THE JOURNAL OF FINANCE • VOL. XXXVII, NO. 1 • MARCH 1982

Information Diversity and Market Behavior

STEPHEN FIGLEWSKI*

ABSTRACT

The paper addresses two major issues raised by information diversity in a speculative market. First, we analyze what property of an investor's information leads to an expected speculative profit and show that independence is more important than accuracy. Second, we consider whether the market price must become fully efficient, in the sense that every investor's information is accurately discounted, when traders use it rationally as an information source. We prove that for any information structure there is a unique equilibrium weighting of investor beliefs at which the price is fully efficient and also every trader's expected profit is zero. Except for special structures, however, this equilibrium need not be attained in finite time.

MUCH OF MODERN CAPITAL market theory depends heavily on the assumption that investors know the expected values, variances, and covariances of asset returns. Although the assumption of homogeneous expectations is a conventional one and necessary for many theoretical results, it is not one which can be offered with much confidence. Real world investors have access to differing information, and may also disagree on the correct interpretation of widely held information. When asset returns are assumed to be governed by a known probability distribution, many features of trading in real-world markets are eliminated from consideration. There is no role for information gathering and no way for an investor to profit from having superior information. Indeed, this term has no real meaning since all relevant information is contained in the known returns distribution. Speculative trading is also hard to explain without divergence of opinion. A speculator who thought the market price impounded all available information would never trade against it.

With heterogeneous information we lose some of the neat results of standard capital market theory, but new questions can be addressed about how information diversity among investors affects the markets in which they participate. One issue concerns the impact of informational differences on investor profits. Investors who are better informed than others or are better able to evaluate widely held information should be able to earn higher returns. All effort devoted to security analysis and market research depends on the belief that this is so. But what makes one investor's information "superior" to another's, beyond the circular observation that those investors with consistent profits must be better informed than those with losses? Do investors with inferior information, however measured, continually lose money until they are squeezed out of the market, or can an

* Associate Professor of Finance, New York University Graduate School of Business, New York, New York 10006. I would like to thank Vijay Bawa, Roy Radner, Kenneth Garbade, and Michael Brennan for helpful discussions and suggestions.

investor with low quality information that is not available to others, and therefore not discounted by the market, use it to make a profit?

Turning to the issue of market efficiency, a competitive market price reflects the demands of all investors, and each investor's demand function is based on his subjective evaluation of the information available to him. When the market aggregates investors' diverse beliefs into a single price, that price embodies all information currently available, to some degree. If the market does this "efficiently," arriving at a price that fully discounts all of the information, there are strong implications about the social optimality of competitive financial markets, as well as about appropriate investor behavior in these markets.

Market efficiency has been a central focus of much empirical as well as theoretical work in finance. The accumulated evidence from efficiency tests seems to show that widely held information, such as that contained in the history of past prices and public announcements of events like stock splits and antitrust suits, is fully discounted in the current market price. But studies of investors who are likely to have access to relevant information that is not widely available show that they are able to make excess returns and that markets do not fully reflect *all* information. In Fama's [2] terminology, markets are efficient in the "semi-strong" but not necessarily in the "strong" sense. Information that is not held by all investors, and therefore leads to heterogeneous beliefs among them, can sometimes be used to produce speculative profits.

Theoretical models based on the assumption of divergent expectations frequently encounter the problem that the market becomes too efficient. Since a market aggregates and reflects investors' information, the market price itself becomes an important source of information, revealing something about other investors' beliefs. In a fully efficient market, the price reveals *all* information. When rational investors recognize this fact and use the market price in place of their own information in making decisions, there are no longer any effective information differences among them and the model collapses to homogeneous expectations. The model treated by Williams [9] is an example of this phenomenon. Asset prices are assumed to follow a diffusion process with constant, but unknown mean. Investors begin with heterogeneous prior beliefs about mean returns, but over time the behavior of the market price reveals the true mean with increasing precision, and the posterior distributions converge to the homogeneous expectations solution where the mean is known.

Grossman [4, 5] shows that under general conditions a speculative market price can become a "sufficient statistic" for the union of all investors' information sets. When the price is a sufficient statistic, the probability distribution of security returns conditional on the market price alone is the same as the probability distribution conditional on all available information. The price will reflect all relevant information, despite initial differences among investors.

Fully revealing speculative market prices have some inconvenient implications. There is again no return to superior information. This creates a paradox when information is costly to obtain. If the market reveals all information for free, there is no reward for an investor to engage in information production. But if information is not produced, the market cannot reveal it. For traders to earn a return on resources devoted to producing information, the market must not reveal all of their information.[1] Even without costly information, no rational investor will speculate against a market which reflects all information perfectly. In a purely speculative market, trading disappears and the market breaks down.

Grossman's existence theorem is an important theoretical result in that it shows that a well specified heterogeneous expectations market model will normally possess a fully revealing informational equilibrium. However, the implications of this result are less strong than they might appear at first glance. While a fully revealing equilibrium *exists*, the important question is whether the market *achieves* it. If trading starts at a point where the market is less than fully efficient, will convergence to full information efficiency necessarily occur, and will it occur in the short run or only asymptotically? If convergence does not take place immediately, how does a speculative market behave when it is not in full informational equilibrium and there are speculative profits available for traders with better information?

This paper examines these issues in the context of a general model of a speculative market in which investors are assumed to have regular access to diverse information. In the next section we develop and analyze the behavior of a short run model in which the market price is not fully revealing. The main results of the section are summarized in two theorems. Theorem 1 shows that an investor has an expected profit from trading on his information when the covariance between his error in forecasting the future price and the market's error is less than the variance of the market's error. Trading profits do not require accurate information so much as information that is uncorrelated with that included in the market price.

We then consider the market's aggregation of information. The market weights an individual's expectations by the size of his investment. This depends on risk aversion, perceptions of the level of risk, and the investor's confidence in his prior information. Theorem 2 reveals the connection between the market's weighting of investors' information, trading profits, and information efficiency. There is a unique vector of weights that depends only on the structure of information differences in the investor population, for which expected profits are zero for all investors. This is the only point at which the market is fully efficient.[2]

In Section II we discuss ways in which the market might converge to a permanent state of full efficiency. Over the short run, the market will only be fully efficient in special cases. For example, if the market population consists of many investors, all of whom have independent information, the market approaches maximum efficiency in every period, regardless of its weighting of individual expectations.

We then consider mechanisms by which convergence to fully revealing prices

[1] Grossman discusses this problem and suggests that one way the market price might fail to be a sufficient statistic is if there is noise unrelated to information differences entering the price determination, for example, if asset supplies were uncertain. In a paper directly addressing this problem, Grossman and Stiglitz (6) develop a model with costly information in which the price is not fully revealing and long run equilibrium involves both equality of asset supply and demand, and entry and exit from the information production business. Hellwig [7] and Diamond and Verrecchia [1] also address this issue.

[2] However, unlike the Grossman equilibrium concept, full efficiency need not be self sustaining at this point.

might occur over time for general information structures. In models in which the market evolves because of learning by investors, and in models based on wealth redistribution from traders with "bad" information to those with "good" information, convergence will be a function of the stochastic behavior of the market price. Full efficiency occurs at best in the long run, and not at all in some cases.

The last section contains concluding comments.

I. The Model

In order to focus specifically on information differences, we shall construct a model of a purely speculative market, in which speculation on price changes is the only motive for trading, and the return for investors as a group is zero, even though each trader possesses information which he believes will allow him personally to make a profit. Most real-world futures markets fit this description fairly well.

The market operates in many periods. It is actually a sequence of markets in which trading takes place at the beginning of each period in one-period futures contracts calling for delivery of one unit of the underlying good at the end of the period. The difference between the beginning of period futures market price and the end of period spot price determines trading profits. In this section we analyze the market's behavior within a single period. For simplicity the subscript t will be left out. Within one period we can take as exogenous certain parameters, such as an investor's evaluation of the quality of his information, which may actually be a function of the market's behavior in earlier periods. These parameters become endogenous in a multiperiod framework. Discussion of the model's asymptotic properties is deferred to the next section.

At the beginning of each period investors receive information about what the end of period spot price P^* will be. The information sets $\{\Phi_i\}$ are produced by a stochastic information generating process which possesses a certain degree of regularity. For example, one trader may regularly receive more accurate information than another, or two investors may get very similar information.

Investor i's information Φ_i reveals the true end of period price up to a random error. His forecast of P^* given his information is P_i, which can be written

$$P_i = P^* + \epsilon_i \tag{1}$$

The forecast error ϵ_i is assumed to have mean zero, to rule out persistent expectations biases which the investor should be able to detect and eliminate over time, and to be normally distributed. The variance of ϵ_i is a natural measure of how good trader i's information is. We define V_i^2 to be trader i's estimate of his error variance, leaving open the possibility that his estimate may not be equal to the true variance.

In developing the model it is convenient to focus on forecast *errors* rather than on the forecasts themselves. The regular information differences among investors will be defined in terms of the stochastic structure of their errors in estimating P^*. Let $\underset{\sim}{P}$ and $\underset{\sim}{\epsilon}$ be N-vectors of the P_i and ϵ_i and $\underset{\sim}{e}$ be a vector of 1's. N, the number of traders, is assumed to be large. Throughout the paper, vectors will be denoted by an underline. Then,

$$\underline{P} - P^*\underline{e} = \underline{\epsilon} \qquad \underline{\epsilon} \sim N(\underline{0}, \Omega) \tag{2}$$

The vector of forecast errors is multivariate normally distributed with mean vector $\underline{0}$ and covariance matrix Ω. We assume Ω has full rank, implying that no investor's forecast can be consistently expressed as a linear combination of the others. Forecast variance is a measure of information quality—better information leads to a smaller variance. The covariance between two traders' errors measures information diversity, since the errors made by investors with nearly homogeneous information will be highly correlated.

After the information is released, the market opens and a market clearing price is established by tâtonnement, so that investors know what the market price P_M will be before they trade, and thereby obtain additional information about other traders' estimates of P^*. They use this information to update their expectations and all trading is ultimately based on forecasts conditioned on both Φ_i and P_M.

Investor i knows that the market price reflects information held by other investors which he does not possess, but he also believes he has information that P_M does not accurately reflect which will allow him to earn a speculative profit trading against the market. There is also public information that is reflected in both P_i and P_M, as well as "unavailable information," i.e., residual uncertainty which neither contains. This implies that the individual's error $(P_i - P^*)$ will be correlated with the market's error $(P_M - P^*)$.

Jaffe and Winkler [8] have derived the investor's optimal posterior forecast for the case where his error and the market's error are jointly normal and nonindependent. Investor i's posterior distribution conditional on both Φ_i and P_M is normal, with mean ρ_i and variance σ_i^2, where

$$\rho_i = \alpha_i P_i + (1 - \alpha_i)P_M \tag{3}$$

After observing the market price, an investor's posterior forecast is a linear combination of his prior forecast and P_M. Both α_i, the weight he places on his prior, and σ_i^2 are functions of trader i's estimates of his own error variance V_i^2, the market's error variance σ_M^2, and the covariance between the two errors C_{iM}.[3] Again we allow the possibility that these estimates may be inaccurate, e.g., trader i might think his covariance with the market is lower than it really is. While α_i is a fixed parameter within a given period, in a multiperiod framework, it will vary over time as investors observe the behavior of prices and revise their estimates of V_i^2, σ_M^2, and C_{iM}.

Next, we consider the individual trader's investment decision. Investors calculate their demands by maximizing utility functions defined on the mean and

[3] The expressions for these parameters can be written

$$\alpha_i = \frac{\sigma_M^2 - C_{iM}}{\sigma_M^2 + V_i^2 - 2C_{iM}}$$

$$\sigma_i^2 = \frac{\sigma_M^2 V_i^2 - C_{iM}^2}{\sigma_M^2 + V_i^2 - 2C_{iM}}$$

See Jaffe and Winkler [8, Appendix 1].

variance of end of period wealth.[4] No initial investment is required to enter into a futures contract; wealth only changes hands after P^* is known at the end of the period. This is consistent with practice in real futures markets, where the initial "margin" is really only a security deposit held by the exchange and required equally from both buyer and seller. Let W_{i0} and W_{i1} be the ith investor's beginning and end of period wealth, and n_i be the number of futures contracts purchased. We allow n_i to take on any real value, with negative values corresponding to short sales.

A trader determines his demand by solving the problem.

$$\text{Max}_{n_i,\underline{m}_i}\ U_i(E[W_{i1}], \text{Var}[W_{i1}]) \tag{4}$$

where $W_{i1} = (1 + r)W_{i0} + n_i(P^* - P_M) + \underline{m}'_i(\underline{Q}_1 - (1 + r)\underline{Q}_0)$, and r is the riskless rate of interest; $\underline{m}$ is the vector of demands for the other available risky assets; and $\underline{Q}_0$, $\underline{Q}_1$ are their period 0 and period 1 prices. To avoid complications from interactions between the positions taken in this market and the other assets in investor i's portfolio, we assume the price movement $(P^* - P_M)$ is distributed independently of $\underline{Q}_1$ and $\underline{Q}_0$. We also assume that investors are all "small" relative to the total market so that they treat the effect of n_i on P_M as being negligible.

From trader i's posterior distribution, the term $(P^* - P_M)$ has expected value $(\rho_i - P_M)$ and variance σ_i^2. Solving the first order condition from (4) gives the demand function

$$n_i = \frac{-U_{i1}}{2U_{i2}} \cdot \frac{\rho_i - P_M}{\sigma_i^2}$$

where the second subscripts indicate partial derivatives. The first term is the reciprocal of trader i's level of absolute risk aversion, which will be denoted $1/a_i$. We also substitute for ρ_i from (3), yielding

$$n_i = \frac{\alpha_i}{a_i\sigma_i^2}(P_i - P_M) \tag{5}$$

The demand function (5) is shown in Figure 1. It is linear in P_M and has its intercept at $P_M = P_i$. For a given expected price change demand increases as trader i becomes less risk averse, as the perceived level of risk is reduced, and as the weight he places on his prior information is increased.

By summing the demands of all investors, we obtain the market's demand function. Since there is no net supply of contracts, the market clears at the price of P_M at which net demand is zero.

$$\sum_i n_i = \sum_i \frac{\alpha_i}{a_i\sigma_i^2}(P_i - P_M) = 0, \quad \text{or}$$

$$P_M = \frac{\sum_i \frac{\alpha_i}{a_i\sigma_i^2} P_i}{\sum_j \frac{\alpha_j}{a_j\sigma_j^2}}$$

[4] The assumptions about the nature of investors' information sets lead to normally distributed returns. Assuming mean-variance utility functions imposes no additional restrictions on the model.

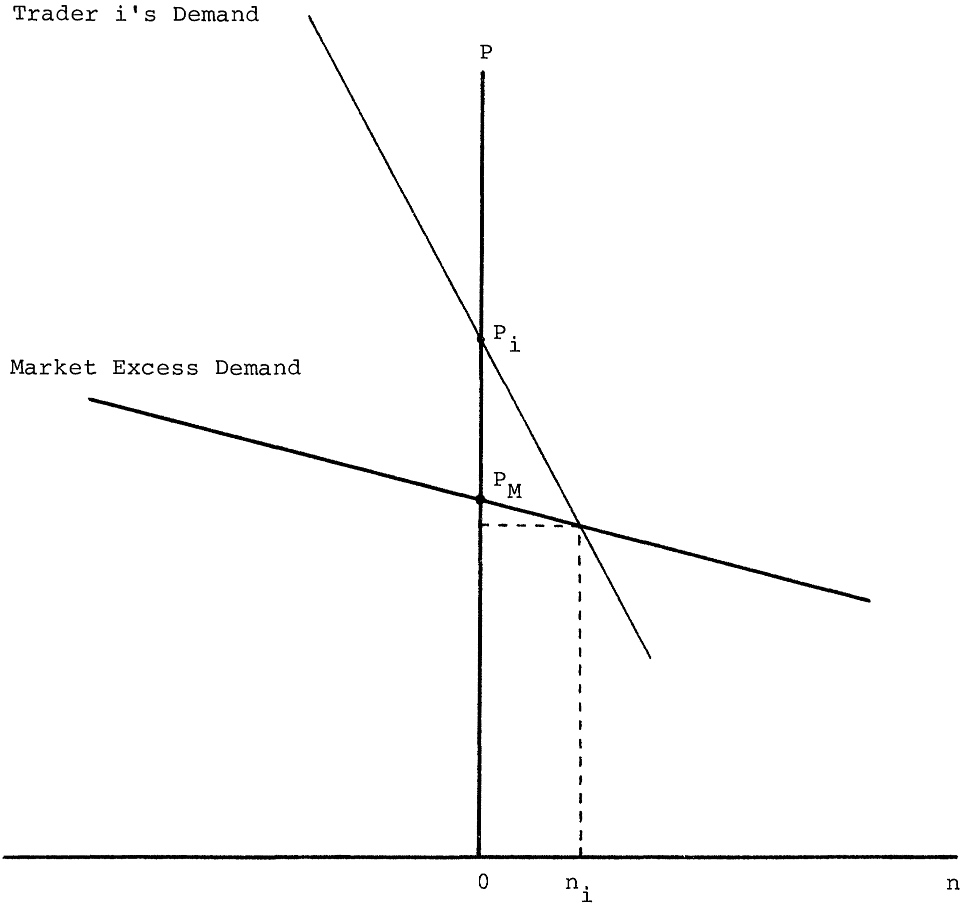

Figure 1. Individual and Market Excess Demand Functions.

Letting $X_i = \dfrac{\alpha_i/a_i\sigma_i^2}{\sum_j \alpha_j/a_j\sigma_j^2}$, and $\underline{X}$ be the vector of the X_i, the market price can be written simply as

$$P_M = \underline{X}'\underline{P} \tag{6}$$

The market price is a weighted average of the individual investors' prior price forecasts, where the weights are proportional to the weight the investor himself places on his prior, and inversely proportional to risk aversion and the degree of perceived risk. Note also that the weights sum to 1, i.e., $\underline{X}'\underline{e} = 1$.

Investor profits are determined by the price change over the period $(P_M - P^*)$, which may be thought of as the forecast error of the market price. Using (6) we have

$$\epsilon_M = (P_M - P^*) = (\underline{X}'\underline{P} - P^*) = \underline{X}'\underline{\epsilon} \tag{7}$$

$$\text{and} \quad E[\epsilon_M] = 0,$$

$$\operatorname{Var}[\epsilon_M] = \underline{X}'\Omega\underline{X}$$

The market's "error" ϵ_M has expected value zero over all realizations of $\underline{\epsilon}$—the market price is an unbiased forecast of the future price—and its variance and covariance with other variables depend on the vector of investor weights $\underline{X}$.

The market excess demand function is shown as the heavy line in Figure 1. In the situation depicted here, trader i expects a future price above the current market price, so he will buy forward an amount n_i. If the realized P^* is above P_M, he and all other investors who expected higher prices will have a profit of $(P^* - P_M)$ per contract. Notice that trader i's forecast does not have to be particularly accurate for him to make a profit. P^* may be much closer to P_M than to P_i, but as long as the difference is in the *direction* predicted by trader i, he will make money.

An individual trader's profit in a single period π_i is given by

$$\pi_i = n_i(P^* - P_M).$$

Substituting from Equation (5),

$$\pi_i = \frac{\alpha_i}{a_i\sigma_i^2}(P_i - P_M)(P^* - P_M) \tag{8}$$

By rewriting $(P_i - P_M)$ in terms of the forecast errors and taking the expected value, we derive trader i's expected profit:

$$E[\pi_i] = \frac{\alpha_i}{a_i\sigma_i^2}E[(\epsilon_i - \epsilon_M)(-\epsilon_M)]$$

$$E[\pi_i] = \frac{\alpha_i}{a_i\sigma_i^2}(\mathrm{Var}[\epsilon_M] - \mathrm{Cov}[\epsilon_i, \epsilon_M]) \tag{9}$$

This proves

THEOREM 1: An individual trader has a positive expected profit from trading on his information whenever the covariance between his forecast error and the market's error is less than the variance of the market's forecast error.

In Equation (9) each investor's expected profit or loss is a function of the relation between his information and that contained in the market price, which combines all investors' information. The market tends to redistribute wealth away from traders whose errors in forecasting P^* are too highly correlated with the market's error, and towards those with smaller covariance. Equation (9) reveals what makes one investor's information superior to another's. It is not so much accurate information—the own variance of trader i's forecast does not enter here—as it is *independent* information that leads to trading profits.

Notice that the existence of trading in this market requires that not all investors know their true expected profits as given by (9). An investor only takes a position if he expects a positive return. Since the market is a zero-sum game, this would imply some other investor had an expected loss. But no investor who *knew* his expected return was negative would willingly participate in the market, so there could be no trading. Trading arises when investors do not all know the true values of $\mathrm{Var}[\epsilon_M]$ and $\mathrm{Cov}]\epsilon_i, \epsilon_M]$ and each believes *his* information is not fully discounted in the market price.

In order to address questions of information aggregation, we must examine the behavior of profits for the market as a whole. Let $\underline{\pi}$ be the vector of the π_i, and let A be the diagonal matrix whose ith diagonal element is $(\alpha_i/a_i\sigma_i^2)$. Then the vector equation for trader profits can be written

$$\underline{\pi} = A(\underline{P} - P_M\underline{e})(P^* - P_M)$$

Rewriting this in terms of the vector of forecast errors gives

$$\underline{\pi} = A(\underline{\epsilon} - \underline{X}'\underline{\epsilon}\underline{e})(-\underline{X}'\underline{\epsilon})$$

$$\underline{\pi} = A(\underline{X}'\underline{\epsilon}\underline{\epsilon}'\underline{X}\underline{e} - \underline{\epsilon}\underline{\epsilon}'\underline{X}) \tag{10}$$

Taking the expected value over all realizations of $\underline{\epsilon}$, we have

$$E[\underline{\pi}] = A(\underline{X}'\Omega\,\underline{X}\underline{e} - \Omega\,\underline{X}) \tag{11}$$

Equation (11) gives the expected value of the profit vector as a function of traders' individual demand parameters α_i, a_i, and σ_i^2, the vector of market weights $\underline{X}$, and the error covariance matrix. For a given information structure Ω, an investor's expected profit or loss depends on the current weighting of his and all other investors' estimates of P^*. Thus, for some values of $\underline{X}$ he may be trading at an expected loss while he will earn an expected profit with a different $\underline{X}$. However, there is a unique stochastic equilibrium point $\underline{X}^{eq}$ at which the expected profit for every trader is zero. This will be shown below in the proof of Theorem 2.

The information efficiency of the market also varies with changes in the weighting vector, since from (7) the error variance of the market price is $\underline{X}'\Omega\underline{X}$. Thus, P_M may be a more or less accurate estimate of P^*, depending on which investors' price expectations receive the greatest weight. Since all information available to the market is contained in the set of investor price forecasts $\{P_i\}$, full informational efficiency for this market corresponds to a set of weights $\underline{X}^{eff}$ for which the variance of ϵ_M is a minimum.

The relation between trading profit equilibrium and information efficiency is summarized in Theorem 2.

THEOREM 2: (a) There is a unique, nondegenerate vector of market weights $\underline{X}^{eq}$ at which the expected value of profits for every trader is zero; (b) there is a unique vector of market weights $\underline{X}^{eff}$ at which the market is fully efficient and the variance of ϵ_M is minimized; and (c) these two vectors are identical and equal to the following expression which depends only on the information structure in the market population.

$$\underline{X}^{eq} = (\underline{e}'\Omega^{-1}\underline{e})^{-1}\Omega^{-1}\underline{e} = \underline{X}^{eff} \tag{12}$$

Proof:

(a) The zero-profit point $\underline{X}^{eq}$ is obtained from (11) by setting the term in parentheses equal to $\underline{0}$. This yields

$$\underline{X}'\Omega\underline{X}\underline{e} = \Omega\underline{X}.$$

Premultiplying both sides by $\underline{e}'\Omega^{-1}$, we have $\underline{e}'\Omega^{-1}\underline{e}(\underline{X}'\Omega\,\underline{X}) = \underline{e}'\underline{X}$. Since $\underline{e}'\underline{X} =$

1 and $(\underline{X}'\Omega\underline{X})$ is a nonzero scalar, this yields

$$(\underline{X}'\Omega\underline{X}) = (\underset{\sim}{e}'\Omega^{-1}\underset{\sim}{e})^{-1}$$

Substituting this in the first equation, premultiplying by Ω^{-1} and rearranging gives the first equality in Equation (12).

This derivation shows that $\underline{X}^{eq}$ is the unique zero-profit weighting with all investors participating. There are other such vectors for degenerate cases in which $X_i = 0$ for one or more investors. These are obtained in the same way from (12) using a reduced Ω matrix with the appropriate rows and columns deleted.

(b) and (c) The efficiency weights $\underline{X}^{eff}$ solve the constrained minimization problem

$$\underset{\underline{X}}{\text{Minimize}}\ \mathrm{Var}[\epsilon_M] = \underline{X}'\,\Omega\,\underline{X}$$

$$\text{subject to} \qquad \underline{X}'\,\underset{\sim}{e} = 1$$

From the Lagrangian equation, the first order conditions for a minimum are

$$\text{F.O.C. } \underset{\sim}{0} = 2\Omega\underline{X} - \lambda\underset{\sim}{e}$$

$$\underset{\sim}{0} = \underline{X}'\underset{\sim}{e} - 1$$

Premultiplying the first condition by $\tfrac{1}{2}\underset{\sim}{e}'\Omega^{-1}$ and rearranging gives

$$\underset{\sim}{e}'\underline{X} = \frac{\lambda}{2}\,\underset{\sim}{e}'\Omega^{-1}\underset{\sim}{e} = 1$$

$$\lambda = 2(\underset{\sim}{e}'\Omega^{-1}\underset{\sim}{e})^{-1}$$

Substituting this in the first order condition and premultiplying by Ω^{-1} yields the second equality in (12),

$$\underline{X}^{eff} = (\underset{\sim}{e}'\Omega^{-1}\underset{\sim}{e})^{-1}\Omega^{-1}\underset{\sim}{e} \qquad \text{Q.E.D.}$$

Market efficiency and trading profits are closely connected in a market where investors have heterogeneous information. A competitive speculative market produces a market price in which each investor's expectation is weighted according to the size of his investment. This is not necessarily the weight which would lead to his information being accurately discounted. If the covariance between an investor's estimation error and the market's error is smaller than the variance of the market's error, his information has too little weight in the market price, and by Equation (9) he earns expected trading profit. Market efficiency, measured by the variance of the market's error, would be improved if the market weight for this trader were increased and every other investor's weight reduced proportionately. This would increase his covariance with the market and decrease the market's variance, thus diminishing his expected profit. It is also possible for an investor to carry too great a weight, leading to an expected loss. Theorem 2 shows that there is only one market weighting at which no trader has an expected profit or loss from trading on the basis of his information, and this is also the point at which every investor's individual information is accurately discounted and the market is fully efficient. Notice that while the market price at this point is the full information price, the market is not necessarily in full efficiency equilibrium. That would require that investors were aware the market price was fully revealing and were acting accordingly. It may also happen that the efficiency weights $\underline{X}^{eq}$ could arise "accidentally" from the right combination of risk aversion and other parameters. In that case there would be no presumption that the market will remain fully efficient in the next period.

The equilibrium market weights which yield zero expected profit for all investors depend only on Ω. For some information structures, it is possible for X_i^{eq} to be negative. This will be the case if the covariance between traders i's forecast error and the errors of other investors is so high that when their weights are near the equilibrium values, $\text{Cov}[\epsilon_i, \epsilon_M] > \text{Var}[\epsilon_M]$ even if the market places no weight at all on P_i. Traders with negative equilibrium weights are clearly at a disadvantage trading in this market and would probably be better off dropping out or even trading against their information.[5] Note, though, that this is not necessarily true since the equilibrium weight is of relevance only when the market is close to full efficiency. For other values of $\underline{X}$ even a trader with X_i^{eq} negative can have an expected profit.

II. The Prospects for a Fully Efficient Market

Full informational efficiency in our model corresponds to a unique set of market weights on individual price expectations such that the market price is the minimum variance estimate of the future price. In a single period (the "short run") the market weights $\underline{X}$ will not in general be the efficiency weights $\underline{X}^{eq}$ given in (9), and the market will be less than fully efficient. This means that some investors' information will be underweighted, leading to positive expected profits, while others will have expected losses. However, there may be ways in which the market can approach full efficiency in the short run for particular information structures, or in the long run, if we allow learning or some other way for the market weights to converge over time to $\underline{X}^{eq}$. In this section we shall discuss several such possibilities.

Independent Information

In the general case, informational efficiency depends on the market's weighting of expectations, but for specific information structures the dependence may disappear. A trivial example of this is homogeneous expectations, where any weighting leads to the same price, which by definition contains all information.

A more important special case occurs when investors have independent information. Here Ω becomes a diagonal matrix,

$$\Omega_{IND} = \begin{matrix} V_1^2 & 0 & \cdots & 0 \\ 0 & V_2^2 & & \vdots \\ \vdots & & \ddots & \\ 0 & \cdots & & V_N^2 \end{matrix} \tag{13}$$

[5] Some trading systems apparently try to do this by trading opposite to what the "odd-lot" traders are doing. In this way, a negative weight is applied to the information of those investors who are felt to be the least informed.

From (12) trader i's equilibrium weight is

$$X_i^{eq} = \frac{1/V_i^2}{\sum_j 1/V_j^2} \tag{14}$$

All investors have zero expected profits when each individual prior price estimate enters the market price with a weight that is inversely proportional to its forecast variance. As expected, better information leads to a higher efficiency weight. But notice that regardless of forecast accuracy, all equilibrium weights are positive. No trader with independent information, even though of inferior quality, is driven out of the market, since with zero weight his covariance with the market's error is also zero and by (9) he has an expected profit.

An important feature of the independent information case is the effect of an increase in the number of traders. The market is most efficient at $\underline{X}^{eq}$, where its forecast variance is $\underline{X}^{eq\prime}\Omega\underline{X}^{eq}$. Substituting from (14) we have

$$\underline{X}^{eq\prime}\Omega\underline{X}^{eq} = (\underline{e}'\Omega^{-1}\underline{e})^{-1} = \frac{1}{\sum_{J=1}^{N} 1/V_j^2} \tag{15}$$

Adding another investor unambiguously improves the market's maximum efficiency. Moreover, the market's minimum forecast variance goes to zero as the number of investors becomes large. In the limit, despite heterogeneous information, or rather because of it, since it is the independence assumption that leads to this result, the market price is equal to the true final price with probability 1 when the market is fully efficient. This result holds in any neighborhood of $\underline{X}^{eq}$.

THEOREM 3: When investors have independent information, the market's forecast variance approaches arbitrarily close to zero in any neighborhood of $\underline{X}^{eq}$ for N sufficiently large.

Proof:

Define a neighborhood of $\underline{X}^{eq}$ as $\underline{X} = \underline{X}^{eq} + \underline{d}$, where $\underline{d}$ is vector of deviations. Let k be some positive constant such that:

$$k \geq \text{Maximum} \frac{|d_i|}{X_i^{eq}}$$

The market's variance at weights $\underline{X}$ is given by

$$\underline{X}'\Omega\underline{X} = \underline{X}^{eq\prime}\Omega\underline{X}^{eq} + 2\underline{d}'\Omega\underline{X}^{eq} + \underline{d}'\Omega\underline{d}$$

Since $\underline{X}'\underline{e} = 1 = \underline{X}^{eq\prime}\ \underline{e}$, then $\underline{d}'\underline{e} = 0$. Thus, the second term in this expression, which can be written $2\underline{d}'\underline{e}\frac{1}{\sum_j 1/V_{j'}^2}$ is identically zero.

Using the definition of k we have

$$0 < \underline{X}'\Omega\underline{X} \leq (k^2 + 1)\underline{X}^{eq\prime}\Omega\underline{X}^{eq}$$

and

$$\lim_{N\to\infty} \underline{X}'\Omega\underline{X} = 0. \qquad \text{Q.E.D.}$$

With independent information the market approaches full efficiency, and in fact becomes a perfect predictor of the future price when the number of investors becomes large. Thus, for this special case, even in the short run, the market price is fully revealing almost regardless of the weights on individual information.

Independent information is not likely to be an adequate description of the information structure of a real-world speculative market. Although large numbers of investors participate, the variance of market prices is not zero. Rather, forecast errors tend to be relatively highly correlated across investors when prices are affected by events like crop failures which are generally unexpected. Without independence, the "Law of Large Numbers" result of Theorem 3 no longer holds, and full efficiency does not occur in the short run.

Long Run Efficiency Through Wealth Redistribution

Unless the Ω matrix has a special structure, the market will normally be less than fully efficient in the short run. There remains the possibility that over time, as investors trade in the market and those with underweighted information earn trading profits while those with too great a weight experience losses, the market will converge to the full information homogeneous expectations solution in the long run.

One common heuristic argument for long run efficiency is based on the view of a speculative market as a sort of Darwinian process. Since the market redistributes wealth away from relatively uninformed investors to those with better information, over time the uniformed tend to be eliminated and the market comes to be dominated by the traders with the best information.

For this mechanism to work in our model, the market weights must depend on past trading profits. For example, absolute risk aversion may be related to wealth so that an investor alters the size of his position for a given expected price change as he experiences profits and losses in the market. The assumption of constant absolute risk aversion which rules out this effect has played an important role in earlier models in which the price became a sufficient statistic. When decreasing absolute risk aversion is assumed, a trader who earns a profit in one period because the market weighted his information too lightly will have more wealth and (generally)[6] a greater weight in the next period. As his market weight approaches the informationally efficient value, *ceteris paribus*, the market becomes more efficient.

The behavior of the market weights as wealth is transferred from investors with "bad" information to those with "good" information was examined for a special case of this model in Figlewski [3]. In that model, as in the present one,

[6] This intuitively plausible market behavior need not always occur. An individual's market weight is given by

$$X_i = \frac{\alpha_i / a_i \sigma_i^2}{\sum_j \alpha_j / a_j \sigma_j^2}$$

If absolute risk aversion decreases with wealth, trading profits increase the numerator unambiguously. But it is possible that a given redistribution of wealth can increase the denominator sufficiently that X_i is reduced in spite of the trading profit.

the stochastic nature of the system rules out permanent convergence to a fully revealing equilibrium. At each point in time, $\underline{X}_{t+1}$ depends on the current market weights $\underline{X}_t$ and the particular realization of the random vector $\underline{\epsilon}_t$. Since $\underline{X}^{eq}$ has measure zero in the space of possible weightings, its probability of occurrence is zero. Even if full efficiency were to be achieved in one period, the market will remain there in the next only if the realization of $\underline{\epsilon}$ *is* $\underline{0}$. Thus, wealth redistribution through trading is not sufficient in itself to produce full efficiency, even over the long run.

Long Run Convergence Through Learning

The most plausible way for full efficiency to occur is if investors can learn the relevant properties of the market's information weighting function by observing its behavior over time. By (9), an investor can calculate the true expected value and variance of his speculative profit if he knows $\mathrm{Var}[\epsilon_M]$ and $\mathrm{Cov}[\epsilon_i, \epsilon_M]$. Assuming he is aware that P_M and P_i are unbiased, he may try to estimate these parameters from past data using

$$\hat{\sigma}_M^2 = \sum_{t=1}^{T} (P_M - P^*)_t^2/T$$
$$\hat{C}_{iM} = \sum_{t=1}^{T} (P_i - P^*)_t(P_M - P^*)_t/T \tag{16}$$

as his estimates. If $\mathrm{Var}[\epsilon_M]$ and $\mathrm{Cov}[\epsilon_i, \epsilon_M]$ were constant parameters, it is clear that (16) would yield the correct values in the long run. But since ϵ_M depends on the market weights $\underline{X}_t$, the probability distribution from which each period's observation in (16) is drawn changes over time. Further, each X_i depends on past ϵ_i's so the market weights evolve in a way that cannot be known through observation of the market price alone. Learning one's true expected profit, if it is possible at all, is a complex process. The investor must estimate the current values of $\mathrm{Var}[\epsilon_M]$ and $\mathrm{Cov}[\epsilon_i, \epsilon_M]$ from past data drawn from probability distributions which differ from the current one and from each other. It is evident that any estimate will depend on the realizations of $\underline{\epsilon}$ in the sample used. This implies that if full convergence is to be achieved with probability 1, it can only occur in infinite time. Finally, if the information structure Ω also changes over time in an unknown manner, as for example, investors enter or leave the market, convergence through observing the market's behavior in previous periods is clearly impossible.

The following theorem, stated without formal proof, summarizes these general results.

THEOREM 4: In the above model with a general information structure, permanent convergence to full efficiency (a) cannot occur through wealth redistribution resulting from trading profits and losses alone; (b) cannot occur through learning if Ω varies randomly over time; and, (c) cannot occur through learning in less than infinite time when Ω is fixed.

While the market possesses an equilibrium point at which the price fully reveals all investors' information, its normal condition will be one of less than full efficiency. Some traders will have expected profits according to Equation (9) from information that the market has not fully discounted, but no trader will know exactly what his expected speculative profit is. Informational efficiency will vary from period to period with the market's weighting of individual information and the average level of efficiency will depend on the nature of information differences and other characteristics within the investor population.

III. Concluding Comments

This paper has concentrated on two major issues raised by information diversity. One is the idea that a trader with good information should be able to use it to make a trading profit. We showed that "good" information must be defined relative to the market price. It is the relation between an investor's own information and the information already contained in the price which determines whether he has an expected trading profit. The covariance condition of Equation (9) shows that a trader profits not because he has more, or better, information than the market, but because he has *some* information the market has not fully discounted and which tells him which way the market will move.

The second issue is whether we should expect a competitive speculative market to aggregate investor's diverse opinions into a single price which accurately discounts the available information. The answer to this question has important implications for the social value of a competitive financial market versus centralized price setting, as is done for many asset prices like interest rates and some exchange rates, or versus no market at all, as in the case of onions futures.

Our model showed that a market with heterogeneous information will almost certainly not be fully efficient in the short run. The market price reflects each investor's price expectation weighted by the amount of his investment. Since this depends on parameters unrelated to the quality of his information, it is unlikely that the market weights in a given period will be those which accurately discount all information.

More important is whether a speculative market converges to full efficiency in the long run. Like earlier models, this model possesses a point of maximum efficiency at which all information is accurately reflected in the price. At this point expected profits are zero for all investors and the market breaks down. However, while such an equilibrium exists, in general there may be no mechanism by which the market can converge to it, except over the very long run. Only in particular instances, such as the case of independent information in which the market's weighting becomes unimportant, is a fully revealing equilibrium approached.

Thus, in this model, the natural state of a speculative market is one in which the market price reveals some, but not all, of the traders' diverse information. Each one takes a speculative position based on his prior information and his estimate of the information already discounted in the market price. While an investor may be wrong about his expected trading profits, it is not irrational for him to trade against the market.

REFERENCES

1. Douglas Diamond and Robert Verrecchia. "Information Aggregation in a Noisy Rational Expectations Economy." *Journal of Financial Economics*, forthcoming.
2. Eugene Fama. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of*

Finance 25 (May 1970), 333–417.

3. Stephen Figlewski. "Market 'Efficiency' in a Market with Heterogeneous Information." *Journal of Political Economy* 86 (August 1978), 581–97.
4. Sanford Grossman. "On the Efficiency of Competitive Stock Markets Where Traders Have Diverse Information." *Journal of Finance* 31 (May 1976), 573–85.
5. ———. "Further Results on the Informational Efficiency of Competitive Stock Markets." *Journal of Economic Theory* 18 (1978), 81–101.
6. ——— and Joseph Stiglitz. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70 (June 1980), 393–408.
7. Martin Hellwig. "On the Aggregation of Information in Competitive Markets." *Journal of Economic Theory* 22 (June 1980), 447–98.
8. Jeffrey Jaffe and R. Winkler. "Optimal Speculation Against an Efficient Market." *Journal of Finance* 31 (March 1976), 49–61.
9. Joseph T. Williams. "Capital Asset Prices with Heterogeneous Beliefs." *Journal of Financial Economics* 5 (1977), 219–39.