HERD BEHAVIOR AND INVESTMENT

by

David S. Scharfstein
Jeremy S. Stein

WP #   2062-88   August 1988

HERD BEHAVIOR AND INVESTMENT

by

David S. Scharfstein

Jeremy C. Stein*

First Draft: June 1988

This Draft: August 1988

*Sloan School of Management, MIT, and Harvard Business School, respectively. We thank Robert Gertner, Bengt Holmstrom, Chi-Fu Huang, Jim Poterba, Julio Rotemberg, Garth Saloner, participants in the NBER Summer Institute, and as always, Bob Gibbons, for helpful comments.

Herd Behavior and Investment

Abstract

This paper presents a theoretical model of herd behavior in investment. Managers mimic the investment decisions of other managers, in the process ignoring their private information about the attractiveness of various alternatives. Although this behavior is inefficient from a social standpoint, it is rational from the perspective of managers who are concerned about their reputations in the labor market. We discuss applications of the model to corporate investment, the stock market, and decisionmaking within firms.

I.   Introduction

A basic tenet of classical economic theory is that investment decisions reflect agents' rationally formed expectations; decisions are made using all available information in an efficient manner. A contrasting view is that investment is also driven by group psychology, which weakens the link between information and market outcomes. In The General Theory, Keynes (1936) expresses skepticism about the ability and inclination of "long-term investors" to buck market trends to ensure full efficiency. In his view, investors may be reluctant to act according to their own information and beliefs, fearing that their contrarian behavior will damage their reputations as sensible decision-makers:

> ". . . . it is the long-term investor, he who most promotes the public interest, who will in practice come in for most criticism, wherever investment funds are managed by committees or boards or banks. For it is in the essence of his behavior that he should be eccentric, unconventional and rash in the eyes of average opinion. If he is successful, that will only confirm the general belief in his rashness; and if in the short-run he is unsuccessful, which is very likely, he will not receive much mercy. Worldly wisdom teaches that it is better for reputation to fail conventionally than to succeed unconventionally."

Thus Keynes suggests that professional managers will "follow the herd" if they are concerned about how others will assess their ability to make sound judgements. There are a number of settings in which this kind of herd behavior might have important implications. One example is the stock market, for which the following explanation of the pre-October 1987 bull market is often repeated: The consensus among professional money managers was that price levels were too high--the market was, in their opinion, more likely to go down rather than up. However, few money managers were eager to sell their equity holdings. If the market did continue to go up, they

were afraid of being perceived as lone fools for missing out on the ride. On the other hand, in the more likely event of a market decline, there would be comfort in numbers--how bad could they look if everybody else had suffered the same fate?

The same principle can apply to corporate investment, when a number of companies are investing in similar assets. In Selling Money, Gwynne (1986) documents problems of herd behavior in banks' lending policies towards LDC's. Discussing the incentives facing a credit analyst, he writes:

> "Part of Herrick's job--an extremely important part as far as the bank was concerned--was to retrieve information about the countries in which the bank did business. But this function collided head on with what Herrick was actually doing out there...His job would never be measured by how correct his country risk analysis was. At the very least, Herrick was simply doing what hundreds of other larger international banks had already done, and any ultimate blame for poor forecasting would be shared by tens of thousands of bankers around the world; this was one of the curious benefits of following the herd."

This paper presents a theoretical model that is consistent with the kind of herd behavior discussed above. Managers mimic the investment decisions of other managers, in the process ignoring their private information about the attractiveness of various alternatives. Although this behavior is inefficient from a social standpoint, it is rational from the perspective of managers who are concerned about their reputations in the labor market.

Our model is a "learning" model, similar in spirit to one studied by Holmstrom (1982). Like us, he considers a situation in which managers use investment decisions to manipulate the labor market's inferences regarding their ability, where ability represents an aptitude for making decisions. This definition of ability contrasts with one where ability adds to physical productivity, as in Holmstrom and Ricart i Costa (1986) as well as

in another part of Holmstrom (1982). The key difference between our model and these others is that ours is most interesting when there is more than one manager, whereas theirs are single-manager models.[1]

We assume that there are two types of managers: "smart" ones who receive informative signals about the value of an investment, and "dumb" ones, who receive purely noisy signals. Initially, neither the managers themselves nor the labor market can identify the types. However, after the managers have made an investment decision, the labor market can update its beliefs, based on two pieces of evidence: 1) whether the manager made a "good" investment (i.e., whether he made money); 2) whether the manager's behavior was similar to or different than that of other managers.

The first piece of evidence will not be used exclusively, since success is only a noisy indicator of ability--on any given draw, all smart managers could get unlucky and receive misleading signals. Hence the second piece of evidence is important as well. The key point is that smart managers will tend to receive correlated signals (since they are all observing a piece of the same "truth"), while dumb ones will not (they simply observe uncorrelated noise). Consequently, if one manager mimics the behavior of others, this will suggest to the labor market that he has received a signal that is correlated with theirs, and is more likely to be smart. In contrast, a manager who takes a contrarian position will (all else being equal) be perceived as more likely to be dumb. Thus even if a manager's private information tells him that an investment has a negative expected value, he may pursue it if others before him have. Conversely, he may refuse investments that he perceives as having positive expected value if others before him have also done so.

These points are further developed in the four sections following this one. In Section II, we present the structure and assumptions of the model. The basic results on the existence of herding equilibria are developed in three propositions in Section III. In Section IV, we elaborate on the implications of our model for corporate investment, the stock market, and decisionmaking within firms. Finally, Section V contains concluding remarks.

II. The Model

The model that is developed in this section applies more literally to the example of corporate investment discussed above than it does to the stock market. We assume that the investments under consideration are available in perfectly elastic supply at a given price. This allows us to avoid explicitly considering the feedback from investment demand to prices, thereby simplifying the analysis considerably. In Section IV, we will discuss at greater length how we think our results carry over to the stock market, where this assumption is clearly not appropriate. For the time being, however, it may help for the reader interested in concreteness to bear in mind the following story: Our "managers" are in charge of capital investment at industrial firms, and they are each considering expanding capacity by building a similar type of plant. The value of the plant will depend on the state of the economy in the future, and each manager receives a signal that gives him some information about this future state. The question we address is the following: how well does the aggregate level of investment reflect all of the available information?

a) **Timing and Information Structure**

The economy consists of two firms, firm A and firm B, run by managers we call A and B respectively. These managers invest sequentially, with A moving first. At date 1, manager A decides whether or not to build the plant. There are two possible states for the economy at date 3: either the economy will be in the "high" state, in which case the plant will yield a profit (net of investment expense and discounting) of $x_H > 0$; or the economy will be in the "low" state, in which case the plant will yield a net profit of $x_L < 0$. The prior probabilities for these two states are α and $(1 - \alpha)$ respectively. The state of the economy is publicly obervable, even if neither manager decides to invest.

In making his decision, manager A has access to a signal, which can take on one of two values: s_G (a "good" signal); or s_B (a "bad" signal). Interpreting this signal is a bit complicated, because the manager does not know if he is "smart" or "dumb." If he is smart, which occurs with prior probability θ, the signal is informative--that is, a good signal is more likely to occur prior to the high state than prior to the low state. Formally, we have:

(1) Prob $(s_G | x_H, \text{smart}) = p;$

(2) Prob $(s_G | x_L, \text{smart}) = q < p.$

If the manager is dumb, however, which occurs with probability $(1 - \theta)$, he receives completely uninformative signals--he is as likely to receive s_G prior to the high state as prior to the low state, so that:

(3) Prob $(s_G | x_H, \text{dumb}) = \text{Prob}(s_G | x_L, \text{dumb}) = z.$

We make the assumption that the ex-ante distribution of signals is the same for both smart and dumb managers--both are equally likely to receive s_G, so that the actual signal received does not communicate any information about types by itself. (In other words, the only difference is that s_G really means something for smart managers--but they don't receive it more often.) This amounts to assuming that: Prob $(s_G|\text{smart})$ = Prob $(s_G|\text{dumb})$, or:

(4) $\quad z = \alpha p + (1 - \alpha)q.$

Given that the manager does not know if he is smart or dumb, a straightforward application of Bayes' law allows us to calculate the probabilities he attaches to the high state after receiving the good and bad signals:

(5) $\quad \text{Prob } (x_H|s_G) = \mu_G = \dfrac{[\theta p + (1 - \theta)z]}{z}\, \alpha;$

(6) $\quad \text{Prob } (x_H|s_B) = \mu_B = \dfrac{[\theta(1 - p) + (1 - \theta)(1 - z)]}{(1 - z)}\, \alpha.$

In order to make the investment problem interesting, we assume that the investment is attractive if a good signal has been received, but not if a bad signal has:

(7) $\quad \mu_G x_H + (1 - \mu_G)x_L > 0 > \mu_B x_H + (1 - \mu_B)x_L$

After manager A has made his investment decision at date 1, manager B makes his decision at date 2. Manager B also has access to a private signal. In addition, he can also observe whether or not firm A has decided to invest. This is valuable information--even in an ideal world (i.e., one

without reputational concerns), firm B's investment decision should be partially influenced by what firm A does. Our main point is that with reputational concerns, firm B's manager pays <u>too much</u> attention to what firm A has done, and too little to his private signal.[2]

Like manager A, firm B's manager can be either smart or dumb. If one manager is smart and the other is dumb, their signals are drawn independently from the binomial distributions given in equations (1) - (3). Similarly, if both are dumb, their signals are drawn independently--so that, for example, the probability that two dumb managers both observe s_G is z^2.

However, if both managers are smart, they are assumed to observe <u>exactly the same signal</u>. Thus the probability that two smart managers both observe s_G when the true state is x_H is p. (As opposed to the p^2 that would prevail if smart managers received independent draws from the distributions described in (1) and (2)). This feature is crucial to our analysis. It can be generalized somewhat, to allow the draws to be imperfectly correlated. However, if the signals of smart managers are drawn independently from the distributions, our results will fail to go through.

Heuristically, we require smart managers' prediction errors to be at least partially correlated with each other. Although this feature may seem slightly unnatural given the current set-up and notation, it amounts to nothing more than saying that there are <u>systematically unpredictable factors affecting the future state that nobody can know anything about</u>. For example, we might model the outcome of the state draw as being driven by the sum of two random variables, u and v. If u + v > 0, then x_H obtains. If u + v < 0, then x_L obtains. Our assumption is equivalent to

allowing smart managers to observe u but not v. (As suggested above, we could generalize to allow each manager to observe u perturbed by independent noise, so long as we retained the unobservability of the common component v.)

The importance of a common component to the prediction error for smart managers follows logic similar to that seen in tournament models of effort elicitation (e.g. Nalebuff and Stiglitz (1983)). If prediction errors are independent, the labor market can, in our model, efficiently update on ability using only individual performance--i.e., whether the manager picked a successful investment. Analogously, in a tournaments model, if the agents face independent shocks to output, optimal contracts evaluate the agent based only on individual absolute performance. However, to the extent that errors are correlated, there will be an informational gain from comparing agents, and looking at relative performance. In our model, one will not wish to evaluate too harshly a manager who picks a bad investment, if his colleague's similar choice suggests that they were both victims of a completely unpredictable factor.

Thus it is the common component to prediction errors that gives this model its bite, by causing some inferential weight to be placed on the similarity of managers' decisions. Perversely, the existence this extra channel of inference actually leads to ex-ante reductions in efficiency--just the opposite of the result seen in the literature on tournaments. This is because here, managers will attempt to actively manipulate their investment decisions in such a way as to bias the inference process in their favor. Even if the market recognizes that they will be engaging in this manipulation, it will continue to exist in equilibrium.[3]

b)   Managerial Objectives

Our next step is to specify the objective functions of the managers. In a first-best world, the managers would seek only to maximize the expected returns on investment, and would invest anytime their information (either from their private signal or from observing the other managers) indicated that investing had positive expected value.

In our model, managers' investment decisions enable the labor market to update its beliefs about their ability. We denote by $\hat{\theta}$ the market's revised assessment of the probability that a manager is smart.

In order to establish a simple relationship between managers' objectives and $\hat{\theta}$, we make several simplifying assumptions. Following Holmstrom and Ricart i Costa (1986), we assume that: 1) the investment game will be replayed once more after date 3; 2) at this point there will be no further reason to build a reputation, so managers invest efficiently; 3) competition leads managers' spot market wages to be set to the economic value of their ability.

It is straightforward to demonstrate that, for a wide range of parameter values in our model, the expected return on the investment opportunity is linear in the manager's ability (as measured by $\hat{\theta}$) if the manager invests efficiently. Hence the spot market wages referred to above will be proportional to $\hat{\theta}$.[4]

We do not explicitly analyze contracting behavior in what follows. Rather, we employ the "no slavery" assumption (used by Holmstrom (1982) and others) to claim that managers cannot be bound to their firms against their will ex-post, so that any long-term contract that would pay some types less than spot market wages in the second go-round of the investment game is infeasible.

Since their future wages are thus linear in $\hat{\theta}$, managers will have some incentive to generate high values of $\hat{\theta}$, rather than to invest efficiently in the first go-round. Of course, it is still possible that short-term incentive contracts could serve at least partially to align managerial and firm interests, by specifying a profit-contingent wage in the first go-round of the investment game. Thus, in principle, it seems reasonable to believe that managers would be induced to act so as to maximize some weighted average of expected profits and their reputations. However, this more general formulation leads to the same basic conclusions that obtain if managers care only about reputation--although naturally, more weight on expected profits will tend to attenuate the inefficiencies. For the sake of starkness and notational simplicity, we leave expected profits out of the managerial objective function. Later, in discussing our results we will briefly touch on how they would be altered by adding expected profits back into managerial consideration.[5]

The last assumption we make is that managers are risk neutral, so that their objective function simplifies to maximizing expected wages. This is equivalent to maximizing the expected value of $\hat{\theta}$.

III. Herding Equilibria

a) Comparison with Efficient Investment Decisions

In order to economize on notation, we set $p = 1 - q$. We also set $\alpha = \frac{1}{2}$. Taken together, these simplifications imply that $z = \frac{1}{2}$ from equation (4). For the time being, we leave the sign of $(x_H + x_L)$ unspecified, so that the investment can have an ex-ante expected value that is either positive or negative.

As a benchmark, we first derive the optimal decision rules in a first-best world with no reputational concerns. Manager A would invest if and only if he observed s_G--this is a consequence of the assumption in equation (7). Thus manager B can infer manager A's signal from his investment decision.

If manager B observes s_B after firm A has invested, he knows that the total information set is (s_G, s_B). Given our symmetry assumptions this implies that the probability of the high state, $\text{prob}(x_H|s_G, s_B) = \frac{1}{2}$. Thus the investment decision will hinge on the sign of $(x_H + x_L)$--if this quantity is positive, manager B will invest, and if not, he won't.

Similarly, if firm A doesn't invest, and manager B observes s_G, the investment decision turns on the same criterion of whether $(x_H + x_L) > 0$. Clearly, in the first best, the order in which the information arrives is irrelevant to manager B's decision. If one manager observes s_G and the other sees s_B, manager B's decision will be the same regardless of whether the s_G signal was received by him or by manager A.

With reputational considerations, the decision rules will be different. When manager A observes s_G and invests, manager B will also invest, regardless of his signal and regardless of the sign of $(x_H + x_L)$. Hence if this signal is s_B and $x_H + x_L < 0$, the investment will be inefficient. Conversely, if firm A doesn't invest, firm B never will either, which is inefficient when manager B observes s_G and $x_H + x_L > 0$. Now the order in which information arrives is important to firm B's decision--the same total information set of (s_G, s_B) can lead it to invest or not to invest, depending on whether the signal s_G is received by the first mover firm A or not.

b)   Equilibria with Reputational Concerns

We now examine the equilibria that exist when managers seek to maximize the expected value of $\hat{\theta}$. For now, we focus on the decision rules for the manager of firm B--all the "continuation" equilibria that we look at will have the manager of firm A behaving efficiently, by investing if and only if he observes s_G. Later, we will establish that this efficient behavior by firm A's manager is part of an equilibrium for the overall game.

We develop our results through a series of propositions.

Proposition 1:   There does not exist any continuation equilibrium in which manager B's investment decision depends on the signal he observes. Thus the only possible equilibria are those where manager B mimics manager A regardless of the signal, or where manager B does the opposite of manager A regardless of the signal.

The proof will be by contradiction. We start by conjecturing the existence of the "separating" equilibrium described above.[6] We then determine the updating rules the labor market would use to calculate $\hat{\theta}$ in such an equilibrium. Finally, we show that given these updating rules, rational managers will not wish to behave as posited in the equilibrium.

The revised ability assessments will be a function of the labor market's conjectures about the signals observed by the managers as well as the realized state of the world. Of course, only the managers observe their signals, but in the putative separating equilibrium, there is a one-to-one mapping from signals to actions. Thus for example, suppose the separating equilibrium calls for each manager to invest if and only if he

observes s_G. Then if manager A doesn't invest and manager B does, the market believes that manager A observed s_B and manager B observed s_G, and it can do its updating based on these beliefs.

As noted above, the main focus of our analysis is manager B. Continuing with the above example, suppose the high state was realized. How would the market revise its prior about the manager's ability? Let (s_B, s_G, x_H) denote this event, and let $\hat{\theta}(s_B, s_G, x_H)$ be the revised prior. By Bayes' rule, one can show that:

$$(8) \quad \hat{\theta}(s_B, s_G, x_H) = \frac{\frac{1}{2}p\theta(1 - \theta)}{\frac{1}{2}p\theta(1 - \theta) + \frac{1}{4}(1 - p)\theta(1 - \theta) + \frac{1}{4}(1 - \theta)^2}$$

$$= 2\theta p/(1 + \theta)$$

The explanation for this result is as follows. There are three possible configurations of managerial ability that could give rise to this event: (dumb, smart); (smart, dumb); and (dumb, dumb). Note that (smart, smart) is not possible since if this were the case, both managers would have received the same signal, by virtue of our assumption that smart managers' signals are perfectly correlated.

We wish to know the probability of (dumb, smart) conditional on the event (s_B, s_G, x_H). If the configuration of talent is (dumb, smart), which occurs with ex-ante probability $\theta(1 - \theta)$, the probability of (s_B, s_G) in state x_H is $\frac{1}{2}p$. This explains the numerator in (8). The denominator gives in addition the probabilities of $(s_B, s_G|x_H)$ if the configuration is (smart, dumb) and (dumb, dumb). These are $\frac{1}{2}(1 - p)$ and $\frac{1}{4}$ respectively. By symmetry, it is straightforward to show that $\hat{\theta}(s_G, s_B, x_L) = 2\theta p/(1 + \theta)$ also.

The derivation of the other updating rules follow along similar lines. They are listed below:

(9) $\hat{\theta}(s_B, s_G, x_L) = \hat{\theta}(s_G, s_B, x_H) = 2\theta(1 - p)/(1 + \theta)$;

(10) $\hat{\theta}(s_B, s_B, x_H) = \hat{\theta}(s_G, s_G, x_L) = \dfrac{2\theta(1 - p)(1 + \theta)}{4\theta(1 - p) + (1 - \theta)^2}$;

(11) $\hat{\theta}(s_B, s_B, x_L) = \hat{\theta}(s_G, s_G, x_H) = \dfrac{2\theta p(1 + \theta)}{4\theta p + (1 - \theta)^2}$.

Now, in order for our posited equilibrium to actually hold together, it must be that managers find it in their interest to behave as assumed. For example, suppose that manager A has observed s_B and has not invested. Given the updating rules above, can it ever be rational for manager B to invest upon observing s_G, but not invest upon observing s_B? Or will one type of manager "break" the equilibrium by deviating and attempting to fool the market into thinking that he has received a different signal?

To answer these questions, the following probabilities must be calculated:

(12) Prob $(x_H | s_B, s_G) = \frac{1}{2}$;

(13) Prob $(x_H | s_B, s_B) = \dfrac{4\theta(1 - p) + (1 - \theta)^2}{4\theta + 2(1 - \theta)^2}$.

We can now check the rationality conditions that must hold for the equilibrium to be viable. One of these is:

(14) $\hat{\theta}(s_B, s_G, x_H)$ Prob $(x_H | s_B, s_G) + \hat{\theta}(s_B, s_G, x_L)$ Prob $(x_L | s_B, s_G)$

$\geq \hat{\theta}(s_B, s_B, x_H)$ Prob $(x_H | s_B, s_G) + \hat{\theta}(s_B, s_B, x_L)$ Prob $(x_L | s_B, s_G)$.

Inequality (14) represents the requirement that if manager B receives signal s_G, he prefers to invest (and identify himself as someone who had observed s_G), rather than not invest (and masquerade as someone who had received signal s_B). Direct substitution from equations (8) - (12) establishes that the inequality is violated--if manager B receives s_G, he will wish to deviate by mimicking firm A and not investing. A symmetric argument establishes that if firm A <u>has</u> invested, there is also no separating equilibrium. In this case, if manager B observes s_B, he will deviate by mimicking firm A and also investing.

There are also potentially "perverse" separating equilibria, where a manager B invests if and only if he observes s_B, rather than s_G. It can be easily demonstrated using the same lines of reasoning that such equilibria are also not viable. This completes the proof of Proposition 1.

It is worth examining the updating rules in equations (8) - (11) to gain some intuition for the forces that prevent the efficient equilibrium from existing. Two main points emerge from these equations:

First, $\hat{\theta}(s_B, s_G, x_H) > \hat{\theta}(s_B, s_G, x_L)$; and $\hat{\theta}(s_G, s_G, x_H) > \hat{\theta}(s_G, s_G, x_L)$. Holding the investment decision of manager A fixed, manager B is indeed compensated for making "absolutely" good decisions--for investing prior to a realization of x_H, as opposed to investing prior to a realization of x_L.

Second, however, the investment decision of manager A does have an important externality effect. Holding the correctness of the investment decision fixed, there is a higher payoff to manager B for imitating manager A. That is, $\hat{\theta}(s_G, s_G, x_H) > \hat{\theta}(s_B, s_G, x_H)$; and $\hat{\theta}(s_G, s_G, x_L) > \hat{\theta}(s_B, s_G, x_L)$.

Proposition 1 is a direct consequence of this second effect. Because of the payoff to imitation, even if the new information makes it more likely that contradicting manager A will be the economically correct decision, manager B will prefer to mimic A. As a result, decisions cannot be made contingent on signals, and there cannot be an equilibrium where manager B takes advantage of his private information.

As was emphasized in the previous section, the result depends on our assumption that prediction errors are correlated across smart managers. If the signals of smart managers are independent, Proposition 1 no longer holds. This can be demonstrated by calculating the updating rules that would prevail in a world of independent signals. Denoting these rules by $\hat{\theta}^i(\ )$, and using the same Bayesian logic as before, we can derive:

(15) $\hat{\theta}^i(s_B, s_G, x_H) = \hat{\theta}^i(s_G, s_G, x_H) = \hat{\theta}^i(s_B, s_B, x_L) = \hat{\theta}^i(s_G, s_B, x_L)$

$$= \frac{2\theta p}{2\theta p + (1 - \theta)};$$

(16) $\hat{\theta}^i(s_B, s_G, x_L) = \hat{\theta}^i(s_G, s_G, x_L) = \hat{\theta}^i(s_B, s_B, x_H) = \hat{\theta}^i(s_G, s_B, x_H)$

$$= \frac{2\theta(1 - p)}{2\theta(1 - p) + (1 - \theta)}.$$

According to equations (15) and (16), in the case of independent signals, the labor market's assessment of firm B's manager is unrelated to the investment decision of firm A. All that matters is absolute correctness--investing before state x_H leads to a more favorable ability assessment than not investing before state x_H. As a result, the inequality in (14) is satisfied (with equality) and it is possible to sustain the efficient equilibrium in which manager B's investment decisions can depend on his private signal.

If the independence assumption is relaxed, the updating rules become a function of firm A's investment decision, inequality (14) is violated, and Proposition 1 holds. Thus perfect correlation of smart manager signals is not necessary for our results--all that is needed is some correlation of prediction errors.

Having established that continuation equilibria with signal-contingent decisions by manager B do not exist, we now turn our attention to the equilibria that can be supported in our model.

Proposition 2: There exists a continuation equilibrium in which manager B always mimics manager A, investing if and only if A does. This herding equilibrium is supported by the following "reasonable" out of equilibrium beliefs: i) if manager B deviates by investing when A has not, the labor market conjectures that he observed s_G; and conversely, ii) if manager B deviates by not investing when A has, the labor market conjectures that he observed s_B.

In order to prove Proposition 2, it is necessary to show that manager B will always find it optimal to behave as prescribed, given the beliefs

posited. Let us consider only the case where firm A has already not invested; the other case works exactly the same way.

If manager B follows firm A by also not investing, his revised ability assessment will be simply equal to θ--there is no revision from the prior because the equilibrium is a "pooling" one, with both s_G and s_B recipients choosing the same action.

In order for a manager B who observes s_B to not deviate, it must be the case that:

(17) $\theta \geq \hat{\theta}(s_B, s_G, x_H)$ Prob $(x_H | s_B, s_B) + \hat{\theta}(s_B, s_G, x_L)$ Prob $(x_L | s_B, s_B)$

Inequality (17) is the requirement that the payoff θ to pooling exceed the payoff to an s_B recipient from deviating and investing, given that out of equilibrium beliefs are such that he will be taken to be an s_G recipient if he deviates. Direct substitution from (8), (9) and (13) verifies that the inequality is satisfied.

In order for a manager B observing s_G to not deviate, the following must hold:

(18) $\theta \geq \hat{\theta}(s_B, s_G, x_H)$ Prob $(x_H | s_B, s_G) + \hat{\theta}(s_B, s_G, x_L)$ Prob $(x_L | s_B, s_G)$

Comparison of (17) and (18) shows that it is relatively more tempting for manager B to deviate by investing after observing s_G, as opposed to s_B. (It is in this sense that the out-of-equilibrium conjecture that a deviator has seen s_G is "reasonable.") Nonetheless, (18) reduces to:

(18') $\theta \geq \theta/(1 + \theta)$;

which is clearly satisfied. Thus Proposition 2 is proved, and we have established the existence of a herding continuation equilibrium.

It should be pointed out that there is another, perverse continuation equilibrium in which the decisions of manager B do not depend on his signal. In this equilibrium, manager B always contradicts manager A, investing if and only if A has not. This equilibrium can only be supported by the following "unreasonable" beliefs off the equilibrium path: if manager B deviates by investing when the equilibrium calls for him not to, it is because he has observed s_B; and if he deviates by not investing when the equilibrium requires investment, it is because he has observed s_G.

There are a number of reasons why this equilibrium is not likely to be economically relevant. The first, noted above, is the unreasonableness of the out-of-equilibrium beliefs, where investing is associated with the signal s_B and not investing is associated with s_G. A second is the inefficiency of the contradiction equilibrium relative to the herding equilibrium--while manager B's information is not used in either, the herding equilibrium at least has firm B behaving correctly based on the information in firm A's actions.

The multiplicity of equilibria stems from our assumption that investment decisions do not directly affect the manager's utility, but are nothing more than a means of conveying information to the labor market about the signal that the manager has observed. If the labor market (perversely) interprets investment to mean that the manager has observed s_B, then the manager may well invest if he wishes to convince the labor market that he has seen s_B. It follows that in the current formulation of our model, we cannot pin down exactly what actions will be taken. However,

we can pin down how much information is revealed in equilibrium: in both equilibria, managers' actions do not depend on their private signals and hence convey no information.[7]

This reasoning suggests that the model can be altered slightly so as to leave the herding equilibrium as the unique outcome. Suppose that managers do not invest directly; instead, they report their signals to the "owners" of the firm, who then make investment decisions so as to maximize profits. Proposition 1 then can be interpreted as saying that there is no equilibrium where manager B can be relied on to make informative reports. Given that he cannot learn anything from his manager, the owner of firm B will then have to rely on the only available information, the action of firm A. The unique profit-maximizing decision for the owner of firm B is thus to always mimic firm A.

In sum, then, the contradiction equilibrium is probably not a sensible one. If one dismisses it, the herding equilibrium is left as the unique continuation equilibrium of the game. It remains only to establish that the efficient behavior on the part of manager A that has been assumed to this point is part of an overall equilibrium.

Proposition 3: There exists an equilibrium for the overall game where manager A invests if and only if he receives s_G, and where manager B always mimics manager A regardless of B's signal.

In the proposed equilibrium, there is no information inherent in manager B's actions, since he always does the same thing, regardless of his signal. Thus manager A can only be evaluated absolutely--his revised

ability $\hat{\theta}$ can be a function of just his action and the realized state. The updating rules are therefore identical to those given for the two manager, independent signal case in equations (15) and (16). That is:

(19) $\hat{\theta}^A(s_G, x_H) = \hat{\theta}^A(s_B, x_L) = 2\theta p/(2\theta p + (1 - \theta))$; and

(20) $\hat{\theta}^A(s_G, x_L) = \hat{\theta}^A(s_B, x_H) = 2\theta(1 - p)/(2\theta(1 - p) + (1 - \theta))$.

In order for manager A to be willing to invest after observing s_G and not invest after observing s_B, the following two conditions must hold:

(21) $\hat{\theta}^A(s_G, x_H)\mu_G + \hat{\theta}^A(s_G, x_L)(1 - \mu_G) \geq \hat{\theta}^A(s_B, x_H)\mu_G + \hat{\theta}^A(s_B, x_L)(1 - \mu_G)$;

(22) $\hat{\theta}^A(s_B, x_H)\mu_B + \hat{\theta}^A(s_B, x_L)(1 - \mu_B) \geq \hat{\theta}^A(s_G, x_H)\mu_B + \hat{\theta}^A(s_G, x_L)(1 - \mu_B)$.

From equations (5) and (6), we can obtain the simplification: $\mu_G = (1 - \mu_B) = \theta p + \frac{1}{2}(1 - \theta) > \frac{1}{2}$. The inequalities can then be verified, which proves that manager A's behavior is part of an equilibrium.

c) <u>Discussion</u>

The herding equilibrium derived above will be inefficient relative to the first-best for some configurations of private information, except in the knife-edge case where $x_H + x_L = 0$. For example, if $x_H + x_L > 0$, then the equilibrium will be inefficient when manager A observes s_B, manager B observes s_G; and neither firm invests.

If managers place some weight on expected profits in their objective functions, these inefficiencies can disappear, depending on the value of $x_H + x_L$. Continuing with the above example, for any fixed weight on

profits, if $x_H + x_L$ is large enough, the efficient signal-dependent equilibrium will exist. On the other hand, if $x_H + x_L$ is small enough (but still above zero) even a very large weight on profits will not be sufficient to restore the efficient equilibrium. In this sense, if managers care some about expected profits, the more egregious inefficiencies associated with herd behavior can be alleviated, though not eliminated in all situations.

It should be emphasized that our results concerning herd behavior generalize to any number of managers. The logic behind this fact can be seen by considering a third manager C, who moves after A and B, but before the state of the economy is realized. Since we have established that in a two-manager herding equilibrium, manager B's actions are independent of his signals, manager C learns nothing from observing what firm B has done. He learns only from observing firm A, where actions do depend on the signal received. Thus manager C is in exactly the same position as manager B before him, and the same arguments can be used to show that he too cannot make his actions contingent on his private information. This line of reasoning can be applied repeatedly to any number of subsequent managers.

A final point to note is that herd behavior can have adverse effects on the incentives for information acquisition. Although it was not explicitly included in our model, one can imagine a preliminary stage where managers have to expend effort to become informed. Since herding considerations will diminish the value of any information acquired, there will be less motivation to gather information in the first place.

IV. Implications of the Model

The theoretical model developed above has implications in a number of different areas. In order to give a feeling for some of the potential applications, we discuss a few examples.

a) Corporate Investment

Bank lending to LDC's was mentioned earlier as an apparent instance of herd behavior in corporate investment. A recent paper by Morck, Shleifer and Vishny (1988) suggests that the problem may be more widespread. They study the effectiveness of boards of directors in dealing with poorly managed firms. Their principal empirical finding is that top management firings are primarily associated with poor performance of a firm relative to its industry, rather than with industry-wide failures. They interpret these results as evidence that boards have a difficult time assigning blame to their managers for mistaken strategies, when other firms in the industry are following similar strategies.

Excessive investment in businesses that should be shrinking is a commonly cited example of such an industry-wide mistake -- for instance, in the oil industry, managers kept up costly explorations for additional reserves in spite of falling oil prices (see, e.g., Jensen (1986)). The Morck-Shleifer-Vishny evidence, which implies that managerial talent assessment seems to be done on a relative basis, provides a plausible story for how such ill-conceived investment plans can be sustained across an entire industry. Any given manager will, even if he feels the investment to be excessive, be more positively disposed towards it when all his colleagues are doing the same thing.

b)   <u>The Stock Market</u>

Herd behavior by money managers could provide a partial explanation for excessive stock market volatility. By mimicking the behavior of others (i.e., buying when others are buying, and selling when others are selling) rather than responding to their private information, members of a herd will tend to amplify exogenous stock price shocks. In a sense, the ideas developed here can be thought of as providing the "microfoundations" for stock market phenomena that are often thought to stem from psychological sources such as "groupthink," mass euphoria or panic, etc.

It should be pointed out, however, that the model of this paper does not fit perfectly into a stock market setting, due to the assumption of perfectly elastic supply and the consequent lack of a market clearing price. Adding pricing considerations would complicate the formal analysis considerably. Nonetheless, we think that our basic insights do carry over to the stock market. At any given level of prices, major money managers are likely to have an idea about the extent to which their competitors are "in" the market. If this is the case, there is the possibility that money managers will mimic each others' asset allocation strategies -- upon observing that manager A has 50% of his assets in stocks and 50% in bonds, manager B may aim for a similar portfolio composition, even when his private information suggests that current price levels are too low or too high. Thus one testable implication of our model is that the asset allocation decisions of professional money managers should be more closely correlated over time than the decisions of equally active private investors who are unconcerned about their reputations.

Shiller and Pound (1986) present some evidence that can be viewed as consistent with the existence of herd behavior in the stock market. They surveyed institutional investors to determine the factors that went into

their decision to buy a particular stock. Purchase of stocks that had recently had large price runups tended to be motivated by the advice of others (other investment professionals, newsletters, etc.). This contrasted with more stable stocks, where fundamental research (a systematic search procedure for a security with certain characteristics) played a more important role. This seems to suggest that the comfort inherent in following common wisdom can lead professional money managers to invest in stocks where fundamentals might dictate otherwise.

c) **Decisionmaking within Firms**

Recent work on the theory of the firm by Sah and Stiglitz (1985, 1986) has emphasized the role of managers as information filters. They argue that firms may organize themselves internally in such a way as to take maximum advantage of the fact that different managers will tend to have imperfectly correlated errors of judgment. Thus there may be benefits to having decisions made by committees, or through a vertical chain of command where projects can be rejected at various points along the way.

Our model points up certain limitations that may be inherent in group decisionmaking, and also offers some new insights about how organizational structure can facilitate the decisionmaking process. As a stylized example, consider the case of a capital budgeting committee meeting, where the gathered managers are supposed to vote in turn on a proposed investment project. Ideally, the point of having several managers vote is to gather a wide range of information. Unfortunately, if career concerns are present, this may not work well. Once the first manager has voted, the others may simply echo his choice, regardless of their private beliefs. Thus a false consensus is achieved, and the information of the other managers is wasted.

One way to mitigate this problem would be to have those with the stronger career concerns vote first. If the committee consists of young and old executives, the young ones should be asked to voice their opinions before the old ones, since there is presumably more uncertainty about their ability and hence a greater temptation for herd behavior. Old executives, on the other hand, should be more willing to express contrarian views. More generally, this line of reasoning implies an advantage to a "bottom up," rather than "top down" organization of information flow within a firm. To the extent that new ideas or project suggestions can be passed upstream for approval, this may result in better decisionmaking than if the ideas are originated at a high strategic planning level and then are passed downstream for line manager input.

Our model is consistent with observed differences in the decisionmaking processes of Japanese and U.S. firms. Abegglen and Stalk (1985) document the following three facts: (i) labor mobility is lower in Japan than in the U.S.; (ii) earnings differentials between the highest and lowest paid employees are lower in Japan; and (iii) decisionmaking is more participative in Japan--more employees are consulted in the process. According to our model, the first two observations should imply the third: to the extent that reputational concerns in the labor market are not as important in Japan, valuable information can be more readily elicited from all employees.[8]

V. Conclusions

Herd behavior can arise in a variety of contexts, as a consequence of rational attempts by managers to enhance their reputations as decisionmakers. This behavior introduces an element of extrinsic arbitrariness (or "animal spirits") to market outcomes. As we have seen,

these outcomes can depend critically on the order in which managers with different pieces of information make their investment decisions. Since this ordering is completely an artifice of the model, the fundamental economic conclusion is that these are effectively multiple equilibria -- there can be more than one outcome associated with a given economy-wide information set.

Other recent papers have also examined models with the Keynesian feature of multiple equilibria in investment -- see, e.g., Weitzman (1982), Shleifer (1986), and Shleifer and Vishny (1988). These authors emphasize the roles of increasing returns and/or aggregate demand spillovers as sources of multiplicity. In contrast, we find that there may be indeterminacy in investment even in the absence of such factors.

References

Abegglen, James and Stalk, George, <u>Kaisha: The Japanese Corporation</u>. New York: Norton, 1985.

Crawford, Vincent and Sobel, Joel, "Strategic Information Transmission," <u>Econometrica</u>, 50, (1982), 1431-1451.

Fudenberg, Drew and Tirole, Jean, "A Signal-Jamming Theory of Predation," <u>Rand Journal of Economics</u> 17 (1986), 366-376.

Gwynne, Samuel, <u>Selling Money</u>. New York: Weidenfeld and Nicolson, 1986.

Hart, Oliver and Holmstrom, Bengt, "The Theory of Contracts," in Truman Bewley, ed., <u>Advances in Economic Theory, Fifth World Congress</u>, Cambridge University Press, 1987.

Hellwig, Martin, "On the Aggregation of Information in Competitive Markets," <u>Journal of Economic Theory</u> 22 (1980), 477-498.

Holmstrom, Bengt, "Managerial Incentive Problems: A Dynamic Perspective," in <u>Essays in Economics and Management in Honor of Lars Wahlbeck</u> Helsinki: Swedish School of Economics, 1982.

_____ and Ricart i Costa, Joan, "Managerial Incentives and Capital Management," <u>Quarterly Journal of Economics</u> 101 (1986), 835-860.

Jensen, Michael, "Agency Costs of Free Cash Flow, Corporate Finance and Takeovers," _American Economic Review Papers and Proceedings_ 76 (1986) 323-329.

Keynes, John Maynard, _The General Theory of Employment, Interest and Money_ London: Macmillan, 1936.

Milgrom, Paul and Roberts, John, "An Economic Approach to Influence Activities in Organizations," _American Journal of Sociology_, 94, (1988) S154-S179.

Morck, Randall, Shleifer, Andrei and Vishny, Robert, "Alternative Mechanisms for Corporate Control," mimeo 1988.

Nalebuff, Barry and Stiglitz, Joseph, "Information, Competition and Markets," _American Economic Review Papers and Proceedings_ 73 (1983) 278-283.

Sah, Raaj Kumar and Stiglitz, Joseph, "Human Fallibility and Economic Organization," _American Economic Review Papers and Proceedings_ 75 (1985) 292-297.

_____ and _____, The Architecture of Economic Systems: Hierarchies and Polyarchies," _American Economic Review_ 76 (1986) 716-727.

Shiller, Robert and Pound, John, "Survey Evidence of Diffusion of Interest Among Institutional Investors," NBER Working Paper #1851, 1986.

Shleifer, Andrei, "Implementation Cycles," Journal of Political Economy 94 (1986) 1163-1190.

_____ and Vishny, Robert, "The Efficiency of Investment in the Presence of Aggregate Demand Spillovers," forthcoming, Journal of Political Economy, 1988.

Weitzman, Martin, "Increasing Returns and the Foundations of Unemployment Theory," Economic Journal 92 (1982) 787-804.

[1]In Holmstrom's (1982) example where talent is related to the ability to make good decisions, there can be inefficiencies with only one manager. This follows from his assumption that the outcome of a potential investment project is unobservable when the investment is not undertaken. Our model differs in that the state of the world that determines investment profitability is always observable. Hence there would be no inefficiencies in a single-manager setting.

[2]An analogy to noisy rational expectations models of the stock market (such as Hellwig (1980)), where all the players move simultaneously, may be helpful. In these models, traders also put some weight on information drawn from the investment decisions of others (which are reflected in the stock price) and some weight on their private information. The question we address here is whether too little weight is put on private information in making decisions.

[3]The basic idea about manipulation of the learning process was first developed by Holmstrom (1982). Fudenberg and Tirole (1986) apply the concept and refer to it as "signal jamming."

[4]This will be the case if the manager's investment decision in the second go-round of the investment does not depend on θ; the manager invests if and only if he observes the good signal. This condition is analogous to that posited for the first go-round in inequality (7).

[5]Our simplified formulation would be most literally applicable to situations where the state is publicly observable, but cannot be verified by the courts, so that profit-contingent contracts are not feasible. (For more discussion of this point, see Hart and Holmstrom (1987).) This may be a reasonable assumption for certain types of jobs where there are no easily describable performance measures, but it will not be valid in other cases, e.g., portfolio management.

[6]It is important to keep in mind that the only private information of the manager is about the signal he observes, not about his ability. Hence the "separation" is with respect to this signal.

[7]The structure of our model is similar in many respects to Crawford and Sobel's (1982) "cheap talk" model of strategic information transmission. In their framework, an informed party (the "sender" sends a costless message to another party (the "receiver") who then takes an action that affects the utility of both parties. Similarly, in our model, the informed manager (sender) makes an investment decision (sends a costless message) and the labor market (receiver) pays a wage that affects both parties' utilities. What matters in these models is not the actual wording of the messages, but rather how they are interpreted by the receiver. Hence, it is impossible to determine exactly what words will be sent, though it is possible to determine how much information is contained in these words.

[8]These facts are cited in Milgrom and Roberts (1988) to support a different theory, one based on the costs of influence activities.