Market "Efficiency" in a Market with Heterogeneous Information

Stephen Figlewski
New York University

It is commonly felt that a financial market achieves informational efficiency as traders with the best information and the most skill make profits at the expense of those with inferior information or ability and come to dominate the market. This paper develops a model of a speculative market in which this redistribution of wealth among traders with different information and ability can be studied. In the short run the market tends toward increased efficiency, but in neither the short nor the long run is full efficiency likely. The average deviation from efficiency is shown to depend on traders' characteristics such as the quality and diversity of their information and their risk aversion.

From the time of Adam Smith, economists have extolled the virtues of the competitive price system as a mechanism for allocating scarce resources among competing uses. Given the proper assumptions, the free operation of the competitive market can be shown to result in a Pareto-optimal allocation of goods. But for certain goods whose characteristics are not completely known, the market has the additional role of aggregating the available information about these characteristics. A share of IBM stock represents a claim on future earnings whose total value cannot be known in the present. The market price for IBM represents an aggregate opinion about the company's future prospects based on whatever information may be currently available to the participants in the market. And in general, if we define a speculative market broadly as a market for a good demanded not (entirely) for its own sake but for resale (or potential resale) in the future, the current price in a speculative market will always have at least a component which is the market's estimate of the future price. In most

I am indebted to Glenn Loury, Steven Sheffrin, Jerry Hausman, and an anonymous referee for many valuable suggestions and to Robert Solow for an extraordinarily insightful conjecture which proved to be the basis of this essay.

[*Journal of Political Economy*, 1978, vol. 86, no. 4]

cases, of course, speculative prices serve as signals which affect productive decisions in the rest of the economy. Equity prices affect firms' investment decisions, commodity futures prices determine storage and production, and so on. We would like to know how information is actually incorporated into a market price and especially whether the information-processing function of a competitive speculative market has optimality properties like those which apply to the allocative function. Is there an informational "invisible hand" which leads a competitive speculative market to make the optimal use of the information that society has available?

One answer to this question comes in the form of the "efficient-markets hypothesis." Under this hypothesis, a competitive speculative market is typically asserted to be "efficient" in the sense that the current market price always "fully reflects" all available information or the current price, plus normal profits, is the "best estimate" of the future price. Although there is clearly some kind of optimality property involved, such terms as "fully reflects" or "best estimate" are sufficiently imprecise that there is a wide latitude among economists on what "efficiency" should mean exactly. Fama (1970) distinguishes three degrees of market efficiency, "weak," "semistrong," and "strong," according to what type of information is fully reflected in the market price. A market is weakly efficient if the current price always completely discounts the information contained in the history of past market prices. The semistrong form of efficiency widens the scope to include all publicly available information. In addition to the history of past prices, the market accurately evaluates such things as dividend declarations, crop reports, and, we might expect, *Wall Street Journal* articles. Finally, the strong form of efficiency occurs when the market accurately discounts all information, including that held only by small numbers of market participants.

Obviously there is a large difference between weak and strong efficiency in terms of the optimality properties they imply for information processing in a decentralized market. The social value of a mechanism for aggregating information depends on its ability to generate price signals that accurately reflect all of society's information. Thus only a market that is efficient in the strong sense can really be said to have the optimality properties we would like. Throughout this paper, "efficiency" will be taken to mean strong efficiency.

If we try to imagine the mechanism by which a speculative market would achieve and maintain informational efficiency, we are led, like Cootner, to the following kind of story:

> Given the uncertainty of the real world, the many actual and virtual investors will have many, perhaps equally many, price forecasts. . . . If any group of investors was consistently better than average in forecasting stock price, they would accumulate

> wealth and give their forecasts greater and greater weight. In the process, they would bring the present price closer to the true value. Conversely, investors who were worse than average in forecasting ability would carry less and less weight. If this process worked well enough, the present price would reflect the best information about the future in the sense that the present price, plus normal profits, would be the best estimate of the future price. [1967, p. 80]

In this view market efficiency is a condition that is achieved in the long run as wealth is redistributed from investors with inferior information to those with better information. Of course in the short run, before this has a chance to happen, the distribution of wealth and the distribution of information quality may be very different. Since the market weights traders' information by "dollar votes," not quality, a trader with superior information but little wealth may have his information undervalued in the market price, and the market will be inefficient. However, there will exist some distribution of wealth—we might call it the efficient-market distribution—for which the dollar-vote weights are identical with information-quality weights and each trader's information is accurately reflected in the market price.[1] If the distribution of wealth ultimately converges to this efficient-market distribution, in the long run the market does become informationally efficient. The purpose of this paper is to develop a model of a speculative market in which the convergence can be analyzed. We find that in the short run the distribution of wealth tends to move toward the efficient-market distribution. A trader whose information was undervalued has an expected profit, and one whose information was overvalued has an expected loss. But in the longer run, random fluctuations in wealth resulting from the inability to forecast prices perfectly lead to deviations from the efficient-market wealth distribution and consequently from market efficiency. In general the market price is not the best estimate of the future price, given the currently available information. From the long run or "steady-state" distribution of wealth we can calculate the average efficiency of the market, as measured by the average variance of the market's forecast error, and compare it with the efficient-market variance. The difference can be thought of as the efficiency cost of processing information through a decentralized market rather than a centralized authority. A series of examples will give some idea of how this cost depends on the underlying parameters of the market such as the traders' risk aversion, the disparity in their forecasting abilities, and so on.

[1] The existence of the efficient-market wealth distribution undoubtedly requires some regularity conditions on traders' demands. In the model presented below, these are satisfied, and the efficient market distribution exists and is unique.

The question of the informational efficiency of a decentralized financial market has already been raised in a different manner in a series of recent papers by Grossman (1975, 1976) and Grossman and Stiglitz (1975, 1976). They show that when information is costly to obtain it cannot be true that the market price will accurately reflect all available information. If the market price revealed all information for free, it would not pay anyone to invest in information gathering individually. But if no one gathered information, there would be none for the market to reveal. Thus an informationally efficient market is incompatible with costly information. Grossman and Stiglitz's solution to this difficulty is to expand the traditional concept of market equilibrium to one of "informational equilibrium." In addition to the standard equilibrium condition that supply equals demand in every period, there is the further condition that the market price must reveal just enough of the costly information that participants are indifferent between becoming "informed" or remaining "uninformed." In full equilibrium, the market price does not reveal all the information, and traders who buy information do earn a higher return in the market. But the extra return is just sufficient to offset the cost of the information, and the expected return, including the cost of information, is equal for informed and uninformed traders.

An important feature of the Grossman and Stiglitz models is that the redistribution of wealth among bad and good forecasters which Cootner (1967) talks about is not a factor. In all of their models, investors are assumed to have constant absolute risk-aversion utility functions which have the property that the demand for a given risky asset is independent of wealth. Even if the good forecasters do accumulate wealth over time, this does not lead to a heavier weighting of their forecasts in the market price. Adjustment to informational equilibrium occurs not because of redistribution of wealth among traders but because of entry into and exit from the information-collection business. Thus deviation from market efficiency arising from wealth redistribution, which will be a principal feature of the model I analyze below, is over and above the informational inefficiency that Grossman and Stiglitz discuss.

The foregoing discussion has revolved around the question of how a competitive financial market processes information without considering specifically what "information" consists of. In the Grossman and Stiglitz framework, information is a datum which allows a trader to reduce his forecast variance of the next period price. If every trader possessed the information, they would have identical expectations about the future price. Clearly, this view of information is rather restrictive. We can easily think of things which not all market participants would consider even to be information at all. For example, the news that IBM had just completed a perfect "head and shoulders" top would be information to some, not to others. More generally, even if all traders accept a certain piece of information such as an earnings report as being important, they will not necessarily agree completely on its implications. It is not possible to separate the impact of elementary information such as news releases, crop reports, etc., from the subjective evaluation of this information by the participants in the market. Thus, rather than dealing with differences in "information," it will be more convenient to work with differences in forecasting ability—bearing in mind that access to elementary information is a major determinant of forecasting ability. The operational definition of an efficient market, then, is one in which the market price at any time (plus normal profits) is the best, that is, minimum variance, estimate of the future price, given the individual forecasts of all the market participants.

In the next section I develop a model of a speculative market with two types of traders who have differing information. I derive the stochastic difference equation which describes the redistribution of wealth among the two groups and analyze the market's short-run behavior. In the following section, I approximate the difference equation by a discrete Markov model and analyze its long-run properties. Although it is not possible to derive the steady-state wealth distribution analytically, several illustrative examples show how the market behaves with different values of the underlying parameters. This allows me to draw some tentative conclusions about how the informational efficiency of a decentralized market should be affected by heterogeneous information, differences in the quality of traders' information, risk aversion, and so on. The final section gives a summary and conclusion.

I. The Model

A competitive market weights a trader's information by the size of his investment, so a market's informational efficiency depends on the distribution of wealth among its participants. In this section I develop a model of a speculative market in which the interaction of information and wealth and the resulting effects on market efficiency can be analyzed.

The market is made up of equal numbers of two types of traders, a and b. All a traders are alike, as are all b traders, but members of the two groups may differ in price expectations, risk aversion, predictive ability, and wealth. The assumption of just two types of traders is made purely for expositional convenience. The model developed below can readily be extended to n traders with no change in the basic results. Although there are only two groups, we will assume that each trader views himself as trading in a perfectly competitive market. Otherwise we would have the problem that the a traders could solve back from the observed market price to obtain the b group's information and vice versa. In a market with more than two groups, it would not be possible to determine the information held by every other participant from the market price alone.

The market is for claims on an asset which pays a random return P_t^* at the end of each period. The net supply of claims is zero, so that the only way for one trader to buy a claim is for another to sell one short. All trading takes place between the two groups, since all members of a given group are identical. Two features of this setup lead to considerable simplification. First, P_t^* is independent of the operation of the market, so we have avoided the "beauty contest" problem which Keynes talked about with respect to the stock market. Second, this is a zero-sum game, so that the "normal profits" to a trader are zero, and also the analysis is not complicated by changes in the scale of the market over time. Both characteristics are fairly closely approximated by a typical futures market in which P_t^* represents the spot price in the contract month.

At the outset of each period, the two groups receive information about P_t^*. Next, the market opens and an equilibrium market price is achieved by a *tâtonnement* process. (Expectations may be revised at any point up to equilibrium, so that the market clearing price P_t is part of the information set upon which expectations are ultimately based.) At the equilibrium the demands of the a traders, n_t^a, are exactly offset by the (algebraic) demands of the b traders, n_t^b. Finally the market closes, P_t^* is revealed, and there is a net transfer of $n_t^a(P_t^* - P_t)$ from b to a traders.[2] [Of course $n_t^a(P_t^* - P_t)$ may be negative, so that b traders receive money from a traders.]

There will only be a wealth transfer when P_t^* differs from P_t, that is, when the market price is an inaccurate forecast of the future price. But the market's forecast error is just a combination of the traders' forecast errors, so the wealth redistribution in period t is a function of the traders' individual errors in forecasting P_t^*. We will derive expressions for n_t^a and $(P_t^* - P_t)$ in terms of the traders' characteristics such as forecasting ability and risk aversion and their random forecast errors. The latter drive the model, and their known distribution allows us to derive an equation for the stochastic process governing the redistribution of wealth within the market.

We now consider the expectations formation of the two groups. (In what follows the subscript t will be dropped for simplicity when it is not essential. There should be no confusion about what period the variables refer to.) No information about P^* is available before the beginning of the period, so traders come into the period with noninformative (flat) prior distribu-

[2] A question arises when a price change is so large that one trader cannot cover his losses. Treating the possibility of bankruptcy explicitly would greatly complicate the model. Instead, we will assume that in this market (as in many actual markets) trading takes place through a well-capitalized clearing corporation which insures all trades. Thus traders can transact without fear that their contracts will not be fulfilled. In any case, a trader's acceptable level of risk exposure depends on his risk aversion. In this model, if traders are sufficiently risk averse the probability of a bankruptcy becomes arbitrarily small.

tions on P^*. Before the market opens and during the *tâtonnement* process, they receive information sets Φ^a and Φ^b (including P, the market clearing price), so that by the time the market has closed they have formed and traded on the basis of the following posterior distributions on $P^*: f^a(P^*|\Phi^a)$, the a traders' posterior distribution, is normal with mean P^a and variance η_a^2; $f^b(P^*|\Phi^b)$ is normal with mean P^b and variance η_b^2. We will assume that the posterior variances, which measure the quality of traders' information and their ability to forecast from it, are constant over time and equal to the true forecasting variances. The (subjective) expected values of P^* conditional on each group's information are P^a and P^b. We assume that they are unbiased estimates of P^* over all realizations of P^*, Φ^a, and Φ^b. These posterior distributions imply that we can write the traders' prediction errors as

$$\begin{aligned} P^a - P^* &= \xi^a \qquad \xi^a \sim N(0, \eta_a^2), \\ P^b - P^* &= \xi^b \qquad \xi^b \sim N(0, \eta_b^2). \end{aligned} \tag{1}$$

Since ξ^a and ξ^b are not necessarily independent, it will be convenient to split each one into two independent parts, one of which is correlated with the other group's error and the other of which is not.

$$\begin{aligned} P^a - P^* &= \varepsilon^a + \delta \qquad \varepsilon^a \sim N(0, \sigma_a^2); \\ & \qquad\qquad\qquad \delta \sim N(0, \theta^2); \\ P^b - P^* &= \varepsilon^b + \delta \qquad \varepsilon^b \sim N(0, \sigma_b^2); \end{aligned} \tag{2}$$

ε^a, ε^b, δ are mutually independent.

A trader calculates his market demand by maximizing a two-parameter utility function defined on the expected value and variance of end of period wealth. Any portion of a trader's wealth not invested is held in a riskless asset earning zero return, and there are no limitations on borrowing. A trader's expected wealth depends on the expected price change, which in turn depends on his expected value for P^*.

The variance of wealth is determined by the variance of the change in the market price. We assume that traders estimate this variance by observing the market's operation over time. In each period they treat the variance of $(P^* - P)$ as being a constant ϕ^2 equal to the long-run average variance. Notice that ϕ^2 is a characteristic of the market, not of the individual. The risk involved in taking a position in this market is the same for all traders, even though they may have information of differing quality.

To derive his demand, n^a, an a trader solves the following maximization problem:

$$\max U^a[E(W_{t+1}^a), \operatorname{var}(W_{t+1}^a)]. \tag{3}$$

Calculating the expected value and variance conditional on a's information and current wealth, we have

$$\begin{aligned} E(W^a_{t+1}) &= E[W^a_t + n^a(P^* - P)], \\ &= W^a_t + n^a(P^a - P); \\ \text{var}\,(W^a_{t+1}) &= \text{var}\,[W^a_t + n^a(P^* - P)], \\ &= (n^a)^2\,\text{var}\,(P^* - P), \\ &= (n^a)^2\phi^2. \end{aligned}$$

Equation (3) becomes

$$\max_{n^a} U^a[W^a_t + n^a(P^a - P), (n^a)^2\phi^2],$$

and solving the first-order condition for n^a gives $n^a = (-U^a_1/2U^a_2)(P^a - P)/\phi^2$, where subscripts denote partial derivatives.

For a given expected value and variance of price change, the term $-U^a_1/2U^a_2$ determines a trader's desired investment. In order for traders to come to play a larger role in the market as their wealth increases, this term must be increasing in W. That is, the derivative with respect to W must be positive. For convenience, we will assume that traders will desire to invest a fixed fraction of their initial wealth in a given risky opportunity, regardless of the level of wealth. This assumption (which can be weakened considerably) implies that $d/dW(-U^a_1/2U^a_2) = 1/R^a$, where R^a is a positive constant. This gives $-U^a_1/2U^a_2 = W^a/R^a$ and

$$n^a = \frac{W^a}{R^a} \cdot \frac{(P^a - P)}{\phi^2}. \tag{4}$$

The constant R^a is a measure of risk aversion. The larger R^a is, the smaller will be a trader's demand for a given risky investment at any level of wealth. A similar calculation gives $n^b = W^b/R^b \cdot (P^b - P)/\phi^2$. At the market clearing price P, $n^a + n^b = 0$, so P solves $n^a + n^b = W^a/R^a \cdot (P^a - P)/\phi^2 + W^b/R^b \cdot (P^b - P)/\phi^2 = 0$. We will define $C = R^b/R^a$ and set $W^a + W^b = 1$ for convenience.

$$P = \left[W^aP^a + \frac{(1 - W^a)}{C}P^b\right] \Big/ \left(W^a + \frac{1 - W^a}{C}\right).$$

Let

$$X^a = W^a \Big/ \left(W^a + \frac{1 - W^a}{C}\right),$$

$$P = X^aP^a + (1 - X^a)P^b. \tag{5}$$

The market price will be a linear combination of the two groups' predictions with X^a and $(1 - X^a)$ as weights.

Let us now consider the price that an efficient market would produce. An efficient market should aggregate all the information into a price that is a sufficient statistic. That is, given the market price, there should be no further advantage to be derived from having the specific information of the individual market participants. In this case, the efficient-market price, P^{eff}, would be the mean of the posterior distribution on P^* of an omniscient observer who had access to both a and b traders' information.[3] Jaffee and Winkler (1976, p. 57) derive the posterior distribution for this case and find it to be a normal distribution with

$$\text{mean:} \qquad \frac{v(v - \rho)P^a + (1 - \rho v)P^b}{v^2 - 2\rho v + 1},$$

$$\text{variance:} \qquad \frac{v^2(1 - \rho^2)\eta_a^2}{v^2 - 2\rho v + 1},$$

where

$$v = \sqrt{\frac{\eta_b^2}{\eta_a^2}}, \qquad \rho = \frac{\theta^2}{\sqrt{\eta_a^2\eta_b^2}}.$$

If we substitute $(\theta^2 + \sigma_a^2)$ and $(\theta^2 + \sigma_b^2)$ for η_a^2 and η_b^2 and define $k = \sigma_b^2/\sigma_a^2$, these expressions can be simplified. We obtain

$$E[P^*|P^a, P^b] = \frac{k}{1 + k}P^a + \frac{1}{1 + k}P^b = P^{\text{eff}},$$
$$\text{var}\,[P^*|P^a, P^b] = \frac{k}{1 + k}\sigma_a^2 + \theta^2 = \text{var}\,[P^{\text{eff}} - P^*]. \qquad (6)$$

An efficient market should produce a market price which is a linear combination of the forecasts of the two types of traders where the weights vary inversely with the variance of the independent part of their forecasting errors.[4] Notice from (5) that, while the actual market price is also a

[3] We consider an omniscient observer who only deals with the posterior distributions of the traders. Since we have not ruled out the possibility of differences of opinion in evaluating a given set of elementary information, this tacitly involves one of two assumptions. Either the omniscient observer must always agree with a trader's evaluation of his information, or he must have access only to traders' posterior distributions and not their entire information sets.

[4] Another way to derive the efficient-market price is to treat this as a forecasting problem in which what is wanted is the combination of the two estimates of P^* which minimizes the squared forecast error. From normal theory we know the combination will be linear, and since both P^a and P^b are unbiased, their weights must sum to 1.

$$\min_\lambda E\{[\lambda P^a + (1 - \lambda)P^b - P^*]^2\},$$
$$= \min_\lambda E\{[\lambda\varepsilon^a + (1 - \lambda)\varepsilon^b + \delta]^2\},$$
$$= \min_\lambda [\lambda^2\sigma_a^2 + (1 - \lambda)^2\sigma_b^2 + \theta^2].$$

This yields the first-order condition $\lambda = \sigma_b^2/(\sigma_a^2 + \sigma_b^2)$, from which eq. (6) follows directly.

linear combination of P^a and P^b, the weights are not the same as in the efficient-market price. The forecast variance of the actual price will be

$$\begin{aligned} \text{var}\,[P^*|P] &= E[(P^* - P)^2], \\ &= E\{[-X^a\varepsilon^a - (1 - X^a)\varepsilon^b - \delta]^2\}, \\ &= (X^a)^2\sigma_a^2 + (1 - X^a)^2\sigma_b^2 + \theta^2, \end{aligned} \tag{7}$$

which is necessarily $\geq k/(1 + k)\sigma_a^2 + \theta^2$, since the latter is the minimum variance, given the market's information.

With the results above and two more intermediate steps, we can write the equation for the stochastic process describing the redistribution of wealth between a and b traders over time. First we calculate

$$\begin{aligned} P^a - P &= P^a - [X^aP^a + (1 - X^a)P^b], \\ &= (1 - X^a)(P^a - P^b), \\ &= (1 - X^a)(\varepsilon^a - \varepsilon^b) \end{aligned} \tag{8}$$

and

$$\begin{aligned} P^* - P &= P^* - P^a + P^a - P, \\ &= -\varepsilon^a - \delta + (1 - X^a)(\varepsilon^a - \varepsilon^b), \\ &= -X^a\varepsilon^a - (1 - X^a)\varepsilon^b - \delta. \end{aligned} \tag{9}$$

We can now write the equation for W_{t+1}^a. In this model, if we know W^a we also know $(1 - W^a) = W^b$, so that the value of W^a completely determines the distribution of wealth in the market. We have

$$\begin{aligned} W_{t+1}^a &= W_t^a + n_t^a(P_t^* - P_t), \\ &= W_t^a + \frac{W_t^a}{R^a\phi^2}(P_t^a - P_t)(P_t^* - P_t) \end{aligned}$$

from equation (4). Substituting from (8) and (9), this is

$$W_{t+1}^a = W_t^a - \frac{W_t^a}{R^a\phi^2}(1 - X_t^a)(\varepsilon_t^a - \varepsilon_t^b)[X_t^a\varepsilon_t^a + (1 - X_t^a)\varepsilon_t^b + \delta_t]. \tag{10}$$

To analyze the short-run behavior of the process, we will look at the expected value of W_{t+1}^a; the distribution of wealth in the next period given the current distribution W_t^a; and the characteristics of the market participants, R^a, R^b, σ_a^2, σ_b^2, and ϕ^2. Since ε^a, ε^b, and δ are mutually independent, in taking the expected value all of the cross terms are zero.

$$E[W_{t+1}^a|W_t^a] = W_t^a - \frac{W_t^a}{R^a\phi^2}(1 - X_t^a)[X_t^a\sigma_a^2 - (1 - X_t^a)\sigma_b^2]. \tag{11}$$

A short-run equilibrium point occurs when $E[W^a_{t+1}|W^a_t] = W^a_t$, that is, when the second term on the right side of (11) is zero.[5] This happens when

$$W^a_t = 0 \tag{12a}$$

or

$$X^a_t = 1, \text{ implying } W^a_t = 1,\ W^b_t = 0, \tag{12b}$$

or

$$X^a_t\sigma^2_a - (1 - X^a_t)\sigma^2_b = 0, \text{ that is, } X^a_t = k(1 - X^a_t). \tag{12c}$$

There are three equilibria to this process. If one group's wealth goes to zero, it drops out of the market permanently, so two of the three correspond to elimination of one type of trader. These two points are "trapping states" of the Markov process given in (10). Once they are entered, they are never left. The third equilibrium point corresponds to a distribution of wealth at which the weight the market gives P^a in forming the market price is exactly k times the weight on P^b. Comparing this with (6), we see that with those weights the market is efficient. Further, it is clear from (11) that $E[W^a_{t+1} - W^a_t|W^a_t] \gtreqless 0$ as $X^a_t/(1 - X^a_t) \lesseqgtr k$. If the market's weight on P^a is higher (lower) than the efficient market weight, a traders will have an expected loss (gain) in the next period. Thus the equilibrium point in (12c) is an attracting one. (Although we might be tempted to call it a stable equilibrium, we have not ruled out parameter values that would lead to explosive oscillations around it.)

One interesting feature of this result is that traders with inferior information do not get driven out of the market by the better forecasters. An inferior forecaster tends to lose money only as long as his independent information is overvalued in the market price. Once his wealth drops below the efficient-market level, his information becomes undervalued, and he begins to recoup some of his losses. This market is not a "game against nature," in which traders face an unchanging set of probabilities. Since the market weights expectations by wealth, as a trader loses wealth he receives increasingly favorable odds from the other traders. The more a trader has lost in the past, the more likely it becomes that he will win in the next round.

The equilibrium wealth distribution $W^a_{\mathbf{eq}}$ can be obtained by plugging the definition of X^a into (12c).

$$W^a_{\mathbf{eq}} = 1\bigg/\left(1 + \frac{C}{k}\right). \tag{13}$$

[5] This concept of short-run equilibrium is only relevant to the system's behavior in the next period. It does not say anything about long-run equilibrium wealth, a term which more properly refers to the ergodic distribution of W^a.

From (13) we see that $\partial W^a_{eq}/\partial k > 0$. If b traders become relatively poorer forecasters, the a traders' equilibrium wealth increases, not surprisingly. It is a little more curious that $\partial W^a_{eq}/\partial C < 0$. If b traders become more risk averse relative to the a traders, the equilibrium point shifts in their favor. After reflection, however, it is not hard to see why this happens. When b traders are more risk averse, they invest a smaller proportion of their wealth in the market. In order for the weight on P^b to be $1/k$ times the weight on P^a, b traders must have a higher fraction of the total wealth than before.

It should be noted that W^a_{eq} depends not on R^a or σ^2_a, the actual values of risk aversion and independent prediction variance, but only on C and k, the ratios of these variables between groups. It is only relative forecasting ability and risk aversion that count. Also θ^2, the variance of the joint prediction error, does not enter.

We see then that the market we have set up behaves as follows. If either group's wealth becomes zero (or negative), they drop out of the market permanently—meaning that the market ceases to exist, since members of the remaining group will never trade among themselves. When both groups have positive wealth, the distribution of wealth tends in expected value toward that distribution at which the market is an efficient market and the market price is the best estimate of P^*, given the available information. However, this distribution is only a stochastic equilibrium. Because of the randomness of the price movements, the actual distribution of wealth at any time may deviate relatively far from the efficient-market distribution. The market, although it *tends* toward efficiency, will actually *be* efficient only on a set of measure zero. The next question is how far the market deviates from efficiency on average. The answer will give us an idea of the informational efficiency loss of a decentralized market compared with a centralized information-processing authority. We will deal with the long-run behavior of the market in the next section.

II. Market Efficiency in the Long Run

In order to analyze the long-run behavior of the model developed in the previous section, we need to know the steady-state distribution of wealth, which we might write as

$$F^{\infty}(W^a) = \lim_{n \to \infty} F(W^a_{t+n} | W^a_t), \tag{14}$$

where $F(\cdot)$ is the probability density function of W^a. Unfortunately, for the model in equation (10) the analysis is complicated by the existence of the trapping states. Once entered, these states are never left. In the model there are two trapping states, $W^a = 0$ and $W^a = 1$, since if either of the values occurs one of the types of traders becomes bankrupt and the market ceases to operate. By increasing the traders' risk aversion, the probability of bankruptcy in any period or indeed in any finite number of periods can be made arbitrarily small. However, since disturbances are normally distributed, there will always be a positive probability of entering one of the trapping states from any other state, so in the longest run the probability that $W^a = 0$ or $W^a = 1$ is 1, not a very interesting result.

The easiest way to deal with this problem is to change the model slightly to allow the traders to receive an income of y every period. Bankruptcy is no longer permanent in this case, because each trader has an addition to his wealth from outside the market every period. This removes the trapping states. The required changes to the model developed above are relatively minor. Since it is no longer possible to set $W^a + W^b = 1$ for all t, we have $P = [W^aP^a + (W^b/C)P^b]/(W^a + W^b/C)$. Letting

$$X^a = W^a/(W^a + W^b/C) = (W^a/W)/[W^a/W + (1 - W^a/W)/C],$$

where $W = W^a + W^b$, we again have equation (5).

Since we are solving for the fraction of total market wealth held by the a traders, equation (10) must also be altered to take account of the income term. It becomes

$$\frac{W^a_{t+1}}{W_{t+1}} = \frac{1}{W_t + 2y} \left\{W^a_t + y - \frac{W^a_t}{R^a\phi^2}(1 - X^a_t)(\varepsilon^a_t - \varepsilon^b_t)\,[X^a_t\varepsilon^a_t + (1 - X^a_t)\varepsilon^b_t + \delta_t]\right\}. \tag{10'}$$

The next step is to derive the steady-state or ergodic distribution of W^a/W from (10′). Unfortunately, (10′) as it stands cannot be solved analytically to give a "nice" form for $F^\infty(\cdot)$. However, by using a discrete approximation to (10′) it is possible to recast the problem as a standard Markov model with a discrete-state space for which long-run transition probabilities can be easily obtained by numerical methods.

We divide the range of possible values of W^a/W into 21 cells from 0 to 1 inclusive and calculate the transition probability matrix A, where A_{ij} gives the probability that W^a_{t+1}/W_{t+1} will fall in cell j, given that W^a_t/W_t is equal to the midpoint of cell i. To get these probabilities, we approximate the joint density function of ε^a, ε^b, and δ with a discrete distribution in which each variable is allowed to take one of 21 possible values and solve (10′) for each of the resulting $(21)^3$ combinations. Now A_{ij} is the sum of the probabilities associated with each combination of ε^a, ε^b, and δ that make W^a_{t+1}/W_{t+1} fall in cell j when W^a_t/W_t is in cell i. The income term y is taken to be the minimum amount necessary to place a trader in the lowest cell in the period following a bankruptcy.

The long-run distribution of W^a/W can be obtained from A in two ways, either as a row of

$$A^{\infty} = \lim_{n \to \infty} A^n$$

or as the solution to the eigenvector problem $A'\omega = \omega$. Since this distribution is conditional on the initial guess for ϕ^2, the prediction variance of the market price, we then repeat the process with a new estimate of ϕ^2 derived from the calculated ergodic distribution. (In all cases ϕ^2 converges to two-decimal-place accuracy within two iterations.) Table 1 gives statistics on the steady-state distribution of wealth and market efficiency for several choices of market parameters in the model of (10′).

In table 1, run 1 is the basic run with which we will compare the results for other parameter values. The risk-aversion parameter R^a is set equal to 1; a and b traders are assumed to have equally good information and to be equally risk averse, so we set $k = 1$ and $C = 1$; and with $\sigma_a^2 = 1$ and $\theta^2 = 1$, a trader's prediction variance is made up of two equal parts, one of which is due to mistakes he shares with all other participants in the market and the other of which is specific to his own group. Another way of expressing this is to say that the correlation coefficient between the two groups' forecasting errors is .5.

With these parameter values, the mean of the long-run distribution of wealth in this market is 0.50. As we should expect, when traders are equally good forecasters and equally risk averse, on average each group will tend to have half the wealth. However, the standard deviation of the distribution is fairly high. For much of the time, the actual distribution of wealth between the two types of traders will in fact be far from equal.

The next line, "efficient-market" variance, shows the forecasting variance for the optimal combination of P^a and P^b. This figure is calculated from equation (6), and it is the forecasting variance that could be achieved by a centralized information-processing authority who was able to poll the market participants individually to determine their price expectations.

The prediction variance of the actual market price as calculated in equation (7) is a function (through X^a) of the distribution of wealth in the market. We use the long-run wealth distribution to derive the expected actual market variance. (This is ϕ^2 in eq. [10′].) For run 1 we see that the competitive market price has a forecasting variance on average 10.6 percent above the informationally efficient forecast.

In run 2 we see what happens in a market where traders are more risk averse. Here R^a is set equal to 2. Since the traders are still equally good forecasters and equally risk averse, wealth is still equally distributed on average. Further, because the increase in risk aversion leads traders to take smaller positions and expose themselves to fewer large gains and losses, the standard deviation of W^a/W has diminished considerably. The distribution of wealth will be in general much closer to equality than in the

TABLE 1

LONG-RUN MARKET BEHAVIOR UNDER DIFFERENT ASSUMPTIONS ABOUT TRADERS' CHARACTERISTICS

	Run Number						
	1	2	3	4	5	6	7
R^a	1	2	1	1	1	1	1
σ_a^2	1	1	1	1	1	1	1
k	1	1	2	4	1	1	1
C	1	1	1	1	2	1	1
θ^2	1	1	1	1	1	5	0
Mean W^a/W	.50	.50	.64	.74	.40	.50	.50
SD W^a/W	.29	.22	.27	.22	.24	.26	.29
"Efficient-market" variance	1.50	1.50	1.67	1.80	1.50	5.50	.50
Actual market variance	1.66	1.59	1.87	2.05	1.62	5.63	.67
Increase in variance (%)	10.6	6.2	12.4	14.1	8.2	2.3	33.6

previous case. This fact results in an increase in the efficiency of the decentralized market and leads us to expect that a market in which the traders are relatively unwilling to expose themselves to risk will be more informationally efficient, other things being equal, than one whose participants are not so risk averse.

In runs 3 and 4 we see what happens when one group of traders has better information than the other. In run 3, the *b* traders' independent forecasting error has a variance twice as large as that of the *a* traders, and in run 4 it is four times as large. Their overall forecasting variance is therefore 1.5 and 2.5 times larger, respectively. As we expect, when *b* traders become poorer forecasters, the distribution of wealth shifts toward the *a* group. It is also the case that the market becomes less efficient. By the time that σ_b^2 is four times σ_a^2, the market's prediction variance is 2.05—the market makes less efficient use of information than the *a* traders acting alone! Thus we may expect markets in which information is distributed among equally good forecasters to be relatively more efficient than those in which there is a wide disparity in forecasting ability, to the point that long-run market efficiency may be actually diminished by the participation of traders whose information is distinctly inferior.

Run 5 shows what happens when traders differ in their risk aversion. We see that the distribution of wealth shifts in favor of the more risk-averse group. Market efficiency also increases, since in a sense the market's average risk aversion has increased.

Finally, in runs 6 and 7 we examine the effect of changing the relative importance of the shared part of the traders' overall forecast error. The relative magnitudes of θ^2, σ_a^2, and σ_b^2 measure the extent to which traders have heterogeneous information. When θ^2 is small relative to σ_a^2 and σ_b^2, traders' information is practically independent, and prediction variance can be significantly reduced by taking the proper combination. Conversely, when θ^2 is large, most of a trader's forecast error comes from informational deficiencies which he shares with all of the other traders. No combination of forecasts can reduce this element. In the limit, when σ_a^2 and σ_b^2 are 0 and the entire forecast error is shared, traders have identical information, and any individual forecast or combination of forecasts will be the same.

The θ^2 for run 6 is 5, meaning that the correlation between a and b traders' forecast errors is .86. In this case, the market's variance is only 2.3 percent greater than the minimum possible. When $\theta^2 = 0$ and the traders' errors are completely uncorrelated, the market is much less efficient. Its variance is then 33.6 percent above the minimum. This indicates that we should expect a market where traders have diverse information to achieve significantly less efficiency than would be possible if information were processed through a central authority. Conversely, a market should be close to its theoretically possible efficiency when traders have similar expectations and the differences between their forecasts are small relative to the total forecast error.

The results above are all based on an approximation to equation (10′). In order to determine to what extent approximation error may be playing a role, run 3 was redone, using a much finer approximation. The wealth distribution was divided into 41 cells instead of 21, and a 39 × 39 × 39 approximation to the joint distribution of ε^a, ε^b, and δ was used. The mean of W^a/W and standard deviation each changed only by about 0.01, and the estimated actual market variance was identical to three decimal places. Thus we may assume that the approximation used in constructing table 1 was sufficiently fine to capture the true behavior of the system (10′) quite accurately.

III. Conclusion

Discussions of the efficient-markets model seldom specify precisely how the market processes information to produce a price that accurately discounts it. In general the rationale for strongly efficient markets seems to be some kind of weeding-out process over time by which those with inferior information gradually lose money to those whose information is better, with the result that in the long run the market puts the heaviest weight on the best information in forming a market price. In this paper I have examined market efficiency in the context of a simple model in which the market's treatment of diverse information can be seen explicitly. Because the market weights traders' information not by its quality but by "dollar votes," it was shown that neither in the short nor the long run was the market likely to be efficient, in the sense that the market price was the best estimate of the future price, given the available information. In the short run the distribution of wealth tended to move toward the distribution at which the market was efficient, but that precise distribution had zero probability of actually occurring. In the long run, the stochastic nature of price movements played a role, so that for some values of the underlying parameters the distribution of wealth had a high probability of occasionally straying far away from the efficient market value.

The results of this paper depend on two necessary characteristics. These are heterogeneous expectations and a dependence of demands on the level of wealth. Samuelson (1972) proves market efficiency when traders have homogeneous expectations, and Grossman (1976) shows that, without wealth effects on demand, even when traders have different information in the long run the market price will discount all of the information. However, when both features are present, strong market efficiency no longer holds. The results of Section II lead to some tentative conclusions about how the degree of market efficiency should be affected by the underlying characteristics of the market participants. The more risk averse the traders are and the more homogeneous their information, the more efficient we expect the market to be. However, when there is a wide range of forecasting ability or a diversity of expectations among the participants, the market may deviate relatively far from efficiency.

References

Cootner, Paul H., ed. *The Random Character of Stock Market Prices*. Rev. ed. Cambridge, Mass.: M.I.T. Press, 1967.

Fama, Eugene F. "Efficient Capital Markets: A Review of Theory and Empirical Work." *J. Finance* 25 (May 1970): 383–417.

Grossman, Sanford. "The Existence of Futures Markets, Noisy Rational Expectations and Informational Externalities." Stanford Technical Report no. 182, Stanford Univ., September 1975.

———. "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information." *J. Finance* 31 (May 1976): 573–85.

Grossman, Sanford, and Stiglitz, Joseph. "On the Impossibility of Informationally Efficient Markets." Paper presented at the Econometric Society meeting in Dallas, 1975.

———. "Information and Competitive Price Systems." *A.E.R. Papers and Proc.* 66 (May 1976): 246–53.

Jaffe, J. F., and Winkler, R. L. "Optimal Speculation against an Efficient Market." *J. Finance* 31 (March 1976): 49–61.

Samuelson, Paul A. "Proof That Properly Anticipated Prices Fluctuate Randomly." In *The Collected Scientific Papers of Paul A. Samuelson*, edited by Robert C. Merton. Vol. 3. Cambridge, Mass.: M.I.T. Press, 1972.