



Herd Behavior and Investment

David S. Scharfstein, Jeremy C. Stein

The American Economic Review, Volume 80, Issue 3 (Jun., 1990), 465-479.

Stable URL:

<http://links.jstor.org/sici?sici=0002-8282%28199006%2980%3A3%3C465%3AHBAI%3E2.0.CO%3B2-7>

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 American Economic Review is published by American Economic Association. 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/aea.html>.

The American Economic Review
©1990 American Economic Association

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

Herd Behavior and Investment

By DAVID S. SCHARFSTEIN AND JEREMY C. STEIN*

This paper examines some of the forces that can lead to herd behavior in investment. Under certain circumstances, managers simply mimic the investment decisions of other managers, ignoring substantive private information. Although this behavior is inefficient from a social standpoint, it can be 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 decision making within firms. (JEL 026, 522)

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*, John Maynard Keynes (1936, pp. 157–58) expresses skepticism about the ability and inclination of “long-term investors” to buck market trends and ensure efficient investment. 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 judgments. 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*, Samuel Gwynne (1986, p. 58) documents problems of herd behavior in banks' lending policies toward LDCs. Discussing the incentives facing a credit analyst, he

*Sloan School of Management, Massachusetts Institute of Technology, Cambridge, MA 02139 and Harvard Business School, Boston, MA 02163. We are grateful for helpful comments from Robert Gertner, Bob Gibbons, Bengt Holmstrom, Chi-Fu Huang, Robert Merton, Jim Poterba, Julio Rotemberg, Garth Saloner, Andrei Shleifer, Larry Summers, Robert Vishny, Mark Wolfson, and two anonymous referees, and participants in seminars at the NBER Summer Institute, Harvard, Rochester, Columbia, and Northwestern.

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.

The aim of this paper is to develop a clearer understanding of some of the forces that can lead to herd behavior. We find that, under certain circumstances, managers simply mimic the investment decisions of other managers, ignoring substantive private information. Although this behavior is inefficient from a social standpoint, it can be 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 Bengt Holmstrom (1982a). 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 Joan Ricart i Costa (1986) as well as in another part of Holmstrom (1982a). 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

¹In Holmstrom's (1982a) example where talent is related to the ability to make good decisions, there are 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

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 profitable investment; 2) whether the manager's behavior was similar to or different from that of other managers.

If there are systematically unpredictable components of investment value, the first piece of evidence will not be used exclusively, since on any given draw, all smart managers could get unlucky and receive misleading signals. Hence the second piece of evidence is important as well. Holding the absolute profitability of the investment choice fixed, managers will be more favorably evaluated if they follow the decisions of others than if they behave in a contrarian fashion. Thus an unprofitable decision is not as bad for reputation when others make the same mistake—they can share the blame if there are systematically unpredictable shocks.

This "sharing-the-blame" effect arises because smart managers tend to receive correlated signals (since they are all observing a piece of the same "truth"), while dumb ones do not (they simply observe uncorrelated noise). Consequently, if one manager mimics the behavior of others, this suggests 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 is perceived as more likely to be dumb, all else being equal. 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.

world that determines investment profitability is always observable. In the case that we study, there are thus no inefficiencies in a single-manager setting.

These points are further developed in the five sections following this one. In Section I, we present the structure and assumptions of the model. The basic results on the existence of herding equilibria are summarized in three propositions in Section II. Section III looks more closely at some countervailing forces that may offset herding tendencies. In Section IV, we elaborate on the implications of the model for corporate investment, the stock market, and decision making within firms. Finally, Section V contains concluding remarks.

I. 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 moment, however, it may help for the reader interested in concreteness to bear in mind the following story: In our model, managers are in charge of capital investment at industrial firms, and they each are considering investing in a cost-saving technology. The value of the technology will be realized in the future, and each manager receives a signal that gives him some information about the 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, *A* decides whether or not to make the invest-

ment. There are two possible outcomes at date 3: either the “high” state, in which case the investment yields a profit (net of investment expense and discounting) of $x_H > 0$; or the “low” state, in which case the net profit is $x_L < 0$. The prior probabilities of these two states are α and $(1 - \alpha)$, respectively. The outcome is publicly observable, even if neither manager decides to invest.

In making his decision, *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 the low state. Formally, we have:

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

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

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:

$$(3) \quad \text{Prob}(s_G|x_H, \text{dumb})$$

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

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 the manager’s type. This amounts to assuming that: $\text{Prob}(s_G|\text{smart}) = \text{Prob}(s_G|\text{dumb})$, 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:

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

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

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 been received:

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

After *A* has made his investment decision at date 1, *B* makes his decision at date 2. Manager *B* also has access to a private signal. In addition, he can observe whether or not firm *A* has decided to invest. This is valuable information; even in a world 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 *too much* 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

²An analogy to noisy rational expectations models of the stock market (such as Martin 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. Note, however, that in a noiseless stock market, the price reveals all information, so that ignoring private signals is actually optimal. See, for example, Paul Milgrom and Nancy Stokey (1982).

distributions given in equations (1)–(3). Similarly, if both are dumb, their signals are drawn independently—so that, for example, the probability is z^2 that two dumb managers both observe the good signal.³

However, if both managers are smart, they are assumed to observe *exactly the same signal*. Thus the probability that two smart managers both observe s_G when the true state is x_H is p (as opposed to p^2 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 concerning herd behavior fail to go through.

Heuristically, herd behavior requires smart managers' prediction errors to be at least partially correlated with each other. Although this feature may seem slightly unnatural given the current setup and notation, it amounts to nothing more than saying that there are *systematically unpredictable factors affecting the future state that nobody can know anything about*.⁴ 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

³Our assumption that all dumb managers receive independent signals may seem extreme. Many investors who respond to the *same* stimuli—for example, the predictions of technical stock market forecasters—are often thought to be “dumb” for doing so. However, care must be taken in interpreting this fact with respect to our model. We only assume that dumb agents' *private* signals are independent—that is, these agents make independent errors when trying to interpret data *on their own*. The fact that people rely on common *outside* sources for forecasts in a suboptimal fashion is not at all inconsistent with our approach. Indeed, it is essentially our main prediction: privately formed opinions will be disregarded in favor of publicly available forecasts.

⁴Empirical evidence suggest that the assumption of systematically unpredictable components is a reasonable one. For example, Patricia O'Brien (1988) finds that individual security analysts' prediction errors contain common components.

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 effort-based models of relative performance such as in Holmstrom (1982b) and Barry Nalebuff and Joseph Stiglitz (1983). If prediction errors are independent, the labor market can, in our model, efficiently update on ability using only individual performance—that is, whether the manager picked a successful investment. Analogously, in models of relative performance, if the agents face independent shocks to output, optimal contracts evaluate the agent based only on individual absolute performance. However, in these models, if errors are correlated, there is an informational gain from comparing agents. 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. This is the "sharing the blame" effect touched on earlier.

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 of 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 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.⁵

⁵The basic idea about manipulation of the learning process was first developed by Holmstrom (1982a). Drew Fudenberg and Jean Tirole (1986) apply the concept and refer to it as "signal jamming."

B. Managerial Objectives

Our next step is to specify the managers' objective functions. 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 is replayed once more after date 3; 2) at this point there is 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 $\hat{\theta}$.⁶

We do not explicitly analyze contracting behavior in what follows. Rather, we assume (as do Holmstrom, 1982a, and others) that managers cannot be bound to their firms against their will *ex post*. This means that any long-term contract that would pay some types less than spot market wages in the second round of the investment game is infeasible.

Since their future wages are linear in $\hat{\theta}$, managers have some incentive to generate

⁶This 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).

high values of $\hat{\theta}$, rather than to invest efficiently in the first 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 round of the investment game. Thus, in principle, it seems reasonable to believe that managers would be induced to act so as to maximize a weighted average of expected profits and their future compensation. 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 briefly touch on how they would be altered if managers cared about expected profits.⁷

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 $\hat{\theta}$.

It should be noted that managers care only about their *absolute* ability assessment—not about whether they are judged to be more or less able than other managers. We touch on the importance of relative ability in Section III.

II. 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 now, 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 makes his decision based on the two-signal information set (s_G, s_B) . Given our symmetry assumptions, this implies the probability of the high state, $\text{prob}(x_H | s_G, s_B) = \frac{1}{2}$. Thus the investment decision hinges on the sign of $(x_H + x_L)$ —if this quantity is positive, manager *B* will invest, and if not, he will not.

Similarly, if firm *A* does not 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 are different. When manager *A* observes s_G and invests, manager *B* will also invest, regardless of his signal and 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* does not 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 aggregate information 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 of manager *B*—in all the “continuation” equilibria that we look at the manager *A* behaves efficiently, by investing

⁷Our 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 Oliver 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 is not valid in other cases, for example, portfolio management.

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 of 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.⁸ 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 does not 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:

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

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 $\theta(s_G, s_B, x_L) = 2\theta p/(1+\theta)$ also.

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

$$(9) \quad \hat{\theta}(s_B, s_G, x_L) \\ = \hat{\theta}(s_G, s_B, x_H) \\ = 2\theta(1-p)/(1+\theta); \\ (10) \quad \hat{\theta}(s_B, s_B, x_H) \\ = \hat{\theta}(s_G, s_G, x_L) \\ = \frac{2\theta(1-p)(1+\theta)}{4\theta(1-p) + (1-\theta)^2}; \\ (11) \quad \hat{\theta}(s_B, s_B, x_L) \\ = \hat{\theta}(s_G, s_G, x_H) \\ = \frac{2\theta p(1+\theta)}{4\theta p + (1-\theta)^2}.$$

⁸It 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.

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:

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

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

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

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

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* has 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 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 break the efficient equilibrium. 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* is the economically correct decision, manager *B* prefers 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 if signals were independent. Denoting these rules by $\hat{\theta}^i(\cdot)$, and using the same Bayesian logic as before, we can derive

$$(15) \quad \hat{\theta}^i(s_B, s_G, x_H) = \hat{\theta}^i(s_G, s_G, x_H) = \hat{\theta}^i(s_B, s_B, x_L) = \hat{\theta}^i(s_G, s_B, x_L) = \frac{2\theta p}{2\theta p + (1-\theta)};$$

$$\begin{aligned}
 (16) \quad & \hat{\theta}^i(s_B, s_G, x_L) \\
 &= \hat{\theta}^i(s_G, s_G, x_L) \\
 &= \hat{\theta}^i(s_B, s_B, x_H) \\
 &= \hat{\theta}^i(s_G, s_B, x_H) \\
 &= \frac{2\theta(1-p)}{2\theta(1-p) + (1-\theta)}.
 \end{aligned}$$

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 the profitability of the investment—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 generally 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 believes that he observed s_G ; and conversely, ii) if manager B deviates by not investing when A has, the labor market believes 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 is 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 manager B who observes s_B not to deviate, it must be the case that:

$$\begin{aligned}
 (17) \quad & \theta \geq \hat{\theta}(s_B, s_G, x_H) \text{Prob}(x_H | s_B, s_B) \\
 & \quad + \hat{\theta}(s_B, s_G, x_L) \text{Prob}(x_L | s_B, s_B).
 \end{aligned}$$

Inequality (17) is the requirement that the payoff to a manager who observes s_B , pools, and receives θ , exceeds his payoff from deviating and investing, given that out of equilibrium beliefs are such that he will be viewed as having observed s_G if he deviates. Direct substitution from (8), (9), and (13) verifies that the inequality is satisfied.

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

$$\begin{aligned}
 (18) \quad & \theta \geq \hat{\theta}(s_B, s_G, x_H) \text{Prob}(x_H | s_B, s_G) \\
 & \quad + \hat{\theta}(s_B, s_G, x_L) \text{Prob}(x_L | s_B, s_G).
 \end{aligned}$$

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:

$$(18') \quad \theta \geq \theta / (1 + \theta),$$

which is strictly 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 .

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.⁹

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 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 always to 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 $\hat{\theta}^A$ is a function of only 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

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

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

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

⁹The structure of our model is similar in some respects to Vincent Crawford and Joel Sobel’s (1982) “cheap talk” model of strategic information transmission. In their framework, an informed party (the “sender”) transmits a message to another party (the “receiver”), who then takes an action that affects the utility of both parties. Similarly, in our model, one can view the investment decision as a message sent by the informed manager and the wage paid by the outside labor market as the response of the receiver. (However, our model differs from theirs in that the message itself—the investment decision—has real economic consequences to a third party, namely, the manager’s firm.) 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.

must hold:

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

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

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.¹⁰

Finally, it should be emphasized that herding equilibria exist in a model with any number of managers. Consider 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 signal, manager *C* learns nothing from observing what firm *B* has done. He learns only from observing firm *A*, whose 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. Note, however, that other equilibria may also exist with more than two managers. For example, with three managers, the second may deviate from the first if he conjectures that the third manager will join him in a new herd.

¹⁰It should be pointed out that with different parameter values, reputational concerns could induce even manager *A* to behave differently. Suppose that x_H is relatively large. Then efficiency could require manager *A* to invest after seeing s_G even if the posterior probability of success μ_G is small. Inspection of (21) reveals that this efficient equilibrium cannot be sustained if $\mu_G < \frac{1}{2}$; manager *A* will be unwilling to invest and thereby contradict the *ex ante* wisdom that success is very unlikely.

III. Countervailing Forces

Up to this point, the model has been simplified to focus exclusively on the forces that make herd behavior likely. Naturally, there are other considerations which may offset herding tendencies. We discuss four such considerations below: managerial concern for profits; limited liability; wages that depend on relative, rather than absolute talent; and alternative definitions of ability.

The herding equilibrium derived above is generally inefficient relative to the first best for some configurations of private information. For example, if $x_H + x_L > 0$, then the equilibrium is inefficient when manager *A* observes s_B , manager *B* observes s_G , and firm *B* fails to invest. The expected profits lost by firm *B* due to herding, denoted by Π , are equal to $\frac{1}{2}(x_H + x_L)$.

If, however, managers place some weight on expected profits in their objective functions, these inefficiencies can disappear. Denote by R the amount by which the incentive constraint in (14) is violated. R measures the perceived reputational benefit to managers from herding, and it is proportional to

$$\left[\hat{\theta}(s_B, s_B, x_H) - \hat{\theta}(s_B, s_G, x_H) \right] \\ + \left[\hat{\theta}(s_B, s_B, x_L) - \hat{\theta}(s_B, s_G, x_L) \right].$$

Managers who care only about their reputations will always herd, since $R > 0$. But managers who care about profits will have to trade off R against Π . For any given weight on profits in managers' objective functions, a large enough value of Π will restore the efficient signal-dependent equilibrium. Put differently, as the weight on profits increases, the range of parameter values over which there is herd behavior shrinks—the more egregious inefficiencies associated with herd behavior are alleviated, even though herding can still occur when Π is relatively small. This points to a role for short-term incentive contracts. Even if managers cannot be bound to their firms for life (so that they always place some weight on reputation) short-term contracts may shift the focus toward ex-

pected profits, thereby reducing herding tendencies.

Limited liability is another factor that can alter the tradeoff between reputation and profits. It may be that managers have fixed outside labor market opportunities and thus never have to accept a wage less than L . If L exceeds the wage to a manager with ability $\theta(s_B, s_G, x_L)$ there is a floor on how poorly a manager can fare by disagreeing with his peers. Thus limited liability lowers R , the reputational gain from herding. Like profit-based compensation, limited liability shrinks the range of parameter values over which herd behavior is observed. A corollary is that herding may become more or less of a problem as a manager's career progresses. On the one hand, there is apt to be less uncertainty about the manager's ability, which should *reduce* the incentives for herd behavior. On the other hand, later in a successful career, wages are probably higher above the outside alternative L . This latter effect can *increase* the propensity to herd.

Relative ability concerns are a third factor that may offset herding tendencies. In our model, managers only care about having a high absolute ability assessment, $\hat{\theta}$. In certain situations, however, they may also care about how their $\hat{\theta}$ compares with that of other managers. For example, there may be a "superstars" effect present (see Sherwin Rosen, 1981), with top-ranking managers getting a disproportionately high wage. If this is the case, manager B will be more reluctant to mimic manager A , since by doing so he destroys any possibility of being the top-ranked manager.

A final way to attenuate the model's predictions regarding herd behavior would be to introduce a broader definition of ability. As we have defined it, ability is nothing more than an aptitude for making precise forecasts about the outcome of a *given* random variable. But ability might in addition include a knack for *uncovering* new random variables to study. Rather than having to decide simply whether or not to buy a certain asset, managers may also be responsible for finding alternatives to that asset. Under these circumstances, a desire to earn a repu-

tation for "creativity" may deter managers from all choosing to buy the same asset.

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 LDCs was mentioned earlier as an apparent instance of herd behavior in corporate investment. A recent paper by Randall Morck, Andrei Shleifer, and Robert Vishny (1989) suggests that the potential for problems 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 industrywide 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—at least in this segment of the labor market, there seems to be support for the "sharing-the-blame" effect discussed earlier. And as long as this effect is at work, one might expect that managerial behavior would be distorted in the direction of herding.

Herd behavior in corporate investment may have important consequences for the adoption of new technologies. Papers by Joseph Farrell and Garth Saloner (1985), and Michael Katz and Carl Shapiro (1985), have shown that compatibility externalities can lead to bandwagon effects in technology adoption. Our work can be viewed as complementary to theirs: even when compatibility is not an important concern, bandwagons can arise. If the manager of one firm adopts a particular technology, this creates a reputational externality, in the sense that other managers will tend to be biased toward the same technology for reputational reasons.

B. *The Stock Market*

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.

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, 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 percent of his assets in stocks and 50 percent 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.¹¹ 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

¹¹Another subtlety that is present in the stock market case but not in our model is a continuum of investment choices. A portfolio manager whose private information tells him that stocks are overpriced can “partially herd” by putting 45 percent, rather 50 percent of his assets in stocks. This may tend to dampen the aggregate effects of herd behavior.

who are unconcerned about their reputations.

Robert Shiller and John 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 run-ups 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 suggests that the comfort inherent in following common wisdom can lead professional money managers to invest in stocks where fundamentals might dictate otherwise.

C. *Decision Making Within Firms*

Recent work on the theory of the firm by Raaj Kumar Sah and Joseph 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 tend to have errors of judgment that are not perfectly correlated. 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 decision making, and also offers some new insights about how organizational structure can facilitate the decision-making process. As a stylized example, consider the case of a capital budgeting committee meeting, where the 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. However, 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.¹²

One way around this problem is to have managers submit their votes simultaneously, perhaps in writing. However, this may limit valuable exchanges of ideas. An alternative approach is to have those with the stronger reputational concerns vote first. As we noted in the previous section, if the limited liability effect is not too important, reputational concerns will be strongest among young managers, since there is presumably more uncertainty about their ability. Thus if the committee consists of young and old executives, the young ones should be asked to voice their opinions before the old ones. 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 can be passed upstream for approval, this may result in better decision making than if the ideas are originated at a high strategic planning level and then are passed downstream for line manager input.¹³

V. Conclusions

Herd behavior can arise in a variety of contexts, as a consequence of rational attempts by managers to enhance their reputations as decision makers. In addition to reputational concerns, there are other factors that influence herding. One of these is the extent to which there are commonly unpredictable components to investment outcomes: correlated prediction errors lead to the "sharing-the-blame" effect that drives

¹²There are other explanations for this "yes-man" effect. One is discussed in Solomon Asch's (1955) classic study of the effects of group pressure. He found that experimental subjects were reluctant to disagree with others in the group (confederates) despite the fact that the confederates' stated opinions were clearly wrong. Asch interprets the results as evidence in support of the view that individuals have an inherent psychological desire to conform to group norms.

¹³Keitaro Hasegawa (1986) discusses at length the emphasis placed by Japanese corporation on implementing a bottom-up flow of information. This form of organizational design is known there as "ringi."

LIST OF NOTATION

Symbol	Meaning
x_H	Investment proceeds in "high" state
x_L	Investment proceeds in "low" state
α	<i>Ex ante</i> probability of high state
s_G	Good signal
s_B	Bad signal
p	Prob($s_G x_H$, smart)
q	Prob($s_G x_L$, smart)
z	Prob($s_G x_H$, dumb) = Prob($s_G x_L$, dumb)
θ	<i>Ex ante</i> probability that manager is smart
$\hat{\theta}$	Posterior probability that manager is smart
μ_G	Prob($x_H s_G$)
μ_B	Prob($x_H s_B$)
Π	Profit passed up in herding equilibrium
R	Reputational benefit to herding
L	Wage that can be earned in outside opportunity

managers to herd. Also important is the nature of the managerial labor market: herding is more likely to be a problem when managers' outside opportunities are relatively unattractive, and when compensation depends on absolute rather than relative ability assessment.

REFERENCES

- Asch, Solomon, "Opinions and Social Pressure," *Scientific American*, 1955, 193.
- Crawford, Vincent and Sobel, Joel, "Strategic Information Transmission," *Econometrica*, November 1982, 50, 1431-51.
- Farrell, Joseph and Saloner, Garth, "Standardization, Compatibility, and Innovation," *Rand Journal of Economics*, Spring 1985, 16, 70-83.
- Fudenberg, Drew and Tirole, Jean, "A Signaling Jamming Theory of Predation," *Rand Journal of Economics*, Autumn 1986, 17, 366-76.
- Gwynne, Samuel, *Selling Money*, New York: Weidenfeld and Nicolson, 1986.
- Hart, Oliver and Holmstrom, Bengt, "The Theory of Contracts," in Truman Bewley, ed., *Advances in Economic Theory, Fifth World Congress*, New York: Cambridge University Press, 1987.
- Hasegawa, Keitaro, *Japanese-Style Manage-*

- ment: An Insider's Analysis*, Tokyo: Kodansha, 1986.
- Hellwig, Martin**, "On the Aggregation of Information in Competitive Markets," *Journal of Economic Theory*, June 1980, 22, 477-98.
- Holmstrom, Bengt**, (1982a) "Managerial Incentive Problems: A Dynamic Perspective," in *Essays in Economics and Management in Honor of Lars Wahlbeck*, Helsinki: Swedish School of Economics, 1982.
- _____, (1982b) "Moral Hazard in Teams," *Bell Journal of Economics*, Autumn 1982, 13, 324-40.
- _____ and **Ricart i Costa, Joan**, "Managerial Incentives and Capital Management," *Quarterly Journal of Economics*, November 1986, 101, 835-60.
- Jensen, Michael**, "Agency Costs of Free Cash Flow, Corporate Finance, and Takeovers," *American Economic Review Papers and Proceedings*, May 1986, 76, 323-29.
- Katz, Michael and Shapiro, Carl**, "Network Externalities, Competition, and Compatibility," *American Economic Review*, June 1985, 75, 424-40.
- Keynes, John Maynard**, *The General Theory of Employment, Interest and Money*, London: Macmillan, 1936.
- Milgrom, Paul and Stokey, Nancy**, "Information, Trade, and Common Knowledge," *Journal of Economic Theory*, February 1982, 22, 17-27.
- Morck, Randall, Shleifer, Andrei and Vishny, Robert**, "Alternative Mechanisms for Corporate Control," *American Economic Review*, December 1989, 79, 842-52.
- Nalebuff, Barry and Stiglitz, Joseph**, "Information, Competition, and Markets," *American Economic Review Papers and Proceedings*, May 1983, 73, 278-83.
- O'Brien, Patricia**, "Analysts' Forecasts as Earnings Expectations," *Journal of Accounting and Economics*, January 1988, 10, 53-83.
- Rosen, Sherwin**, "The Economics of Superstars," *American Economic Review*, December 1981, 71, 845-58.
- Sah, Raaj Kumar and Stiglitz, Joseph**, "Human Fallibility and Economic Organization," *American Economic Review Papers and Proceedings*, May 1985, 75, 292-97.
- _____ and _____, "The Architecture of Economic Systems: Hierarchies and Polyarchies," *American Economic Review*, September 1986, 76, 716-27.
- Shiller, Robert and Pound, John**, "Survey Evidence of Diffusing of Interest Among Institutional Investors," NBER Working Paper No. 1851, 1986.