



Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations

J. Michael Harrison, David M. Kreps

The Quarterly Journal of Economics, Volume 92, Issue 2 (May, 1978), 323-336.

Stable URL:

<http://links.jstor.org/sici?sici=0033-5533%28197805%2992%3A2%3C323%3ASIBIAS%3E2.0.CO%3B2-Y>

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The Quarterly Journal of Economics is published by The MIT Press. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/mitpress.html>.

The Quarterly Journal of Economics
©1978 The MIT Press

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2002 JSTOR

SPECULATIVE INVESTOR BEHAVIOR IN A STOCK MARKET WITH HETEROGENEOUS EXPECTATIONS*

J. MICHAEL HARRISON AND DAVID M. KREPS

I. Introduction, 323.—II. Formulation, 325.—III. An example, 326.—IV. Consistent price schemes, 328.—V. Back to the example, 332.—VI. Relaxing the assumptions, 333.—VII. Concluding remarks, 334.

I. INTRODUCTION

Consider a common stock that pays dividends at a discrete sequence of future times: $t = 1, 2, \dots$. Taking all other prices and the random process that determines future dividends as exogenously given, we can ask what will be the price of the stock? In a world with a complete set of contingency claims markets, in which every investor can buy and sell without restriction, the answer is given by arbitrage. Let $d_t(x_t)$ denote the dividend that will be paid at time t if contingency x_t prevails, and let $\beta_t(x_t)$ denote the current ($t = 0$) price of a one dollar claim payable at time t if contingency x_t prevails. Then the current stock price must be $\sum_t \sum_{x_t} \beta_t(x_t) d_t(x_t)$. Furthermore, in such a world it makes no difference whether markets reopen after initial trading. If markets were to reopen, investors would be content to maintain the positions they obtained initially (cf. Arrow, 1968).

The situation becomes more complicated if markets are imperfect or incomplete or both. Ownership of the stock implies not only ownership of a dividend stream but also the right to sell that dividend stream at a future date. Investors may be unable initially to achieve positions with which they will be forever content, and thus the current stock price may be affected by whether or not markets will reopen in the future. If they do reopen, a *speculative phenomenon* may appear. An investor may buy the stock now so as to sell it later for more than he thinks it is actually worth, thereby reaping capital gains. This possibility of speculative profits will then be reflected in the current price. Keynes (1931, Ch. 12) attributes primary importance to this phenomenon (and goes on to suggest that it might be better if markets never reopened).

We say that investors exhibit *speculative behavior* if the right to resell a stock makes them willing to pay more for it than they would pay if obliged to hold it forever. This phenomenon will not occur in

* Our research was partially supported by a grant from the Atlantic Richfield Foundation to the Graduate School of Business, Stanford University.

a world with one period remaining (as in the capital-asset-pricing model), in a world where all investors are identical, or in a world with complete and perfect contingency claims markets. Our characterization of speculative behavior is not intended to be all-inclusive, nor is it intended to transcend our partial-equilibrium framework. For different definitions and analyses see Feiger (1976) and Hirshleifer (1975).

In Section II we present an extremely simple model (of the market for a single stock) in which the speculative phenomenon can be sharply seen. The key element of our model is the existence of heterogeneous expectations within the community of potential investors. Beyond this reasonable supposition, very restrictive assumptions are made. Investors are partitioned into a finite number of internally homogeneous classes, each class having (what amounts to) infinite collective wealth. All investors have access to the same substantive economic information (although members of different classes may arrive at different subjective probability assessments on the basis of that information). Members of each class are risk-neutral, so that any income stream is valued at its (subjective) expected present worth. For notational convenience only (see Section V) the discount factor is taken to be known and constant. Most importantly, short sales of the stock in question are forbidden.

In Section III we give a simple numerical example that illustrates the delicate nature of price equilibrium in our model. It is shown how members of one class bid up the price of the stock in anticipation of future opportunities for selling it to members of other classes, at higher prices than they themselves would be willing to pay. It is seen that, if an equilibrium price is to be found, it must exceed what any class would be willing to pay for the stock if obliged to hold it forever.

With this motivation, we return to the general model in Section IV. We adapt Radner's (1972) criterion for price equilibrium to the partial-equilibrium context of our model, and a price scheme (or price system) that meets this criterion is called *consistent*. Roughly, a price scheme is consistent if it does not allow any investor to garner excessive expected return through adroit speculation. Standard mathematical results are cited to establish that consistency is equivalent to a simple martingale-like property for the prices and that there exists a minimal consistent price scheme. Typically, the minimal consistent price will exceed every investor's expected present worth of future dividends. Investors are willing to pay a "speculative premium" because of anticipated capital gains. There do exist nonmi-

nimal consistent price schemes, which are obtained by superimposing the usual sort of Ponzi scheme on the minimal consistent prices.

In Section V we explicitly compute the minimal consistent prices for our example. Section VI discusses informally the generality of our results. The precise results obtained of course depend on the assumptions of risk neutrality, infinite wealth, and no short sales. But we argue that our model is a very good approximation to one where only the short sales restriction remains in force. We further believe that the qualitative phenomenon that we have called speculation would occur in a model with the short sales restriction relaxed, but it would not be so clearly visible.

Section VII contains some miscellaneous concluding remarks concerning connections between our simple model and the standard theories of fundamental and technical analysis and the random walk hypothesis.

II. FORMULATION

We consider a market (for a single stock) where trading takes place at a discrete sequence of times: $t = 0, 1, \dots$. Time $t = 0$ corresponds to the present. Future dividends will be paid by the stock at times $t = 1, 2, \dots$, and we denote these (random) dividends by d_1, d_2, \dots , respectively. All dividends are nonnegative. We assume that dividend d_t will be declared immediately prior to time t trading, the dividend being paid to whoever held the stock between time $t - 1$ and t . For each $t = 1, 2, \dots$ we denote by ξ_t the vector of (new) economic information made available to investors between times $t - 1$ and t . Thus, the total economic information available when trading commences at time t is

$$\chi_t = (\xi_1, \xi_2, \dots, \xi_t), \quad t = 1, 2, \dots$$

The dividend d_t is known at time t and hence is included within the information χ_t . We write $d_t(\chi_t)$ to emphasize this functional dependence (or measurability of d_t with respect to χ_t). We denote by X_t the set of all possible realizations (or the support) of χ_t . A point in X_t (a realization of χ_t) will be generically denoted by x_t . To avoid technical difficulties with conditional expectations, we assume that each X_t is countable. With appropriate care, however, the entire treatment can be extended to Borel X_t . For completeness we define X_0 to be a singleton $\{x_0\}$ and χ_0 to be the (trivially random) vector that is identically x_0 . One may interpret x_0 as the vector of economic information available to investors at time zero. We assume that inves-

tors have no control over operations of the firm whose stock is being traded. Thus, they view the dividend stream $\{d_t\}$, and more generally the economic information process $\{\chi_t\}$, as an exogenous source of uncertainty.

We denote by A the finite set of investor classes. Each investor has a subjective probability distribution for the random process $\{\chi_t\}$, and we assume that two members of the same class assess the same probability distribution, although members of different classes may assess different distributions. We denote by $E^a[\cdot]$ an expected value with respect to the probability distribution shared by members of class $a \in A$.

All investors are assumed to be risk-neutral, discounting future income at rate γ per period. At time t , in contingency x_t , an investor of class a is thus indifferent between a random stream of payments

$$\{y_{t+s}(\chi_{t+s}); s = 0, 1, \dots\}$$

and a certain payment in the amount

$$E^a \left[\sum_{s=0}^{\infty} \gamma^s y_{t+s}(\chi_{t+s}) \mid \chi_t = x_t \right].$$

Here $E^a[\cdot \mid \chi_t = x_t]$ means conditional expectation, given that $\chi_t = x_t$. Each class of investors is assumed large enough to prevent collusion and wealthy enough to buy up all of the stock if it so desires.

Our key assumption is that the stock in question cannot be sold short. This is necessary to prevent our infinitely wealthy investor classes from making what amount to infinite side bets when their probability assessments for future events differ. Combining this with our other assumptions, one sees that an equilibrium for our market must have the following property. In every contingency, all of the stock is bought up by whatever class values it most highly, and the price equals that highest value.

III. AN EXAMPLE

Suppose that there are two investor classes, denoted by superscripts 1 and 2, and that in every period the dividend is either zero or one. Each of the two classes, in making probability assessments for future economic events, believes that the only relevant information in χ_t is the most recent dividend d_t , and the process $\{d_1, d_2, \dots\}$ is perceived by each class as a stationary Markov chain with state space $\{0, 1\}$. Throughout our discussion of this example, we shall say that the market is in *state* d when the most recent dividend is d . We denote

by $q^a(d, d')$ the probability assessed by class a for a transition from state d to state d' . We define the transition matrices,

$$Q^a = \begin{bmatrix} q^a(0, 0) & q^a(0, 1) \\ q^a(1, 0) & q^a(1, 1) \end{bmatrix} \quad \text{for } a = 1, 2,$$

and assume the specific numerical data,

$$Q^1 = \begin{bmatrix} 1/2 & 1/2 \\ 2/3 & 1/3 \end{bmatrix}, \quad Q^2 = \begin{bmatrix} 2/3 & 1/3 \\ 1/4 & 3/4 \end{bmatrix},$$

and

$$\gamma = 0.75.$$

Note that each class is positive that it knows the actual transition matrix, so that no amount of transition frequency data will alter their assessments. This is implicit in the statement that only the current state of the market is judged relevant.

Beginning the search for equilibrium prices, we compute the value to the investors of buying the stock and holding it forever. Given their assumed preference structure, this is the expected present worth of future dividends. Denote by $p^a(d)$ the expected present worth of future dividends to an investor of class a when the state of the market is d . Simple calculations give

$$\begin{aligned} p^1(0) &= 4/3 = 1.33, & p^1(1) &= 11/9 = 1.22, \\ p^2(0) &= 16/11 = 1.45, & p^2(1) &= 21/11 = 1.91. \end{aligned}$$

It is interesting to note that members of class 2 assess a higher expected present worth of future dividends than do members of class 1, regardless of the current state. One might therefore conjecture that members of class 2 will always hold the stock and that the price in each state d will be $p^2(d)$.

But if members of class 1 anticipate these prices, the market cannot be in equilibrium. A class 1 investor can buy stock in state zero with the intention of selling it (for 1.91) the first time that a transition to state one occurs (the first time a dividend is declared). From his perspective the expected present worth of revenues generated by the stock under this plan is

$$[(1/2)(0.75) + (1/2)^2(0.75)^2 + \dots](1 + 1.91) = 1.75,$$

which exceeds the 1.45 purchase price in state zero. Thus, if the price in state one is 1.91 or more, the price in state zero must be at least 1.75.

To understand what is happening here, look first at the two tran-

sition matrices. When the market is in state one, investors of class 2 are optimistic about receiving dividends in the immediate future. This is because they assess a probability of $\frac{3}{4}$ that a dividend will be declared in the next period. Members of class 1 are pessimistic about immediate dividend prospects starting from state one, but they cannot sell short on the basis of their belief. When the market is in state zero, class 1 investors are more optimistic than class 2 investors about a transition to state one, and this opens up for them the possibility of (expected) capital gains. They can hold the stock until a dividend is declared, knowing that class 2 will view this as a positive development. At that point, class 1 can unload the stock at what it believes is an inflated price. Members of class 1 are willing to pay more than 1.45 in state zero not because they foresee a future of many one dollar dividends, but because they foresee an event that members of class 2 will take as a signal of good times ahead.

Returning to the equilibrium story, it now develops that 1.91 is too low a price in state one. Members of class 2 can buy in state one, hold until a transition to state zero occurs, and sell at that point to members of class 1 for (at least) 1.75. This generates for them revenues with an expected present worth of (at least) 2.03. Having gone two layers deep in what is obviously an infinite progression, we now ask in generality where it all stops.

IV. CONSISTENT PRICE SCHEMES

Although dividends are taken as exogenously determined, the situation is quite different with stock prices. Investors collectively determine through their current actions the current stock price, and they realize that they will continue to do so in the future. Since the value they attach to the stock today involves the price it will command in the future, they must currently have a conception (clear or otherwise) of how they will price the stock in the future. To introduce our notion of price equilibrium, we require a few rather formal definitions.

We define a general *price scheme* as a sequence of nonnegative real valued functions $\{p_0, p_1, \dots\}$ such that the domain of p_t is X_t for each $t = 0, 1, 2, \dots$ (Note that p_0 is then just a nonnegative constant.) Such a family of functions provides a potential mechanism for translating any sequence $\{x_1, x_2, \dots\}$ of partial economic histories into a corresponding price sequence $\{p_0, p_1(x_1), p_2(x_2), \dots\}$.

We define a *t-legitimate selling strategy* as a (possibly infinite) integer-valued random variable T , which is optional¹ with respect to

$\{\chi_t\}$ and satisfies $t + 1 \leq T \leq \infty$. Any such strategy T represents a potential plan for the sale of stock held at (just after) time t , with the event $\{T = t + k\}$ corresponding to the set of all circumstances that would cause the investor to sell the stock at time $t + k$. The requirement that T be optional with respect to $\{\chi_t\}$ simply insures that the choice of whether or not to sell at time $t + k$ will be based solely on the information χ_{t+k} that is then available. Finally, we shall say that a price scheme $\{p_0, p_1, \dots\}$ is *consistent* if

(1)

$$p_t(x_t) = \max_{a \in A} \sup_T E^a \left[\sum_{k=t+1}^T \gamma^{k-t} d_k(\chi_k) + \gamma^{T-t} p_T(\chi_T) \mid \chi_t = x_t \right]$$

for all $t = 0, 1, \dots$ and all $x_t \in X_t$, the supremum being over all t -legitimate selling strategies. Our argument that (1) is a natural condition for equilibrium in the market goes as follows. Suppose that a price scheme $\{p_t(\cdot)\}$ is to be followed. Then for each class $a \in A$, the expression,

$$\sup_T E^a \left[\sum_{k=t+1}^T \gamma^{k-t} d_k(\chi_k) + \gamma^{T-t} p_T(\chi_T) \mid \chi_t = x_t \right],$$

represents the maximum expected present worth that an investor from that class can realize from stock held at time t when he follows a legitimate strategy for subsequent sale, given the economic information χ_t available at time t . Thus, the right-hand side of (1) is the maximum amount that the stock is worth to any investor at time t . If this amount were strictly larger than the price $p_t(x_t)$, then members of the maximizing class(es) would compete among themselves to drive the price up. If it were smaller, then whoever held the stock at time t would want to sell but would find no buyer, so the price would have to fall.

This consistency condition is simply a partial-equilibrium version of Radner's (1972) equilibria of "plans, prices, and price expectations." It rests on a heroic assumption of what might be called *perfect contingent foresight*. We are analyzing market operations as if current and future prices were being simultaneously determined in a complete array of imperfect "futures markets." In negotiating a current price for the stock, investors are forming a clear and identical conception as to the price that will prevail in each future contingency. These conceptions, moreover, must be consistent among themselves and with the exogenous data. The extent to which this assumption colors our results should be carefully noted.

As we have seen in our example, the achievement of an equilibrium entails a delicate interplay among the subjective assessments of the various classes, so it is not immediately obvious that consistent price schemes always exist. To show that they do exist, we first simplify the definition of consistency. Because dividends and price schemes are nonnegative by assumption, a standard result in discrete potential theory allows us to drop the supremum over all selling strategies and consider only strategies that buy, hold for one period, and then sell.

PROPOSITION 1. A price scheme $\{p_t\}$ is consistent if and only if, for all t and x_t ,

$$(2) \quad p_t(x_t) = \max_{a \in A} E^a[\gamma d_{t+1}(\chi_{t+1}) + \gamma p_{t+1}(\chi_{t+1}) | \chi_t = x_t].$$

The proof is as follows. Suppose that (2) holds. Then for all $a \in A$, $p_t(x_t) \geq E^a[\gamma d_{t+1}(\chi_{t+1}) + \gamma p_{t+1}(\chi_{t+1}) | \chi_t = x_t]$, so speculation is viewed by all investors as no better than a "fair" game. Doob's optional stopping theorem (cf. Chung, 1974) implies that no nonanticipatory strategy for speculation can do better than "break even." (Because d_t and p_t are nonnegative, unbounded, and even infinite, stopping times are permitted.) That is,

$$p_t(x_t) \geq E^a \left[\sum_{k=t+1}^T \gamma^{k-t} d_k(\chi_k) + \gamma^{T-t} p_T(\chi_T) | \chi_t = x_t \right]$$

for all $a \in A$ and for all optional $T > t$. Equation (1) follows immediately. Conversely, if (1) holds, it is evident that

$$p_t(x_t) \geq \max_{a \in A} E^a[\gamma d_{t+1}(\chi_{t+1}) + \gamma p_{t+1}(\chi_{t+1}) | \chi_t = x_t]$$

for all x_t . Suppose that strict inequality held for some x_t . For that x_t , applying the stopping theorem at time $t + 1$ to any optional $T \geq t + 1$ yields

$$\begin{aligned} p_t(x_t) &> \max_{a \in A} E^a[\gamma d_{t+1}(\chi_{t+1}) + \gamma p_{t+1}(\chi_{t+1}) | \chi_t = x_t] \\ &\geq \max_{a \in A} \sup_T E^a \left[\sum_{k=t+1}^T \gamma^{k-t} d_k(\chi_k) + \gamma^{T-t} p_T(\chi_T) | \chi_t = x_t \right]. \end{aligned}$$

This contradicts equation (1), and thus (1) implies (2).

Economically, Proposition 1 has the following interpretation. Investors can achieve an expected net present value exceeding zero using adroit (but legitimate) trading strategies if and only if they can

do so at some point using a simple strategy that buys, holds for one period, and then sells. (This is not to say that strategies that hold the stock for many periods will always achieve an expected net present value of zero when prices are consistent. See the discussion in Section VII.)

To identify explicitly a consistent price scheme, set $p_t^0 \equiv 0$ and for $n = 1, 2, \dots$, recursively define

$$(3) \quad p_t^n(x_t) = \max_{a \in A} E^a[\gamma d_{t+1}(\chi_{t+1}) + \gamma p_{t+1}^{n-1}(\chi_{t+1}) | \chi_t = x_t].$$

Observe that $p_t^n(x_t)$ is nondecreasing in n and hence approaches some (possibly infinite) limit $p_t^*(x_t)$ as $n \rightarrow \infty$.

PROPOSITION 2. The price scheme $\{p_t^*\}$ is consistent. If $\{p_t\}$ is any other consistent scheme, then $p_t(x_t) \geq p_t^*(x_t)$ for all t and x_t . That is, $\{p_t^*\}$ is the minimal consistent price scheme.

A sketch of the proof follows. To show that $\{p_t^*\}$ satisfies (2) and hence is consistent, let $n \rightarrow \infty$ in equation (3). On the right-hand side, the limit can be brought inside the integral by the monotone convergence theorem. If $\{p_t\}$ is a consistent price scheme, then $p_t(x_t) \geq 0 = p_t^0(x_t)$. It is then straightforward to show by induction that $p_t(x_t) \geq p_t^n(x_t)$ for all n , and thus $p_t(x_t) \geq p_t^*(x_t)$.

It is worth noting that there exist nonminimal consistent price schemes whenever $\{p_t^*\}$ is not identically infinite. For example, one can easily verify that the price scheme $\{p_t\}$ given by $p_t = p_t^* + c\gamma^{-t}$ for $c > 0$ is consistent. The difference between a nonminimal consistent price scheme and $\{p_t^*\}$ is a sort of speculative bubble or Ponzi scheme. For such a speculative bubble to be "viable," the time horizon must be infinite. If there is some natural horizon n such that $d_t \equiv 0$ for $t \geq n$, and if we therefore require that $p_t \equiv 0$ for $t > n$ (as seems reasonable), then $\{p_t^*\}$ is the unique consistent price scheme. (Also, p_t^{n-t} as defined recursively in equation (3) will be p_t^* in this case.)

Because $T \equiv \infty$ is a t -legitimate selling strategy, it follows immediately from (1) that any consistent price scheme $\{p_t\}$ must satisfy

$$p_t(x_t) \geq \max_{a \in A} E^a \left[\sum_{k=t+1}^{\infty} \gamma^{k-t} d_k(\chi_k) | \chi_t = x_t \right].$$

That is, a consistent price must be at least as large as every investor's expected present worth of future dividends. With heterogeneous investors this inequality is typically strict (as in our example). With homogeneous investors, however, our results specialize as follows.

PROPOSITION 3. If there is a single investor class a , then the minimal consistent price scheme $\{p_t^*\}$ is the expected present worth of future dividends (for that class). Moreover, if $\{p_t^*\}$ is finite, another price scheme $\{p_t\}$ is consistent if and only if

$$p_t(x_t) = p_t^*(x_t) + \gamma^{-t}Z_t(x_t),$$

where $\{Z_t\}$ is a nonnegative martingale with respect to $\{\chi_t\}$ and the investors' probability assessments, meaning that

$$E^a[Z_{t+1}(\chi_{t+1})|\chi_t = x_t] = Z_t(x_t) \quad \text{for all } t \text{ and } x_t.$$

Proof. Inductively,

$$p_t^n(x_t) = E^a \left[\sum_{k=1}^n \gamma^k d_{t+k} | \chi_t = x_t \right],$$

so by monotone convergence

$$p_t^*(x_t) = E^a \left[\sum_{k=1}^{\infty} \gamma^k d_{t+k} | \chi_t = x_t \right].$$

The characterization of nonminimal price schemes follows by subtracting equation (2) for $\{p_t^*\}$ from equation (2) for any other consistent $\{p_t\}$.

V. BACK TO THE EXAMPLE

The stationary nature of our example guarantees that the minimal consistent price scheme $\{p_t^*\}$ will be time-independent. That is, it consists simply of prices $p^*(0)$ and $p^*(1)$ for states zero and one, respectively. Thus (2) can be rewritten as

$$\begin{aligned} p^*(0) &= \max\{(3/4)(1/2)p^*(0) + (3/4)(1/2)(1 + p^*(1)), \\ &\quad (3/4)(2/3)p^*(0) + (3/4)(1/3)(1 + p^*(1))\} \\ p^*(1) &= \max\{(3/4)(2/3)p^*(0) + (3/4)(1/3)(1 + p^*(1)), \\ &\quad (3/4)(1/4)p^*(0) + (3/4)(3/4)(1 + p^*(1))\}, \end{aligned}$$

for which the unique solution is $p^*(0) = 24/13 = 1.85$ and $p^*(1) = 27/13 = 2.04$. At these prices the stock is held by class 1 investors in state zero and by class 2 investors in state one. In Section III we described a (monotonically increasing) price iteration procedure of the form

$$p_t^n(x_t) = \max_{a \in A} \sup_T E^a \left[\sum_{k=t+1}^T \gamma^{k-t} d_k(\chi_k) + \gamma^{T-t} p_T^{n-1}(\chi_T) | \chi_t = x_t \right],$$

with $p_t^0 \equiv 0$. This iterative procedure is different from that used to define p_t^* in Section IV, but they can be shown in general to yield the same limit as $n \rightarrow \infty$.

VI. RELAXING THE ASSUMPTIONS

For each $a \in A$, let $\{\gamma_t^a\}$ be a sequence of nonnegative functions such that γ_t^a has domain X_t . Let $\{y_t, y_{t+1}, \dots\}$ be a random stream of payments as in Section II. Our entire treatment is easily generalized to the case where members of class a , observing contingency x_t at time t , value the stream $\{y_{t+s}\}$ at

$$E^a \left[\sum_{s=0}^{\infty} \left(\prod_{k=1}^s \gamma_{t+k}^a(x_{t+k}) \right) y_{t+s}(x_{t+s}) \mid x_t = x_t \right].$$

That is, we can allow for one period discount rates that are particular to the investor's class and also depend on the circumstances in which the income being discounted is received. Our analysis goes through without a hitch. In particular, (2) becomes

$$p_t(x_t) = \max_{a \in A} E^a [\gamma_{t+1}^a(x_{t+1})(d_{t+1}(x_{t+1}) + p_{t+1}(x_{t+1})) \mid x_t = x_t].$$

Such $\{\gamma_t^a\}$ must be taken as exogenous, so this extension does not give the general case of risk aversion. It does, however, allow us to propose a situation, without the risk neutrality and infinite wealth assumptions, for which our model is a good approximation. (The assumption of no short sales remains rigid.) Suppose that all investors are risk averse, expected utility (of consumption) maximizers; that the only other securities available for holding wealth are riskless one-period bonds (which are traded in a perfect market); and that all other prices and incomes are certain. Let $\gamma_{t+1}(x_t)$ be the equilibrium price at time t in contingency x_t of a one dollar bond. (In the spirit of our partial-equilibrium analysis, these are taken as exogenous.) Then if the classes of investors are large, we argue that investors will value streams of income derived from speculation in the stock at approximately their expected present worth, discounted at $\{\gamma_t(\cdot)\}$. At equilibrium, no investor will hold so large an amount of stock that his risk aversion is significant because there are many investors who share the most optimistic subjective belief, among whom the fixed amount of stock outstanding can be divided. The riskless bonds serve to bring everyone's marginal rate of wealth transferral from one period to the next into line with $\{\gamma_t(\cdot)\}$. The bonds also make unnecessary the infinite wealth assumption, since purchase of the stock can be financed

by the sale of bonds. (With appropriate independence assumptions, one can further relax the requirement that the only way to transfer wealth is via the riskless bonds and that other prices and incomes are certain.)

Notice that the short sales assumption is still crucial. If the markets for the stock were perfect, the amount of stock available to be held long would not be fixed, but would increase as members of less optimistic classes sold the stock short. Equilibrium will be reached only when investors take positions sufficiently disparate that their aversion to risk gives them identical "marginal beliefs."

VII. CONCLUDING REMARKS

We remarked in Section IV that the minimal consistent price scheme $\{p_i^*\}$ is in fact *uniquely* consistent in the finite horizon case. In general, $\{p_i^*\}$ is the only consistent price scheme that can be gotten as the limit of consistent price schemes for the natural sequence of approximating finite-horizon problems. Thus, we are inclined to say that $\{p_i^*\}$ is a uniquely reasonable price system. How would a fundamentalist react to this conclusion? The basic tenet of fundamentalism, which goes back at least to J. B. Williams (1938), is that a stock has an intrinsic value related to the dividends it will pay, since a stock is a share in some enterprise and dividends represent the income that the enterprise gains for its owners. In one sense, we think that our analysis is consistent with the fundamentalist spirit, tempered by a subjectivist view of probability. Beginning with the view that stock prices are created by investors, and recognizing that investors may form different opinions even when they have the same substantive information, we contend that there can be no objective intrinsic value for the stock. Instead, we propose that the relevant notion of intrinsic value is obtained through market aggregation of diverse investor assessments. There are fundamentalist overtones in this position, since it is the market aggregation of investor attitudes and beliefs about future *dividends* with which we start. Under our assumptions, however, the aggregation process eventually yields prices with some curious characteristics. In particular, investors attach a higher value to ownership of the stock than they do to ownership of the dividend stream that it generates, which is not an immediately palatable conclusion from a fundamentalist point of view.

Investors, if they are not to underprice the security, must take into account the beliefs, preferences, etc., of their fellow investors. Moreover, they must have some conception of how these diverse be-

liefs, etc., will be aggregated into future prices. Our treatment essentially skirts this very difficult problem through the assumption of perfect contingent foresight. If one drops this utopian assumption, and further introduces such a real-life phenomenon as privileged information, one gets a world in which investors must turn to public information, such as prices and trading volume, to discover what their fellow investors know and how they will react to incoming information. At the risk of gross overstatement, we suggest that this line of reasoning might lead to a "legitimate" theory of technical analysis.

Proponents of the efficient market hypothesis conclude that the rational portfolio strategy (in view of transaction costs and risk aversion) is to buy a well-diversified portfolio and hold it. Quite a different view of "rational" portfolio management emerges from our model. Consider again the example presented in Section III. Investors of class 1 cannot achieve an expected net present value of zero from stock bought in state one, and they can achieve this from stock bought in state zero if and only if they sell at or before the first time a one dollar dividend is declared. In the general model, investors can achieve an expected net present value of zero only from stock bought in certain circumstances and only if they follow certain selling strategies. (The strategy of selling after one period, which leads to much churning of the portfolio, always works.) The strategy of buying in favorable circumstances and holding for many periods typically yields an expected loss. In brief, all investors must actively manage their portfolios in order to expect a proper return.

If our model were broadened to include risk aversion and transaction costs, there would be forces at work that yield a notion of "rational" portfolio management more in accord with the efficient market hypothesis. But the conflict between our model and the hypothesis lies at a deeper, philosophical level. As we understand things, a subjectivist statement of the hypothesis implicitly requires the supposition that rational economic agents having the same information will arrive at the same subjective probability assessments. The hypothesis then says that, although investors do not originally have the same information, prices reflect their information in such a way as to induce actions identical to those they would undertake if they had complete information. (See Beja, 1976, for exposition of this view and some negative results on the ability of prices to perform this function.) Thus, the investors are effectively homogeneous in an "efficient" market, and speculation does not occur. (Some assumption about sufficiently similar risk attitudes is also necessary.) In our model, all investors have complete information from the outset, but still they

arrive at different subjective assessments. Speculation and active portfolio management follow inevitably.

STANFORD UNIVERSITY

NOTES

1. Or nonanticipatory, or a stopping time. See Chung (1974), p. 322.

REFERENCES

- Arrow, K. J., *Essays in the Theory of Risk-Bearing* (Chicago: Markham Publishing Company, 1968).
- Beja, A., "The Limited Information Efficiency of Market Processes," unpublished manuscript, 1976.
- Chung, K. L., *A First Course in Probability*, 2nd ed. (New York: Harcourt, Brace and World, 1974).
- Feiger, G., "What Is Speculation?" this *Journal*, XC (Nov. 1976), 677-87.
- Hirshleifer, J., "Foundations of the Theory of Speculation: Information, Risk, and Markets," this *Journal*, LXXXIX (Nov. 1975), 519-42.
- Keynes, J. M., *The General Theory* (New York: Harcourt, Brace and World, 1931).
- Radner, R., "Existence of Equilibrium Plans, Prices, and Price Expectations in a Sequence of Markets," *Econometrica*, XL, No. 2 (1972), 289-303.
- Williams, J. B., *The Theory of Investment Value* (Cambridge: Harvard University Press, 1938).