980), the Efficient Markets Hypothesis is an idealization that is economically unrealizable, but which serves as a useful benchmark for measuring relative efficiency.

A more practical version of the Efficient Markets Hypothesis is suggested by another analogy, one involving the notion of thermal equilibrium in statistical mechanics. Despite the occasional “excess” profit opportunity, on average and over time, it is not possible to earn such profits consistently without some type of competitive advantage, e.g., superior information, superior technology, financial innovation, etc. Alternatively, in an efficient market, the only way to earn positive profits *consistently* is to develop a competitive advantage, in which case the profits may be viewed as the economic rents that accrue to this competitive advantage. The consistency of such profits is an important qualification—in this version of the Efficient Markets Hypothesis, an occasional free lunch is permitted, but free lunch plans are ruled out.

To see why such an interpretation of the Efficient Markets Hypothesis is a more practical one, consider for a moment applying the classical version of the Efficient Markets Hypothesis to a non-financial market, say the

market for biotechnology. Consider, for example, the goal of developing a vaccine for the AIDS virus. If the market for biotechnology is efficient in the classical sense, such a vaccine can *never* be developed—if it could, someone would have already done it! This is clearly a ludicrous presumption since it ignores the difficulty and gestation lags of research and development in biotechnology. Moreover, if a pharmaceutical company does succeed in developing such a vaccine, the profits earned would be measured in the billions of dollars. Would this be considered “excess” profits, or economic rents that accrue to biotechnology patents?

Financial markets are no different in principle, only in degrees. Consequently, the profits that accrue to an investment professional need not be a market *inefficiency*, but may simply be the fair reward to breakthroughs in financial technology. After all, few analysts would regard the hefty profits of Amgen over the past few years as evidence of an inefficient market for pharmaceuticals—Amgen’s recent profitability is readily identified with the development of several new drugs (Epogen, for example, a drug that stimulates the production of red blood cells), some considered breakthroughs in biotechnology. Similarly, even in efficient financial markets there are very handsome returns to breakthroughs in financial technology.

Of course, barriers to entry are typically lower, the degree of competition is much higher, and most financial technologies are not patentable (though this may soon change) hence the “half life” of the profitability of financial innovation is considerably smaller. These features imply that financial markets should be relatively more efficient, and indeed they are. The market for “used securities” is considerably more efficient than the market for used cars. But to argue that financial markets must be perfectly efficient is tantamount to the claim that an AIDS vaccine cannot be found. In an efficient market, it is difficult to earn a good living, but not impossible.

1.3 Practical Implications

Our research findings have several implications for financial economists and investors. The fact that the Random Walk Hypothesis hypothesis can be rejected for recent US equity returns suggests the presence of predictable components in the stock market. This opens the door to superior long-term investment returns through disciplined active investment management. In much the same way that innovations in biotechnology can garner superior returns for venture capitalists, innovations in financial technology can garner equally superior returns for investors.

However, several qualifications must be kept in mind when assessing which of the many active strategies currently being touted is appropriate for an particular investor. First, the riskiness of active strategies can be very

different from passive strategies, and such risks do not necessarily “average out” over time. In particular, an investor’s risk tolerance must be taken into account in selecting the long-term investment strategy that will best match the investor’s goals. This is no simple task since many investors have little understanding of their own risk preferences, hence consumer education is perhaps the most pressing need in the near term. Fortunately, computer technology can play a major role in this challenge, providing scenario analyses, graphical displays of potential losses and gains, and realistic simulations of long-term investment performance that are user-friendly and easily incorporated into an investor’s world view. Nevertheless, a good understanding of the investor’s understanding of the nature of financial risks and rewards is the natural starting point for the investment process.

Second, there are a plethora of active managers vying for the privilege of managing institutional and pension assets, but they cannot all outperform the market every year (nor should we necessarily expect them to). Though often judged against a common benchmark, e.g., the S&P 500, active strategies can have very diverse risk characteristics and these must be weighed in assessing their performance. An active strategy involving high-risk venture-capital investments will tend to outperform the S&P 500 more often than a less aggressive “enhanced indexing” strategy, yet one is not necessarily better than the other.

In particular, past returns should not be the *sole* or even the *major* criterion by which investment managers are judged. This statement often surprises investors and finance professionals—after all, isn’t this the bottom line? Put another way, “If it works, who cares why?”. Selecting an investment manager this way is one of the surest paths to financial disaster. Unlike the experimental sciences such as physics and biology, financial economics (and most other social sciences) relies primarily on statistical inference to test its theories. Therefore, we can never know with perfect certainty that a particular investment strategy is successful since even the most successful strategy can always be explained by pure luck (see Chapter 8 for some concrete illustrations).

Of course, some kinds of success are easier to attribute to luck than others, and it is precisely this kind of attribution that must be performed in deciding on a particular active investment style. Is it luck, or is it genuine?

While statistical inference can be very helpful in tackling this question, in the final analysis the question is not about statistics, but rather about economics and financial innovation. Under the practical version of the Efficient Markets Hypothesis, it is difficult—but not impossible—to provide investors with consistently superior investment returns. So what are the sources of superior performance promised by an active manager and why have other competing managers not recognized these opportunities? Is it better mathematical models of financial markets? Or more accurate statisti-

cal methods for identifying investment opportunities? Or more timely data in a market where minute delays can mean the difference between profits and losses? Without a compelling argument for *where* an active manager's value-added is coming from, one must be very skeptical about the prospects for future performance. In particular, the concept of a "black box"—a device that performs a known function reliably but obscurely—may make sense in engineering applications where repeated experiments can validate the reliability of the box's performance, but has no counterpart in investment management where performance attribution is considerably more difficult. For analyzing investment strategies, it matters a great deal *why* a strategy is supposed to work.

Finally, despite the caveats concerning performance attribution and proper motivation, we *can* make some educated guesses about where the likely sources of value-added might be for active investment management in the near future.

- The revolution in computing technology and datafeeds suggest that highly computation-intensive strategies—ones that could not have been implemented five years ago—that exploit certain regularities in securities prices, e.g., clientele biases, tax opportunities, information lags, can add value.
- Many studies have demonstrated the enormous impact that transactions costs can have on long-term investment performance. More sophisticated methods for measuring and controlling transactions costs—methods which employ high-frequency data, economic models of price impact, and advanced optimization techniques—can add value. Also, the introduction of financial instruments that reduce transactions costs, e.g., swaps, options, and other derivative securities, can add value.
- Recent research in psychological biases inherent in human cognition suggest that investment strategies exploiting these biases can add value. However, contrary to the recently popular "behavioral" approach to investments which proposes to take advantage of individual "irrationality," I suggest that value-added comes from creating investments with more attractive risk-sharing characteristics suggested by psychological models. Though the difference may seem academic, it has far-reaching consequences for the long-run performance of such strategies: taking advantage of individual irrationality cannot be a recipe for long-term success, but providing a better set of opportunities that more closely matches what investors desire seems more promising.

Of course, forecasting the sources of future innovations in financial technology is a treacherous business, fraught with many half-baked successes and some embarrassing failures. Perhaps the only reliable prediction is that the innovations of future are likely to come from unexpected and

underappreciated sources. No one has illustrated this principal so well as Harry Markowitz, the father of modern portfolio theory and a winner of the 1990 Nobel Prize in economics. In describing his experience as a Ph.D. student on the eve of his graduation in the following way, he wrote in his Nobel address:

. . . [W]hen I defended my dissertation as a student in the Economics Department of the University of Chicago, Professor Milton Friedman argued that portfolio theory was not Economics, and that they could not award me a Ph.D. degree in Economics for a dissertation which was not Economics. I assume that he was only half serious, since they did award me the degree without long debate. As to the merits of his arguments, at this point I am quite willing to concede: at the time I defended my dissertation, portfolio theory was not part of Economics. But now it is.

It is our hope and conceit that the research contained in this volume will be worthy of the tradition that Markowitz and others have so firmly established.