

Issue Number 52

September 1997

A publication of External Affairs — Corporate Research

A Nonrandom Walk Down Wall Street: Recent Advances in Financial Technology

In this issue:

Introduction

*Stock Market Prices and
the Random Walk*

The Martingale Model

The Random Walk Hypothesis

Rejecting the Random Walk

*Implications for Investment
Management*

The Efficient Markets Hypothesis

A Modern View of Efficient Markets

Practical Considerations

In this issue of Research Dialogues, Andrew W. Lo, Harris & Harris Group Professor, Sloan School of Management, Massachusetts Institute of Technology, takes a look at an old issue in investment risk management in a new way. In the '50s and '60s, just as the era of the professional portfolio manager was dawning, financial economists were telling anyone who would listen that active management was probably a big mistake—a waste of time and money. Their research demonstrated that historical prices were of little use in helping to predict where future prices would go. Prices simply took a “random walk.” The better part of wisdom, they advised, was to be a passive investor.

At first, not too many of the people who influence the way money is managed

(those who select managers of large portfolios) listened. But as time went on, it became apparent that they should have. Because of fees and turnover, the managers they picked typically underperformed the market. And the worse an active manager did relative to a market index, the more attractive seemed the low-cost alternative of buying and holding the index itself.

But as luck would have it, just as indexing was gaining ground, a new wave of academic research was being published that weakened some of the results of the earlier research and thereby undercut part of the justification for indexing. It didn't obviate all the reasons for indexing (indexing was still a low-cost way to create diversification for an entire fund or as part of an active/passive strategy), but it did tend to silence the index-because-you-can't-do-better school.

Andrew Lo was, and continues to be, at the forefront of this new research. In 1988, for example, he and MacKinlay (see references, page 7) published evidence that resulted in the rejection of the Random Walk Hypothesis. That evidence is summarized below and fittingly constitutes the first step in Professor Lo's interesting nonrandom walk.

Introduction

If, in January 1926, an investor put \$1 into one-month U.S. Treasury bills—one of the “safest” assets in the world—and continued reinvesting the proceeds in Treasury bills month by month until

December 1996, the \$1 investment would have grown to \$14. If, on the other hand, the investor had put \$1 into the stock market, e.g., the S&P 500, and continued reinvesting the proceeds in the stock market month by month over this same seventy-year period, the \$1 investment would have grown to \$1,370, a considerably larger sum. Now suppose that each month, an investor could tell in advance which of these two investments would yield a higher return for that month, and took advantage of this information by switching the running total of his initial \$1 investment into the higher-yielding asset, month by month. What would a \$1 investment in such a “perfect foresight” investment strategy yield by December 1996?

The startling answer—\$2,303,981,824 (yes, over \$2 billion; this is no typographical error)—often comes as a shock to even the most seasoned professional investment manager. Of course, few investors have perfect foresight. But this extreme example suggests that even a modest ability to forecast financial asset returns may be handsomely rewarded: it does not take a large fraction of \$2,303,981,824 to beat \$1,370! For this reason, quantitative models of the risks and rewards of financial investments—now known collectively as *financial technology* or *financial engineering*—have become virtually indispensable to institutional investors throughout the world.

Of course, financial engineering is still in its infancy when compared with the mathematical and natural sciences. However, it has enjoyed a spectacular

period of growth over the past three decades, thanks in part to the breakthroughs pioneered by academics such as Fischer Black, John Cox, Harry Markowitz, Robert Merton, Stephen Ross, Paul Samuelson, Myron Scholes, William Sharpe, and many others. Moreover, parallel breakthroughs in mathematics, statistics, and computational power have enabled the financial community to implement such financial technology almost immediately, giving it an empirical relevance and practical urgency shared by few other disciplines in the social sciences.

This brief article describes one simple example of modern financial technology: testing the Random Walk Hypothesis—the hypothesis that past prices cannot be used to forecast future prices—for aggregate U.S. stock market indexes. Although the random walk is a very old idea, dating back to the sixteenth century, recent studies have shed new light on this important model of financial prices. The findings suggest that stock market prices do contain predictable components and that there may be significant returns to active investment management.

Stock Market Prices and the Random Walk

One of the most enduring questions of financial economics is whether or not financial asset prices are forecastable. Perhaps because of the obvious analogy between financial investments and games of chance, mathematical models of asset prices have an unusually rich history that predates virtually every other aspect of economic analysis. The vast number of prominent mathematicians and scientists who have plied their considerable skills in forecasting stock and commodity prices is a testament to the fascination and the challenges that this problem poses. Indeed, the intellectual roots of modern financial economics are firmly planted in early attempts to “beat the market,” an endeavor that is still of current interest and is discussed and debated even in the most recent journals, conferences, and cocktail parties!

The Martingale Model

One of the earliest mathematical models of financial asset prices was the *martingale*

model, whose origins lie in the history of games of chance and the birth of probability theory.¹ The prominent Italian mathematician Girolamo Cardano proposed an elementary theory of gambling in his 1565 manuscript *Liber de ludo aleae* (*The Book of Games of Chance*), in which he writes:

The most fundamental principle of all in gambling is simply equal conditions, e.g., of opponents, of bystanders, of money, of situation, of the dice box, and of the die itself. To the extent to which you depart from that equality, if it is in your opponent’s favour, you are a fool, and if in your own, you are unjust.²

This clearly contains the notion of a “fair game,” a game which is neither in your favor nor your opponent’s, and this is the essence of a *martingale*, a precursor to the Random Walk Hypothesis (RWH). If P_t represents one’s cumulative winnings or wealth at date t from playing some game of chance each period, then a fair game is one in which the expected wealth next period is simply equal to this period’s wealth.

If P_t is taken to be an asset’s price at date t , then the martingale hypothesis states that tomorrow’s price is expected to be equal to today’s price, given the asset’s entire price history. Alternatively, the asset’s expected price change is zero when conditioned on the asset’s price history; hence its price is just as likely to rise as it is to fall. From a forecasting perspective, the martingale hypothesis implies that the “best” forecast of tomorrow’s price is simply today’s price, where the “best” forecast is defined to be the one that minimizes the average squared error of the forecast.

Another implication of the martingale hypothesis is that nonoverlapping price changes are uncorrelated at all leads and lags, which further implies the ineffectiveness of all *linear* forecasting rules for future price changes that are based on the price history. The fact that so sweeping an implication could come from so simple a model foreshadows the central role that the martingale hypothesis plays in the modeling of asset price dynamics (see, for example, Huang and Litzenberger [1988]).

In fact, the martingale was long considered to be a necessary condition for an *efficient* asset market, one in which the

information contained in past prices is instantly, fully, and perpetually reflected in the asset’s current price. However, one of the central ideas of modern financial economics is the necessity of some trade-off between risk and expected return, and although the martingale hypothesis places a restriction on expected returns, it does not account for risk in any way.

In particular, if an asset’s expected price change is positive, it may be the reward necessary to attract investors to hold the asset and bear its associated risks. Indeed, if an investor is risk averse, he would gladly *pay* to avoid holding an asset with the martingale property. Therefore, despite the intuitive appeal that the “fair game” interpretation might have, it has been shown that the martingale property is neither a necessary nor a sufficient condition for rationally determined asset prices (see, for example, LeRoy [1973], and Lucas [1978]).

Nevertheless, the martingale has become a powerful tool in probability and statistics, and also has important applications in modern theories of asset prices. For example, once asset returns are properly adjusted for risk, the martingale property does hold (see Lucas [1978], Cox and Ross [1976], and Harrison and Kreps [1979]), and the combination of this risk adjustment and the martingale property has led to a veritable revolution in the pricing of complex financial instruments such as options, swaps, and other *derivative* securities. Moreover, the martingale led to the development of a closely related model that has now become an integral part of virtually every scientific discipline concerned with dynamic uncertainty: the Random Walk Hypothesis.

The Random Walk Hypothesis

The simplest version of the RWH states that future returns cannot be forecast from past returns. For example, under the RWH, if stock XYZ performed poorly last month, this has no bearing on how XYZ will perform this month, or in any future month. In this respect, the RWH is not unlike a sequence of fair-coin tosses: the fact that one toss comes up heads, or that a sequence of five tosses is composed of all heads, has no implications for what the

next toss is likely to be. In short, past returns cannot be used to forecast future returns under the RWH. In the jargon of economic theory, all information contained in past returns has been impounded into the current market price: therefore nothing can be gleaned from past returns if the goal is to forecast the next price *change*, i.e., the future return.

The RWH has dramatic implications for the typical investor. In its most stringent form—the *independently-and-identically-distributed* or IID version—it implies that optimal investment policies do not depend on historical performance, so that a spell of below-average returns does not require rethinking the wisdom of the optimal policy. A somewhat less restrictive version of the RWH—the *independent-returns version*—allows historical performance to influence investment policies, but rules out the efficacy of nonlinear forecasting techniques. For example, “charting” or “technical analysis”—the practice of predicting future price movements by attempting to spot geometrical shapes and other regularities in historical price charts—will not work if the independent-returns RWH is true. And even the least

restrictive version of the RWH—the *uncorrelated-increments* version—has sharp implications: it rules out the efficacy of linear forecasting techniques such as regression analysis.³

All versions of the RWH have one common implication: the volatility of returns must increase one-for-one with the return horizon. For example, under the RWH, the volatility of two-week returns must be exactly twice the volatility of one-week returns; the volatility of four-week returns must be exactly twice the volatility of two-week returns; and so on. Therefore, one test of the RWH is to compare the volatility of two-week returns with twice the volatility of one-week returns. If they are close, this lends support for the RWH; if they are not, this suggests that the RWH is false.

This aspect of the RWH is particularly relevant for long-term investors: under the RWH, the riskiness of an investment—as measured by the return variance—increases linearly with the investment horizon and is not “averaged out” over time (see Bodie [1996] for further discussion). Therefore, any claim that an investment becomes less risky over time must be based on the assumption that the RWH does not hold and

that variances grow less than linearly with the return horizon.

Rejecting the Random Walk

Whether or not the variance ratio is close to 1 depends on what “close” means; Lo and MacKinlay [1988] provide a precise statistical measure in their *variance ratio* statistic. They apply the variance ratio statistic to two broad-based weekly indexes of U.S. equity returns—equal- and value-weighted indexes of all securities traded on the New York and American Stock Exchanges—derived from the University of Chicago’s Center for Research in Securities Prices (CRSP) daily stock returns database.⁴

Despite the fact that the CRSP monthly database begins in 1926, whereas the daily database begins in 1962, Lo and MacKinlay choose to construct weekly returns from the daily database; hence their data span the period from 1962 to 1992.⁵ They focus on weekly returns for two reasons: (1) more-recent data are likely to be more relevant to current practice—the institutional features of equity markets and investment behavior are considerably different now than in the 1920s; (2) since their test is based on variances, it is not the calendar time span that affects the accuracy of their estimates but rather the sample size, and weekly data are the best compromise between maximizing the sample size and minimizing the effects of market frictions, e.g., the bid/ask spread, that affect daily data.

Lo and MacKinlay’s findings are summarized in Table 1. The values reported in the main rows are the ratios of the variance of q -week returns to q times the variance of one-week returns, and the entries enclosed in parentheses are measures of “closeness” to 1, where larger values represent more statistically significant deviations from the RWH.⁶ Panel A contains results for the equal-weighted index and Panel B contains similar results for the value-weighted index. Within each panel, the first row presents the variance ratios and test statistics for the entire 1,568-week sample and the next two rows give the results for the two equally partitioned 784-week subsamples.

Table 1
Variance Ratio Test of the RWH

Sample period	Number nq of base observations	Number q of base observations aggregated to form variance ratio			
		2	4	8	16
A. CRSP NYSE/AMEX Equal-Weighted Index					
621212-921223	1,568	1.27 (5.29)*	1.58 (6.66)*	1.86 (7.15)*	1.94 (5.93)*
621212-771214	784	1.31 (5.99)*	1.67 (7.18)*	1.98 (6.80)*	2.11 (5.41)*
771221-921223	784	1.23 (2.40)*	1.47 (3.01)*	1.70 (3.52)*	1.72 (2.91)*
B. CRSP NYSE/AMEX Value-Weighted Index					
621212-921223	1,568	1.06 (1.54)*	1.12 (1.64)*	1.15 (1.43)*	1.13 (0.86)*
621212-771214	784	1.07 (1.49)*	1.16 (1.73)*	1.22 (1.59)*	1.26 (1.24)*
771221-921223	784	1.06 (0.87)*	1.09 (0.76)*	1.07 (0.48)*	0.98 (0.09)*

Variance ratios for weekly equal- and value-weighted indexes of all NYSE and AMEX stocks, derived from the CRSP daily stock returns database. The variance ratios are reported in the main rows, with the heteroskedasticity-robust test statistics $z^*(q)$ given in parentheses immediately below each main row. Under the Random Walk Null hypothesis, the value of the variance ratio is 1 and the test statistics have a standard normal distribution asymptotically. Test statistics marked with asterisks indicate that the corresponding variance ratios are statistically different from 1 at the 5 percent level of significance.

Table 1 shows that the RWH can be rejected at all the usual significance levels for the entire time period and all subperiods. For example, the typical cutoff value for the measure of closeness is ± 1.96 —a value between -1.96 and $+1.96$ indicates that the variance ratio is close enough to 1 to be consistent with the RWH. However, a value outside this range indicates an inconsistency between the RWH and the behavior of volatility over different return horizons. The value 5.29, corresponding to the variance ratio 1.27 of two-week returns to one-week returns of the equal-weighted index, demonstrates a gross inconsistency between the RWH and the data. The other variance ratios in Table 1 document similar inconsistencies.

The statistics in Table 1 also tell us *how* the RWH is inconsistent with the data: variances grow faster than linearly with the return horizon. In particular, for the equal-weighted index, the entry 1.27 in the “ $q=2$ ” column implies that two-week returns have a variance that is 27 percent higher than twice the variance of one-week returns. This suggests that the riskiness of an investment in the equal-weighted index increases with the return horizon; there is “antidiversification” over time in this case.

Of course, it should be emphasized that this antidiversification applies only to return horizons from $q=2$ to $q=16$. For much longer return horizons, e.g., $q=150$ or greater, there is some weak evidence that variances grow less than linearly with the return horizon. However, those inferences are not based on as many data points and are considerably less accurate than the shorter-horizon inferences drawn from Table 1 (see Campbell, Lo, and MacKinlay [1997, Chapter 2] for further discussion).

Although the test statistics in Table 1 are based on nominal stock returns, it is apparent that virtually the same results would obtain with real or excess returns. Since the volatility of weekly nominal returns is so much larger than that of the inflation and Treasury-bill rates, the use of nominal, real, or excess returns in a volatility-based test will yield practically identical inferences.

Implications for Investment Management

The rejection of the RWH for U.S. equity indexes raises the possibility of improving investment returns by active portfolio management. After all, if the RWH implies that stock returns are unforecastable, its rejection would seem to say that stock returns *are* forecastable. But several caveats are in order before we begin our attempts to “beat the market.”

First, the rejection of the RWH is based on historical data, and past performance is no guarantee of future success. While the statistical inferences performed by Lo and MacKinlay [1988] are remarkably robust, nevertheless *all* statistical inference requires a certain willing suspension of disbelief regarding structural changes in institutions and business conditions.

Second, we have not considered the impact of trading costs on investment performance. While a rejection of the RWH implies a degree of predictability in stock returns, the trading costs associated with exploiting such predictability may outweigh the benefits. Without a more careful investigation of trading costs, it is virtually impossible to assess the economic significance of the rejections reported in Table 1.

Finally, and perhaps most important, suppose the rejections of the RWH are both statistically and economically significant, even after adjusting for trading costs and other institutional frictions; does this imply some sort of “free lunch”? Should all investors, young and old, attempt to forecast the stock market and trade more actively according to such forecasts? In other words, is the stock market *efficient* or can one achieve superior investment returns by trading intelligently?

The Efficient Markets Hypothesis

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 (EMH), one of the most controversial and well-studied propositions in all the social sciences. Briefly, the EMH states that in an efficient market, all available information is fully reflected in current market prices.⁷ Therefore, any attempt to forecast future price movements is futile—any information on which the forecast is based has already been impounded into the current price.

The EMH 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.

One of the reasons for this state of affairs is the fact that the EMH, 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, etc. But then a test of the EMH 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. For example, academics in the 1960s equated the EMH with the RWH, but more recent studies by LeRoy [1973] and Lucas [1978] have shown that the RWH may be violated in a perfectly efficient market.

More important, tests of the EMH 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 pis-

ton engine may be rated at 60 percent efficiency, meaning that on average, 60 percent of the energy contained in the engine's fuel is used to turn the crankshaft, with the remaining 40 percent 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 the 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 EMH is an idealization that is economically unrealizable, but serves as a useful benchmark for measuring relative efficiency.

A Modern View of Efficient Markets

A more practical version of the EMH 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 EMH, an occasional free lunch is permitted, but free lunch plans are ruled out.

To see why such an interpretation of the EMH is a more practical one, consider for a moment applying the classical version of the EMH to a nonfinancial 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 rather the fair economic return that accrues 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 EMH, it is difficult to earn a good living, but not impossible.

Practical Considerations

These recent research findings have several implications for both long-term investors in defined-contribution pension plans and for plan sponsors. The fact that the RWH hypothesis can be rejected for recent U.S. equity returns suggests the presence of predictable components in the stock market. This opens the door to superior long-term investment returns through dis-

ciplined 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 being touted is appropriate for a particular investor. First, the riskiness of active strategies can be very different from that of passive strategies, and such risks do not necessarily “average out” over time. In particular, the investor's risk tolerance must be taken into account in selecting the long-term investment strategy that will best match his or her goals. This is no simple task, since many investors have little understanding of their own risk preferences. Consumer education is therefore 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 worldview. 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 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 performance 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 in-

vestment manager solely by past performance 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 Lo [1994] and Lo and MacKinlay [1990] 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 EMH, 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 statistical 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 it 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 moti-

vation, 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 data sources suggests that highly computation-intensive strategies—ones that could not have been implemented five years ago—which 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 transaction costs can have on long-term investment performance. More sophisticated methods for measuring and controlling transaction costs—methods that employ high-frequency data, economic models of price impact, and advanced optimization techniques—can add value. Also, the introduction of financial instruments that reduce transaction 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 match 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 the future are likely to come from unexpected

and underappreciated sources. No one has illustrated this principle 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."⁸ □

References

- Bodie, Z., 1996, "Risks in the Long Run," *Financial Analysts Journal* 52, 18-22.
- Campbell, J., A. Lo, and C. MacKinlay, 1997, *The Econometrics of Financial Markets*. Princeton, NJ: Princeton University Press.
- Cox, J., and S. Ross, 1976, "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics* 3, 145-166.
- Fama, E., 1970, "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance* 25, 383-417.
- Grossman, S., and J. Stiglitz, 1980, "On the Impossibility of Informationally Efficient Markets," *American Economic Review* 70, 393-408.
- Hald, A., 1990, *A History of Probability and Statistics and Their Applications Before 1750*. New York: John Wiley & Sons.
- Harrison, J., and D. Kreps, 1979, "Martingales and Arbitrage in Multi-period Securities Markets," *Journal of Economic Theory* 20, 381-408.
- Huang, C., and R. Litzenberger, 1988, *Foundations for Financial Economics*. New York: North Holland.
- LeRoy, S., 1973, "Risk Aversion and the Martingale Property of Stock Returns," *International Economic Review* 14, 436-446.
- Lo, A., 1994, "Data-Snooping Biases in Financial Analysis," in H. Russell Fogler, ed.: *Blending Quantitative and Traditional Equity Analysis*. Charlottesville, VA: Association for Investment Management and Research.

Lo, A., and C. MacKinlay, 1988, "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies* 1, 41-66.

Lo, A., and C. MacKinlay, 1990, "Data Snooping Biases in Tests of Financial Asset Pricing Models," *Review of Financial Studies* 3, 431-468.

Lucas, R. E., 1978, "Asset Prices in an Exchange Economy," *Econometrica* 46, 1429-1446.

Roberts, H., May 1967, "Statistical versus Clinical Prediction of the Stock Market," unpublished manuscript, Center for Research in Security Prices, University of Chicago.

Samuelson, P., 1965, "Proof That Properly Anticipated Prices Fluctuate Randomly," *Industrial Management Review* 6, 41-49.

Endnotes

- ¹ The etymology of *martingale* as a mathematical term is unclear. Its French root has two meanings: a leather strap tied to a horse's bit and reins to ensure that they stay in place, and a gambling system involving doubling the bet when one is losing.
- ² See Hald (1990, Chapter 4) for further details.
- ³ See Campbell, Lo, and MacKinlay (1997, Chapter 2) for a more detailed discussion of the various versions of the RWH.
- ⁴ An *equal-weighted* index is one in which all stocks are weighted equally in computing the index, e.g., the Value-Line Index. A *value-weighted* index is one in which all stocks are weighted by their market capitalization; hence larger-capitalization stocks are more influen-

tial in determining the behavior of the index, e.g., the S&P 500 index.

- ⁵ Lo and MacKinlay's (1988) original sample period was 1962 to 1985; it has been extended to 1992 in Table 1, and the qualitative features remain unchanged.
- ⁶ Specifically, the parenthetical entries are *z*-statistics, which are approximately normally distributed with zero mean and unit variance.
- ⁷ Although the EMH has a long and illustrious history, three authors are usually credited with its development: Fama (1970), Roberts (1967), and Samuelson (1965).
- ⁸ Tore Frängsmyr, ed., *Les Prix Nobel* (Stockholm: Norstedts Tryckeri AB, 1990), 302.

