Efficient Irrational Markets

Investors do not need to be rational for markets to be efficient.

Les Gulko

Before we can explain why people commit mistakes, we must first explain why they should ever be right.

—Friedrich A. Hayek
Economics and Knowledge, 1937

Individual rationality lies at the heart of neoclassical economics. Both the rational expectations hypothesis in macroeconomics and the efficient market hypothesis in financial economics are rooted in individual rationality; see, e.g., Fama [1970, 1998]. Empirical evidence, though, suggests that investors often act irrationally; examples are in Benartzi and Thaler [2001], Black [1986], Kahneman and Tversky [1979], and Odean [1998]. And it is a common belief that if investors are not rational, then, market prices are not efficient; see, e.g., Shiller [2000] and Shleifer [2000].

My thesis is that investors do not need to be rational for market prices to be efficient. This article provides both empirical evidence and a theoretical model of an efficient capital market inhabited by irrational investors.

The empirical evidence is based on interest rate forecasts of leading Wall Street economists published in *The Wall Street Journal* since 1981. On the one hand, the prevailing spot rates are better predictors of future interest rates than the experts. This conclusion supports the efficient markets hypothesis. On the other hand, the experts routinely deviate from a superior forecast, the spot rate, and make inferior predictions instead. This behavior is irrational because it ignores useful market information.[1]

Les Gulko is a private money manager in Stamford, CT.
lesgulko@hotmail.com

EXHIBIT 1
Forecast Samples over Time

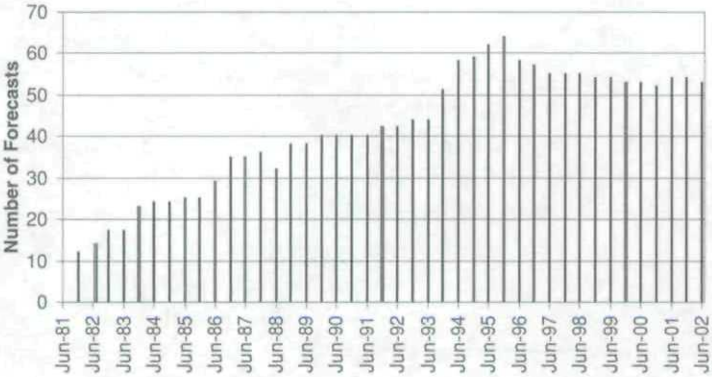

This evidence suggests that price efficiency is not necessarily a product of individual rationality. A system composed of a large number of primitive and imperfect components is able to perform intelligent and reliable functions if the individual components interact and work in concert. Likewise, a securities market inhabited by imperfect investors can function efficiently if the investors communicate and coordinate their decisions.

According to Hayek [1937], investors communicate and coordinate their decisions through market prices. According to Georgescu-Roegen [1971], communication through prices is subject to an "Oedipus effect": Investors observe market prices and then randomly revise their expectations. As a result, investors hold diverse expectations even after observing a superior forecast.

I present a maximum-likelihood model of an efficient market populated by irrational heterogeneous investors. The model produces a distribution of interest rate forecasts that agrees with *The Wall Street Journal* poll. The model also entails the absence of arbitrage, and reproduces the CAPM, the APT, and the Black-Scholes model (see Gulko [1999]). Therefore, financial theory and practice predicated on the efficient markets hypothesis (e.g., index investing) is valid whether or not investors act rationally. Malkiel [2003] offers empirical evidence in support of this view.

EMPIRICAL EVIDENCE

The Wall Street Journal has published semiannual interest rate forecasts of the leading Wall Street economists since December 1981. On the last business day of every June and December, the economists have made predictions of 3-month and 30-year interest rates six months into the future. The individual forecasts are published at the beginning of July and January following the forecast.

Exhibit 1 shows the number I_t of individual forecasts compiled at time t, where $t = 1, 2, \ldots, T$. The sample size increased steadily from under 20 in the early 1980s to over 60 in the mid-1990s, and has remained between 50 and 60 since then. My statistical analysis covers 37 semiannual periods, starting in December 1983, the first time the sample size exceeded 20, and ending in June 2002; thus, T = 37. (In June 2002, the forecasts of the 10-year rate replaced forecasts of the 30-year rate.)

Suppose that at time t the spot rate is r_t and six months later at time $t + 1$ it is r_{t+1}. At time t, I_t experts observe the spot rate r_t and make predictions of the future spot rate r_{t+1}. Let f_{ti} be the individual forecast of expert i at time t. The consensus forecast f_t at time t is defined as the average:

$$f_t = \frac{1}{I_t}\sum_{i=1}^{I_t} f_{ti} \tag{1}$$

In theory, the spot rate r_t absorbs all available information about future interest rates. Hence, r_t represents a *naive market forecast* of the future interest rate r_{t+1}. In order to evaluate the accuracy of the experts' consensus forecast f_t relative to the naive market forecast r_t, a forecast model is estimated by linear regression:

$$r_{t+1} = \alpha r_t + \beta\,(f_t - r_t) + \varepsilon$$

where the historical record of interest rates r_t and r_{t+1} is obtained from Bloomberg; the consensus forecast f_t comes from *The Wall Street Journal*; coefficients α and β are estimated by regression; and ε is a normal random error with zero mean.

This forecast model is similar to that in Muth [1961], except it assumes that the experts use the current spot rate as a reference point. In other words, the forecast model is based on the hypothesis that the experts observe the spot rate r_t and then adjust it by their consensus outlook $(f_t - r_t)$. If this hypothesis is correct, then the α-estimate is approximately one, and the β-estimate determines the accuracy of the experts' foresight as follows:

If $\alpha = 1$ and $\beta > 0$, then the experts are correct on average, about the direction of interest rates.

EXHIBIT 2
Regression Estimation of Forecast Model

Forecast of	No. of f_t observations	α-estimate (t-statistic)	β-estimate (t-statistic)	R^2
3-month rate	37	0.96 (38.4)	0.43 (1.0)	0.77
30-year rate	37	0.96 (59.7)	-1.60 (-3.1)	0.84

If $\alpha = 1$ and $\beta = 0$, then the experts' consensus has zero predictive value.

If $\alpha = 1$ and $\beta < 0$, then the experts are wrong on average about the direction of interest rates.

Exhibit 2 presents the regression estimates of α and β. The α-estimate is approximately one for the forecasts of both 3-month rates and 30-year rates, thus confirming that the experts use the spot rate r_t as a reference point when they form their consensus $(f_t - r_t)$. In the forecast of the three-month rate, the β-estimate is statistically indistinguishable from zero. Therefore, the experts as a group are no more informative than the spot rate r_t, and their consensus view has no predictive value.

In the forecast of the 30-year rate, the β-estimate is negative. Therefore, the experts are generally wrong about the direction of the 30-year interest rate. It is more beneficial to act contrary to their consensus views than to follow them. The regression results do not imply that the current spot rate is a good predictor of the future interest rate; they imply only that the current spot is a better predictor than the experts.

Next, I examine the properties of individual forecasts. Let the *relative forecast* at time t be the difference between a forecast made at time t and the prevailing spot rate r_t. For example, $f_{ti} - r_t$ is the relative individual forecast of expert i at time t, and $f_t - r_t$ is the relative consensus forecast at time t.

Exhibit 3 shows the time series of the descriptive statistics for the relative individual forecasts $f_{ti} - r_t$. The thick solid line plots the relative consensus forecast $f_t - r_t$. The two thinner solid lines above and below $f_t - r_t$ demarcate the band $(f_t - r_t) \pm \sigma_t$, where σ_t is the standard deviation of the relative individual forecasts $f_{ti} - r_t$ at time t. The two dotted lines mark the highest and the lowest values of $f_{ti} - r_t$ at time t. The vertical distance between the dotted lines measures the range of individual forecasts at time t. This range often exceeds 2%.

Large deviations of the individual forecasts from the prevailing spot rate have persisted over time, little affected by the level of interest rates. As interest rates have declined from double digits in the early 1980s to the low single digits in 2002, the range of individual forecasts has remained around 2% since the mid-1980s.

While the experts hold diverse individual outlooks f_{ti}, their consensus forecast f_t tends to converge to the prevailing spot rate r_t. Exhibit 3 shows that the relative consensus forecast $f_t - r_t$, represented by the thick solid line, is close to zero most of the time for both 3-month interest rates and 30-year interest rates.

EXHIBIT 3
Time Series of Forecast Statistics

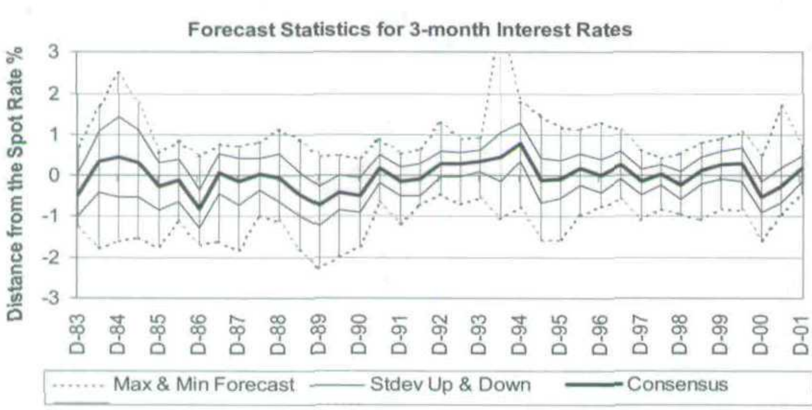

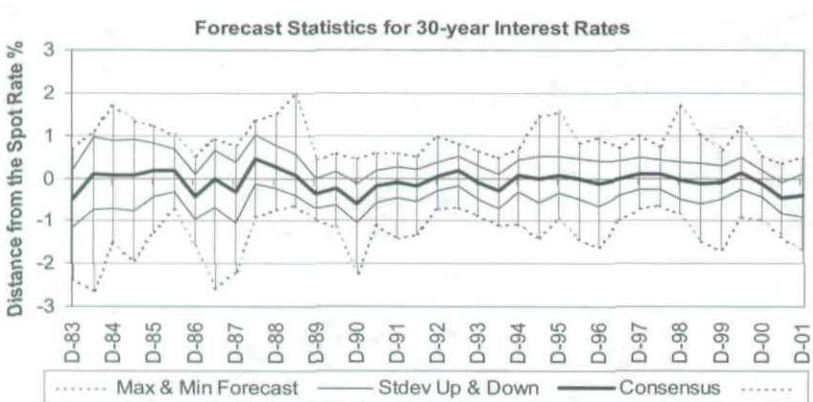

Thick solid line: relative consensus forecast $f_t - r_t$. The two thin solid lines demarcate the band $(f_t - r_t) \pm \sigma_t$ around $f_t - r_t$, where σ_t is the standard deviation of $f_{ti} - r_t$ at time t. The two dotted lines mark the largest and the smallest values of $f_{ti} - r_t$ at time t. The vertical distance between the dotted lines measures the full range of individual forecasts at time t.

EXHIBIT 4

Descriptive Statistics of Random Variable Representing Relative Individual Forecast

Forecast of	Total Number of Observations $\sum I_t$	Mean of $(f_{ti} - r_t)$	Median of $(f_{ti} - r_t)$	Std Deviation of $(f_{ti} - r_t)$
3-month rate	1646	-0.015%	0.010%	0.565%
30-year rate	1649	-0.078%	-0.070%	0.516%

The cross-sectional statistics in Exhibit 4 confirm these tendencies. The mean of $f_{ti} - r_t$ is well under 10 basis points, and the standard deviation is over 50 bp for both 3-month rates and 30-year rates. This means that the experts hold highly diverse views, while their consensus closely matches the prevailing spot. Three observations summarize the empirical evidence:

Observation 1. The experts' consensus forecast has no predictive value, and the spot rate r_t is a better predictor of the future spot rate r_{t+1} than the experts' consensus.

Observation 2. The experts' consensus forecast f_t tends to the prevailing spot rate r_t.

Observation 3. Individual forecasts f_{ti} have a wide dispersion around the spot rate r_t.

IMPLICATIONS

The empirical study supports both price efficiency and individual irrationality. Investors are said to be *rational* only if they fully use available information and *irrational* if they do not fully use available information.[2] Market prices are said to be *efficient* only if they are not predictable and *inefficient* if they are predictable. According to the *efficient markets hypothesis* (EMH), market prices reflect all available information; hence, it is impossible to consistently anticipate price changes (see Fama [1970, 1998]).

By Observation 1, the experts as a group display no predictive ability, and the prevailing spot rates are more informative than the experts' consensus. By Observation 2, the experts' consensus forecast of the future rates tends toward the prevailing spot rate, meaning that the spot rate aggregates the experts' expectations before they are reported to *The Wall Street Journal*. This evidence is consistent with the EMH and modern financial theory.

Therefore, an educated forecast must be close to the prevailing spot. Yet, by Observation 3, individual forecasts f_{ti} often fall far away from the prevailing spot rate r_t. This means that the experts observe the spot, known to be a superior forecast, and then issue poorer forecasts. The experts act irrationally as they ignore useful market information. Aggarwal, Mohanty, and Song [1995] and Schirm [2003] also find—using different criteria—that economic forecasts often lack rationality.

Although the definition of price efficiency does not involve individual rationality, it is usually assumed that price efficiency is a product of individual rationality. Furthermore, any evidence against investors' rationality automatically serves as evidence against the EMH. Yet Observations 1, 2, and 3 suggest that *price efficiency* and *individual irrationality* are not necessarily mutually exclusive—they may coexist in the same market. Then individuals do not need to be rational for market prices to be efficient.

How do prices reach efficiency if investors are not rational? Complex systems with imperfect components can *function properly if the components are interconnected so* that flaws are detected and corrected endogenously. Redundant relay networks and the human nervous system are examples of well-functioning systems with imperfect components.

Both relay networks and the human nervous system consist of a large number of primitive elements—binary relays and neurons, respectively—interconnected so that, even if some of them fail, the entire system may still function well. What makes these systems intelligent and reliable is not the perfection of individual components but rather the multiplicity and interaction of primitive and imperfect components acting in concert.

Likewise, a securities market is an information-processing system consisting of a large ensemble of imperfect elements. What makes the market system efficient is not the high quality of individual elements but instead the multiplicity and interaction of low-quality elements. In this market, price efficiency is not a product of individual knowledge and intelligence; rather, it is a product of coordinated actions of many imperfect agents.

Investors communicate and coordinate their decisions via the price system (Hayek [1945]). The communication protocol consists of three steps. First, investors receive messages by observing market prices. Second, investors use these prices to revise their expectations and valuations. Third, they send messages in the form of buy orders, sell orders, or no orders. Steps 1 and 3 are self-explanatory. To examine how investors revise their expectations, let us consider Observation 3.

Why do the expert forecasts differ from the spot rate r_t, which represents a superior forecast? This paradox can be explained by the Oedipus effect: *Individuals revise their views and actions after they observe the actions of others; the revisions are unpredictable and often contradict the actions observed.* As a result, individuals draw different conclusions from the same evidence. Georgescu-Roegen [1971, p. 335] was the first to recognize this effect in economics:

> The Oedipus effect . . . boils down to this: the announcement of an action to be taken changes the evidence upon which each individual bases his expectations and, hence, causes him to revise his previous plans. Preferences too may be subject to an Oedipus effect. One may prefer a Rolls Royce to a Cadillac but perhaps not if he is told that his neighbor, too, will buy a Rolls Royce.

Likewise, *The Wall Street Journal* forecasters observe the spot rate r_t and then deviate from it, even though they are aware of the informational superiority of market prices and of the verifiable inferiority of the experts' forecasts. The deviations are random both in size and in direction. This leads to a wide dispersion of forecasts and apparent disagreement among experts. The result is that the experts, homogeneous in terms of their professional backgrounds, experiences, and information sets, make highly heterogeneous predictions.

The Oedipus effect reflects a simple fact that competitive people tend to disagree with one another. Therefore, investors always hold diverse views. As non-satiation, or greed, is an intrinsic *individual property* that fuels competition, the diversity of individual views is an intrinsic *group property* that enables competition. No economic competition is possible unless individuals want more than what they own and think differently from others.

Textbook models, though, usually assume that investors share the same views. As a result, these models run up against a dilemma known as the no-trade theorem, discussed in Ross [1987]. According to the no-trade theorem, investors sharing the same expectations never trade.

This absurd conclusion is indicative of flawed assumptions. Financial models that admit heterogeneous beliefs do not suffer from this problem.

MAXIMUM-LIKELIHOOD MODEL

A statistical model of *The Wall Street Journal* polls can be specified as follows. At time t, N experts observe the spot rate r_t and make predictions of the future rate r_{t+1}. Let S be a random variable representing the relative individual forecast, i.e., the deviation of the individual forecast from the prevailing spot rate r_t. The objective is to construct a maximum-likelihood distribution of the relative individual forecast S.

First, suppose S takes on $m = 41$ possible values S_k, $k = 1, 2, \ldots, 41$, uniformly spaced between -2.0% and $+2.0\%$. These values constitute the state space Ω:

$$\Omega = \{-2.0\%, -1.9\%, -1.8\%, \ldots, -0.1\%, 0\%, 0.1\%, \ldots, 1.8\%, 1.9\%, 2.0\%\}$$

Each forecaster makes a prediction by selecting one state S_k, $k = 1, 2, \ldots, m$. Suppose that n_k forecasters select state S_k. Then vector $(n_1, n_2, \ldots, n_m)$, $n_1 + n_2 + \ldots + n_m = N$, represents a distribution of forecasts on the state space Ω. Let $p_k = n_k/N$, $k = 1, 2, \ldots, m$. The forecast poll induces a probability distribution $p \equiv (p_1, p_2, \ldots, p_m)$ on the state space Ω.

Second, assume that market prices are efficient. Then the prevailing spot rate aggregates all relevant information, and the consensus forecast of a large group of experts tends to the prevailing spot rate. Therefore, the relative individual forecast S has a zero mean, $\Sigma\, p_k\, S_k = 0$. This condition is consistent with empirical Observation 2.

Third, assume that the Oedipus effect is at work. Then the relative forecast S has a large variance $\Sigma\, p_k\, S_k^2 \le M^2$, where M^2 is a variance bound. This condition is consistent with empirical Observation 3.

Furthermore, by the Oedipus effect, the experts make state selections at random. Then modeling the forecast poll is identical to distributing N balls across m cells. (This statistical balls and cells model is reviewed in the appendix.) The maximum-likelihood distribution is determined by maximizing the likelihood function:

$$-\sum_{k=1}^{m} p_k \log p_k$$

subject to the normalization $\Sigma\, p_k = 1$, zero-mean, and the variance bound M^2. Motivated by the statistics in Exhibit 4, we set $M = 0.57\%$ for the forecasts of the 3-month interest rate and $M = 0.52\%$ for the forecasts of the 30-year interest rate. The solution is presented in the appendix.

Exhibit 5 plots empirical and theoretical maximum-likelihood distributions of the relative individual forecast

EXHIBIT 5

Empirical and Theoretical Distributions of Relative Individual Forecasts

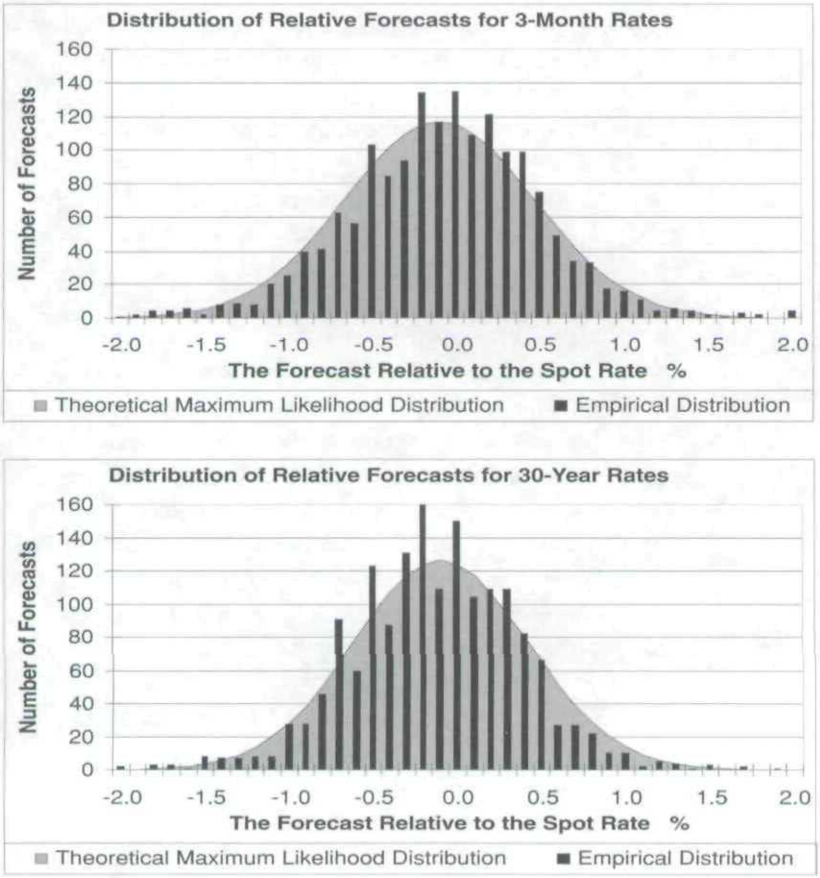

S. The empirical distributions are based on 1,646 forecasts of the 3-month interest rate and 1,649 forecasts of the 30-year interest rate published in *The Wall Street Journal* surveys. The theoretical distributions are qualitatively similar to the empirical ones.[3]

The maximum-likelihood model assumes that market prices are efficient and that investors hold heterogeneous expectations. It does not rely on individual rationality. To the contrary, the investors are represented by pinballs and therefore lack any intelligence or rationality. While economic models usually assume that investors are perfectly rational, the maximum-likelihood model goes to the other extreme, and effectively assumes that investors have zero intelligence. Farmer [2003] argues it is fruitful to explore both extremes.

It is generally accepted in the finance literature that the presence of rational traders is necessary in order to attain price efficiency. For example, Black [1986] proposes a market model with two types of traders. One type is irrational noise traders who help market liquidity but not price discovery. The second type is rational informed traders who carry out price discovery and maintain efficient pricing. By contrast, I suggest that efficient pricing can be achieved even when all traders are irrational.

INTERPRETATION

I interpret the likelihood function $H(\mathbf{p}) \equiv -\sum p_k \log p_k$ as an index of collective market uncertainty about the next price change. Efficient market prices keep investors in a state of maximum uncertainty about the next price change. If $H(\mathbf{p})$ indexes the collective uncertainty, then in efficient markets $H(\mathbf{p})$ must be at its maximum. Therefore, maximizing the likelihood function $H(\mathbf{p})$ amounts to seeking a probability belief that is consistent with the EMH.

In Exhibit 4, the median relative forecast is near zero, meaning that on average 50% of forecasters expect rates to rise and 50% expect rates to fall. A random variable with two outcomes attains maximum uncertainty when the two outcomes occur with probability 50/50. Therefore, on average, spot rates keep the experts in a state of maximum uncertainty about the next up or down shift. Also, the two-state likelihood function $p_1 \cdot \log p_1 + p_2 \cdot \log p_2$ reaches its maximum when $p_1 = p_2 = 0.5$.

The uncertainty index $H(\mathbf{p})$ is a map from the set of plausible probability distributions to the real line as shown in Exhibit 6. The plausible distributions should reflect all relevant consensus information *I*. The information set *I* is specified by the mean and variance constraints for the random relative forecast *S*. The maximum-likelihood probability distribution maximizes the collective uncertainty, conditional on the consensus information *I*.

The maximum-likelihood model represents a statistical filter for inferring the collective market beliefs from the consensus information set *I*. A different information set *I* results in a different maximum-likelihood probability belief. In particular, if the information set *I* includes market prices, then the model filters the collective expectations out of the prices.

The maximum-likelihood model is consistent with modern financial theory. In Gulko [1999], I show that the model entails the absence of arbitrage, the capital asset pricing model, the arbitrage pricing theory, and the Black-

EXHIBIT 6
Market Uncertainty Index

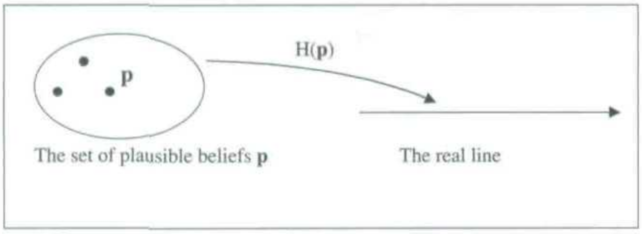

Scholes stock option pricing model without relying on complete markets, homogeneous beliefs, or rational utility-maximizing investors. Thus, the maximum-likelihood model extends modern valuation theory to incomplete markets inhabited by irrational investors with heterogeneous expectations.

A major practical implication of modern valuation theory is the superiority of passive index investing over active investing. If the theory is independent of investors' rationality, then passive index investing is a superior long-term strategy, whether or not investors act rationally. And active investment strategies, designed to profit from irrational investor behavior, might be a red herring. Malkiel [2003] supports this viewpoint with some empirical evidence.

I conclude by noting that the maximum-likelihood model captures the central themes of the capital market described in Hayek [1937, 1945]. First, the Hayek market represents a *spontaneous order* rather than equilibrium. The spontaneous order is a product of random and fuzzy interaction of numerous heterogeneous agents. The spontaneous order—as well as the entire market activity—is governed by random chance rather than design, by ambiguity rather than deterministic relations such as, for example, the equality of supply and demand:

> The concept of equilibrium itself and the methods which we employ in pure analysis have a clear meaning only when confined to the analysis of the action of a single person and ... we are really passing into a different sphere and silently introducing a new element of altogether different character when we apply it to the explanation of the interactions of a number of different individuals. (Hayek [1937, p. 34]).

Next, investors in the Hayek market have heterogeneous knowledge and heterogeneous expectations. No one has perfect information; instead, each one has only "bits of knowledge." Finally, investors exchange information using market prices as a communication medium. Market prices aggregate the bits of limited individual knowledge and disseminate the aggregate information to everyone. As a result, the price system enables many imperfect heterogeneous agents to communicate and coordinate their actions in an efficient manner:

> The whole [economy] acts as one market not because any of its members survey the whole field but because their limited individual fields of vision overlap so that through many intermediaries the relevant information is communicated to all.... [As a result, the market functions as] one single mind processing all the information which is in fact dispersed among all the people involved in the process.... We must look at the price system as such a mechanism for communicating information.... The most significant fact about this system is the economy of knowledge with which it operates, or how little the individual participants need to know in order to be able to take the right action. (Hayek [1945, p. 526])

If the market functions as "one single mind," then a suitable market model can be analogous to a neural network model of the human brain. Statistical methods used to study neural networks may also be useful for studying the random and spontaneous activity taking place in capital markets. The maximum-likelihood model of interest rate forecasts provides a concrete example.

CONCLUSION

The conventional wisdom suggests that the information efficiency of market prices is a product of individual rationality. My thesis is that investors do not need to be rational for market prices to be efficient. An empirical study of interest rate forecasts published in *The Wall Street Journal* finds support for both price efficiency and individual irrationality. Hence, price efficiency and individual irrationality are not necessarily mutually exclusive, and price efficiency is not necessarily a product of individual rationality.

A capital market inhabited by irrational investors can function efficiently if investors communicate and coordinate their actions. According to Hayek [1937], investors communicate and coordinate their decisions through market prices. According to Georgescu-Roegen [1971], communication and coordination through prices is subject to an Oedipus effect: Investors observe market prices, and then randomly revise their expectations. Therefore, investors hold diverse expectations even after they have observed a superior forecast.

I have presented a maximum-likelihood model of an efficient market populated by irrational investors. The model produces a distribution of interest rate forecasts that agrees with *The Wall Street Journal* polls. From the viewpoint of theory, the model represents a formal expression of the efficient markets hypothesis and extends modern valuation theory to markets with irrational investors. From a practical viewpoint, this means that passive indexing is superior to active investing, irrespective of investors' rationality, while active strategies, designed to profit from behavioral biases, might not measure up to their promise.

APPENDIX

Maximum-Likelihood Distribution

The objective is to find the maximum-likelihood distribution of N (distinguishable) balls across m cells, when N is much greater than n. This problem is discussed in many texts; for example, see Feller [1957, p. 38].

The number of (equally probable) ways to obtain distribution $(n_1, n_2, \ldots, n_m)$, where $n_1 + n_2 + \ldots + n_m = N$, is

$$W = \frac{N!}{n_1! \times n_2! \times \ldots \times n_m!}$$

The probability of observing distribution $(n_1, n_2, \ldots, n_m)$ is proportional to W. Therefore, the maximum-likelihood distribution maximizes W. The same distribution maximizes $\log(W)$ because $\log(W)$ monotonically increases with W. Next, $\log(W)$ can be simplified as follows:

$$\log(W) = -N \Sigma_k p_k \log p_k \text{ as } N \to \infty$$

where $p_k = n_k/N$, $k = 1, 2, \ldots, m$, and $\Sigma_k p_k = 1$. The proof of this equality uses Stirling's formula $n! = (2\pi)^{1/2} n^{n+1/2} e^{-n}$ as $n \to \infty$.

Therefore, when N is a large number, the maximum-likelihood distribution is determined by maximizing:

$$-\sum_{k=1}^{m} p_k \log p_k$$

subject to normalization, zero mean, and the variance bound M^2:

$$\sum_{k=1}^{m} p_k = 1$$

$$\sum_{k=1}^{m} p_k S_k = 0$$

$$\sum_{k=1}^{m} p_k S_k^2 \le M^2$$

where S_k, $k = 1, 2, \ldots, m$, are the states of the state space Ω. There is a unique maximum because $-\Sigma_k p_k \log p_k$ is a concave function on a convex set. Consider the Lagrangian:

$$L = -\sum_{k=1}^{m} p_k \log p_k + \lambda_0 \sum_{k=1}^{m} (p_k - 1) + \lambda_1 \sum_{k=1}^{m} p_k S_k + \lambda_2 (\sum_{k=1}^{m} p_k S_k^2 - M^2)$$

where λ_0, λ_1, and λ_2 are the Lagrange multipliers determined from the constraints.

From the first-order conditions:

$$\frac{\partial L}{\partial p_k} = -\log p_k - 1 + \lambda_0 + \lambda_1 S_k + \lambda_2 S_k^2 = 0$$

Therefore, the solution is

$$p_k^* = \frac{\exp(\lambda_1 S_k + \lambda_2 S_k^2)}{\sum_{k=1}^{m} \exp(\lambda_1 S_k + \lambda_2 S_k^2)} \qquad k = 1, 2, \ldots, m$$

If the variance constraint is not binding then $\lambda_2 = 0$ and the solution is a uniform distribution on Ω. In this case, the variance of S is $(1/m)\Sigma_k S_k^2 = (1.18\%)^2$.

Therefore, if $M \ge 1.18\%$ the variance constraint is not binding, and the solution is the uniform distribution on Ω. If $M < 1.18\%$, the variance constraint is binding, and the solution is a discrete normal distribution on Ω.

ENDNOTES

Kim Lambertson provided excellent research and administrative assistance. Suggestions of a referee improved the article, which also benefited from presentations at the Courant Institute at New York University, the Boston QWAFAFEW, the University of Connecticut, and the Global Finance Conference.

[1]The literature examining the consensus forecasts of economic indicators has often questioned the rationality of the forecasts. See Aggarwal, Mohanty, and Song [1995] and Schirm [2003].

[2]More generally, investors are said to be rational if and only if they fully use available information in order to assign probabilities to the future states of the world and then maximize the expected values of their von Neumann-Morgenstern

mize the expected values of their von Neumann-Morgenstern utility functions.

[3]Despite a good qualitative fit, the maximum-likelihood distributions fail the χ^2 test for goodness of fit. The fit can be improved by refining the state space Ω and by including additional information, such as higher moments of S.

REFERENCES

Aggarwal, R., S. Mohanty, and F. Song. "Are Survey Forecasts of Macroeconomic Variables Rational?" *Journal of Business*, 68(1) (January 1995), pp. 99-109.

Benartzi, Shlomo, and Richard H. Thaler. "Naïve Diversification Strategies in Defined Contribution Savings Plans." *American Economic Review*, 91(1) (2001), pp. 79-98.

Black, Fischer. "Noise." *Journal of Finance*, 41(3) (1986), pp. 529-543.

Fama, Eugene E. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance*, May 1970, pp. 383-417.

——. "Market Efficiency, Long-Term Returns, and Behavioral Finance." *Journal of Financial Economics*, 39 (September 1998), pp. 283-306.

Farmer, J. Doyne. "Explaining Statistical Properties of Markets via Random Low-intelligence Agents." Presented at the Mathematical Finance Seminar at the Courant Institute of Mathematical Science, New York University, January 2003.

Feller, William. *An Introduction to Probability Theory and its Applications*, vol. 1, 2nd ed. New York: Wiley, 1957.

Georgescu-Roegen, Nicholas. *The Entropy Law and the Economic Process*. Cambridge: Harvard University Press, 1971.

Gulko, Les. "The Entropic Market Hypothesis." *International Journal of Theoretical and Applied Finance*, 2(3) (1999), pp. 293-329.

Hayek, Friedrich A. "Economics and Knowledge." *Economica* NS 4 (February 1937), pp. 33-54.

——. "The Use of Knowledge in Society." *American Economic Review*, 35 (September 1945), pp. 519-530.

Kahnemann, Daniel, and Amos Tversky. "Prospect Theory: An Analysis of Decision under Risk." *Econometrica*, 47(2) (1979), pp. 263-291.

Malkiel, Burton G. "Passive Investment Strategies and Efficient Markets." *European Financial Management*, 9(1) (March 2003), pp. 1-10.

Muth, J. "Rational Expectations and the Theory of Price Movements." *Econometrica*, 29 (July 1961), pp. 315-335.

Odean, Terrance. "Are Investors Reluctant to Realize Their Losses?" *Journal of Finance*, 53(5) (1998), pp. 1775–1798.

Ross, Stephen A. "Finance." *The New Palgrave: A Dictionary of Economics*. New York-London: W.W. Norton, 1987.

Schirm, David C. "A Comparative Analysis of the Rationality of Consensus Forecasts of US Economic Indicators." *Journal of Business*, 76(4) (2003), pp. 547-561.

Shiller, Robert J. *Irrational Exuberance*. Princeton: Princeton University Press, 2000.

Shleifer, Andrei. *Inefficient Markets: An Introduction to Behavioral Finance*. Oxford: Oxford University Press, 2000.