

November 13, 2000

Conversation, Observational Learning, and Informational Cascades

H. Henry Cao*

David Hirshleifer**

*Haas School of Business, UC Berkeley, 545 Student Services Bldg # 1900, Berkeley, CA 94720-1900; cao@haas.berkeley.edu

**Fisher College of Business, The Ohio State University, 740A Fisher Hall, 2100 Neil Avenue, Columbus, OH 43210-1144; hirshleifer.2@osu.edu

We thank Steven Durlauf and seminar participants at UC Berkeley, University of Chicago, and University of Wisconsin for helpful comments.

Abstract

**Conversation, Learning
and Informational Cascades**

We offer a model to explain why groups of people sometimes converge upon poor decisions and are prone to fads, even though they can discuss the outcomes of their choices. Models of informational herding or cascades have examined how rational individuals learn by observing predecessors' actions, and show that when individuals stop using their own private signals, improvements in decision quality cease. A literature on word-of-mouth learning shows how observation of *outcomes* as well as actions can cause convergence upon correct decisions. However, the assumptions of these models differ considerably from those of the cascades/herding literature. In a setting which adds 'conversational' learning about both the payoff outcomes of predecessors to a basic cascades model, we describe conditions under which (1) cascades/herding occurs with probability one; (2) once started there is a positive probability (generally less than one) that a cascade lasts forever; (3) cascades aggregate information inefficiently and are fragile; (4) the ability to observe past payoffs can reduce average decision accuracy and welfare; and (5) delay in observation of payoffs can improve average accuracy and welfare.

1 Introduction

Historians and social observers have often commented on the ‘folly,’ ‘fickleness,’ ‘mindlessness,’ or ‘madness’ of mass behavior. Such judgements are sometimes arrived at too cheaply with the benefit of hindsight. Nevertheless, the prevalence of fads, and of shifts in behavior without obvious justification, suggests that the problem of error-prone group behavior is real. From a rational perspective, this problem is surprising. In principle, when large numbers of rational individuals communicate their imperfectly correlated signals, decisions should improve greatly.

A possible explanation for the idiosyncrasy of mass behavior is that individuals aggregate information poorly. The literature on informational herding or cascades (see, e.g., Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992)) has studied an informational externality problem that arises when individuals learn by observing others.¹ In simplest form, the cascades/informational herding model assumes that individuals’ payoffs are not interdependent so that there is no strategic interaction between players. The basic idea of cascades/herding theory is that an individual who acts rationally based on information gleaned from observation of predecessors makes his own action less informative to later observers than it might otherwise be. An individual may find it optimal to choose an action consistent with the choices or experience of others regardless of his own (possibly opposing) private signal. In such a situation he is said to be in an informational herd or cascade, and his action is uninformative to later individuals. If individuals are *ex ante* identical, and the only sources of information available to an individual is his private signal and past actions, then all subsequent individuals are also in an informational cascade, and information aggregation ceases. In a path-dependent fashion, the system reaches a precarious equilibrium in which the great mass of individuals make relatively ill-informed decisions, and are therefore prone to shift their behavior in the event of even a modest external shock.²

¹The importance of imitation and observational influence has been documented by numerous studies in fields such as psychology, sociology, zoology, political science, and economics.

²For a review of informational cascades/herding theory, social learning, and applications see Bikhchandani, Hirshleifer and Welch (1998). A recent literature has studied consequences of informational cas-

Existing models of informational cascades/herding essentially almost rule out conversation. These models assume that individuals are unable to observe or discuss outcomes (the payoff consequences of past decisions). A justification for a modelling focus on observation of past actions is the problem of credibility: ‘actions speak louder than words.’ Furthermore, conversation (and email) requires time and attention. However, in many situations it is clearly going too far to discount the effect of words entirely. Ellison and Fudenberg (1993, 1995), and Banerjee and Fudenberg (1999) discuss the importance of word-of-mouth communication in everyday choices among restaurants or auto mechanics, and empirical literature describing how the diffusion of innovation is associated with conversation between potential adopters. Examples include the diffusion of agricultural innovations and the adoption of computer technology. Several studies have found evidence that verbal communication between individuals exerts a strong influence on purchasers’ judgements about a variety of different kinds of consumer products and business factor inputs; see, for example, the citations in Herr, Kardes and Kim (1991). A survey by Shiller and Pound (1989) asked individual investors what first drew their attention to the company whose stock they had purchased most recently. Almost all investors named sources which involved direct personal interaction. Shiller (2000; chapter 8) cites several other studies indicating that conversation plays a crucial role in security investor decisions. Shiller also discusses the importance of the reactions of early viewers of new movies and the resulting word-of-mouth ‘buzz’ that cascades/herding for a variety of social phenomena, such as bank runs, stock market crashes, revolutions, investment decisions, corporate strategic choices, conformity in organizations, aggregation of information in debate, macroeconomic fluctuations, mate choices, the effects of entrepreneurs and social misfits on society, option exercise, public attitudes toward environmental hazards, purchase strategies of buyers, and the pricing of stock offerings; see Banerjee (1992), Bikhchandani, Hirshleifer and Welch (1992), Welch (1992), Chen (1993), Corb (1993), Lee (1993, 1999), Chamley and Gale (1993), Lohmann (1994, 1997), Caplin and Leahy (1994), Khanna and Slezak (1995), Zhang (1997), Khanna (1998), Nelson (1998), Grenadier (1999), Kennedy (1999), Perktold (1999), Avery and Zemsky (1999), Kuran and Sunstein (1999), Taylor (1999), and Smith and Sorenson (2000). Anderson (1994) and Anderson and Holt (1997) study the formation of cascades/herding experimentally. Gibson and Hoglund (1993) review animal imitation and discusses possible zoological applications of cascades/herding.

they create.³

On the other hand, if conversation perfectly conveyed *all* information, we would all know as much as the wisest person in the world. In a thought-provoking essay, Shiller (1995) suggests that there are severe failures in learning through conversation, that these failures may be responsible for poor decisions by society as a whole, and that these problems explain puzzling phenomena in financial markets. We take as given here certain barriers to communication, and explore whether these barriers clog the information pipeline so tightly that large numbers of individuals make ill-informed decisions.

A recent literature on *word-of-mouth learning* has provided important insights about situations in which individuals learn by observing information about past payoffs as well as actions—but not the private signals—of predecessors. The assumption that signals are not communicated provides an attractive starting point for modelling conversation. It is often easier to convey information about the objective payoff consequences of past actions than about the internal reasons for past actions (i.e., private information signals). Subjective information is harder to convey credibly because it is not directly verifiable, is less salient than payoffs, and is often hard to describe succinctly.

One of the interesting contributions of the word-of-mouth learning papers is to describe settings in which only a modest amount of observability, communication and rationality are needed to generate efficient outcomes. The paper in this literature most relevant for this one is that of Banerjee and Fudenberg (1999), who examine the decisions of fully rational decisionmakers.⁴ Banerjee and Fudenberg (1999) analyze a model consisting of successive

³According to Shiller, “It is well known that the advice of movie critics has less impact than the mass effect of such word of mouth.”

⁴Ellison and Fudenberg (1993) and (1995) examine a setting in which a continuum of individuals make simultaneous decisions in each generation. Players observe a subset of past actions and signals about past payoffs, and make choices using simple rules of thumb that are potentially consistent with rationality. Players can converge to the efficient choice even for rules of thumb that are quite naive. Ellison and Fudenberg (1993) consider a static decision environment, whereas Ellison and Fudenberg (1995) examine the effects of stochastic environmental shifts. There is also a more distantly related literature in which individuals adopt whatever choice they hear about without regard to other information signals.

generations of continua of rational optimizing agents. Individuals choose simultaneously based on sampling and observation of the choices and payoffs of a subset of members of the previous generation. The aggregate state of the population evolves deterministically. If each player samples two or more others, the payoff observation is sufficiently informative, the starting point of the system is known to all, and there are no reporting biases, Banerjee and Fudenberg show that the long run outcome is efficient.

In contrast, under the assumptions of the cascades/herding literature, decisions are inefficient. Since the two strands of literature approaches make several different assumptions, it is not obvious whether introducing communication/observation of payoffs to a setting that allows for path-dependent outcomes and possible cascades would lead to efficient outcomes. The purpose of this paper is to explore this issue. We examine whether the ability to learn about the payoffs of predecessors results in a substantial probability of inefficient decisions by a substantial fraction of the population, and fragile social outcomes.

We do so by combining certain features of the cascades/herding and the word-of-mouth learning literatures that we believe are often reasonable approximations of reality. As in cascades/herding papers and Banerjee and Fudenberg (1999), we examine the decision problems of rational optimizing agents. As in the cascades/herding literature, individuals also receive private signals about value that are unrelated to the actions of predecessors. This can be viewed as information derived from direct study of the project rather than information derived from past payoffs. Thus, in our model a choice between actions conveys information of two different types to later decisionmakers— the information derived from the payoff outcomes, and the private information implicit in the action choice. The model is populated with discrete individuals acting sequentially, which causes the system to evolve stochastically in a path-dependent fashion. Furthermore, as in cascades/herding settings, in our model individuals observe the full history of decisions and payoffs, rather than a small subsample. However, as in the word-of-mouth learning literature, individuals observe signals about the payoff outcomes of past decisions.

Our setting therefore allows much greater observability or communication than either the cascades/herding or the word-of-mouth learning literatures (payoffs not just actions;

all of history rather than a subsample). One might suspect that the opportunity to observe/communicate numerous past payoffs as well as actions would sooner or later dislodge all information blockages, leading to correct choices. However, we describe circumstances in which cascades/herding are in the long run ill-informed (i.e., are mistaken with positive probability). Depending on the early sequence of actions and payoffs, a point can be reached where with positive probability (generally less than one) the remaining individuals choose one action only a finite number of times. In other words, at some point the project is never chosen again. If this happens, uncertainty about the quality of the project chosen infinitely often disappears, but uncertainty about the alternative project persists forever.⁵

Just as in the basic cascades/herding literature without payoff observation, cascades/herding can be fragile in the sense that introducing a small external shock can change the behavior of large numbers of individuals. Furthermore, we provide examples in which the ability to observe past payoffs in addition to actions leads to less accurate decisions and average welfare than if only past actions are observable. Intuitively, the observation of past payoffs can trigger inefficient cascades/herding earlier. Payoff information is helpful to an individual who observes it, but this can render his decision less informative to subsequent individuals. Finally, we show that delay in observation of payoffs can improve average decision quality and welfare. Thus, our model is consistent with both learning through conversation about (or observation of) outcomes, and also the stylized facts about conformity and social change discussed at the start of the text.

A distinct literature involving observation of past outcomes studies variations of the multi-arm bandit problem of Rothschild (1974). Multi-arm bandit models generally focus on problems of information acquisition rather than aggregation of information. In Rothschild (1974), there is only one decision maker who plans an optimal policy of experimentation. Bolton and Harris (1999) examine a multi-arm bandit problem with multiple individuals who learn by observing the outcomes of others. All information that is generated is aggregated efficiently. There are no private signals. Inefficient outcomes occur because of an externality

⁵Indeed, in the special case where individuals observe past payoffs perfectly and a project's past payoff is a perfect indicator of its future payoff, a cascade once formed lasts forever with certainty.

problem. Individuals do not experiment sufficiently because they do not care about the informational benefit to later individuals.⁶

In contrast with the multi-arm bandit literature, in the word-of-mouth learning papers with a continuum of decisionmakers there is no problem of insufficient experimentation. In these models every project is taken by an infinite number of individuals in the initial round. Experimentation cannot add to the initial information pool. In a stable decision environment, if the information generated thereby were aggregated efficiently, efficient actions would result. Thus, the problem in the word-of-mouth learning literature is one of imperfect information aggregation rather than insufficient experimentation.

Like the multi-arm bandit literature, in our analysis there is a problem of information generation. This arises because there are only finitely many individuals, so that at each stage only a single action is taken. However, unlike the multi-arm bandit literature, in our model each individual makes only a single decision, so individuals do not intentionally design strategies to generate new payoff information. Furthermore, in our analysis individuals possess private signals which are aggregated in a socially suboptimal fashion.⁷ Thus, in our setting there are potential inefficiencies both in generating and in aggregating information signals.⁸ Our analysis therefore brings together aspects of both the social learning and the multi-arm bandit literatures.⁹

The remainder of the paper is structured as follows. Section 2 provides a simple example to illustrate the inefficient cascades when there is conversation. Section 3 examines the main

⁶Aoyagi (1998) examines a two-armed bandit problem played in parallel by multiple players. Players observe the actions of others but not payoffs. Aoyagi finds that players eventually settle on the same arm with probability one.

⁷Indeed, if private signals were aggregated optimally, in the long run only the efficient action would be chosen.

⁸Caplin and Leahy (1993) examine a setting where potential industry entrants learn indirectly from the actions of previous entrants by observing industry market prices. However, in this setting observation of the entry decision itself does not provide any information because firms do not possess any private information until after they have entered.

⁹See also Schlag (1998), who compares the success of imperfectly rational decision rules in a multi-arm bandit setting.

model with transient shifts in payoffs and/or noisy observation of payoffs. Section 4 provides an example to illustrate how observation of payoffs can reduce welfare and decision accuracy. Section 5 concludes the paper. All proofs are in Appendix.

2 Conversation and Cascades: A Simple Example

This section provides a simple example where the start of a cascade/herding ends all experimentation. If such a cascade starts, one of the choice alternatives will never be tried. Thus, inefficient cascades can occur with positive probability, and once started last forever (in the absence of exogenous shocks). In the general model of Section 3, inefficient cascades can occur even if there is a great deal of experimentation. Because payoffs are subject to transient shocks or are observed with noise, we will see there that inefficient cascades can start even when all choice alternatives have been tried any number of times; and that inefficient cascades have a positive probability of being temporary and a positive probability of lasting forever.

Assume there is a sequence of individuals, each of whom decides what project to adopt from a set of projects. Each individual observes the decisions of all those ahead of him. Each individual privately observes an identically distributed and conditionally independent signal about *expected* profitability of different action choices (projects). The expected gain from each project is the same for all individuals. The realized value of a given project is subject to transient idiosyncratic shocks, i.e., two different individuals may gain different amounts from the same project.

In addition, depending on the scenario, the individual may observe a signal about the payoff outcomes resulting from the action choices of previous individuals. Individuals are ordered sequentially; the order is exogenous and known to all. Since our setting allows for observation of more than in the original cascades/herding models, we need to generalize slightly the definition of a cascade.

Definition: *An informational cascade occurs if, based on his observation of others, an*

individual's action does not depend on his private signal.

This definition generalizes the notion of an informational cascade to apply in a setting where individuals can observe more than just past actions, such as the payoff outcomes resulting from past action choices. In the special case of our setting where there is no observation of a signal about past payoff outcomes, Bikhchandani, Hirshleifer and Welch (1992) showed that if the number of individuals is sufficiently large, the probability that an informational cascade/herding eventually starts approaches one. (See also the related results of Banerjee (1992) and Welch (1992).)

Definition: *An observation structure is **long-run inefficient** if in the long run the choices made are on average inferior to those made when all information is aggregated optimally.*

Implicit in the phrase ‘on average’ is both a probabilistic expectation and averaging over large numbers of individuals. In a scenario in which signals are aggregated perfectly, in the long run everyone makes the correct decision, so if the structure is long-run efficient then almost surely almost everyone chooses the better action.

If cascades/herding lead to a positive probability that wrong decisions are made in the long run, then decisions are long-run inefficient, or in the terminology of Bikhchandani, Hirshleifer and Welch (1992), idiosyncratic. They also show that cascades/herding are fragile. We define fragility formally as follows:

Definition: *An information regime is **fragile** at a given point in time if an exogenous additional public disclosure with precision less than or equal to the information signal possessed by a single individual will with positive probability shift the behavior of the next individual.*

An important feature of the model is that individuals learn only about the payoffs of alternatives that are actually adopted. As in a multi-arm bandit model, this causes the arrival of new information to depend on past action choices.¹⁰ In addition, we allow payoffs to be subject to transient shocks and/or observation of payoffs to be noisy.¹¹ We show that

¹⁰In contrast to traditional multi-arm bandit models, we examine how the actions of an individual affect others.

¹¹Actual conversation is noisy for several different reasons. Shiller (1995) emphasizes the imperfection of conversation. This may arise from discreteness in conceptual categorizing, from difficulty in putting ideas

owing to action dependence, cascades/herding can form in a regime with observation of past payoffs even if these payoffs are observed perfectly and are perfect indicators of the value of action alternatives.

The basic idea is very simple. Suppose that there are two projects A and B . The payoff of project A is either $v_A = 0$ or 2 , with probabilities μ and $1 - \mu$; the payoff of project B is known to all, 1 . Individuals receive a series of signals about the value of project A and the distribution of these value signals conditional on the value of project is i.i.d. Each signal has two realizations, H and L and $\Pr(H|v_A = 2) = p, \Pr(H|v_A = 0) = 1 - p$. Let $v_A(S)$ denote the posterior mean when a sequence of signals S is observed. We assume that μ is such that the prior mean payoff $v_A(0)$ of project A is less than 1 but the posterior mean after receiving one H signal, $v_A(H)$ is greater than 1 . Since the payoff of project B is a constant, there is nothing to learn about project B .

If the first individual obtains an L signal, which suggests that project B is best, his action reveals the signal he receives. Clearly, the next individual will choose project B if he receives an L signal. The second individual will also choose project B even if he receives an H signal. This is because the posterior mean of project A after observing an LH sequence, $v_A(LH)$, is the same as the prior mean, $v_A(0)$, which by assumption is less than the payoff of project B . So if individual 1 chooses B , a cascade on project B forms; the next individual chooses B no matter what signal he receives. Observability of past payoffs does nothing to break this cascade, because observation of non-stochastic payoffs on B tells successors nothing about the profitability of project A . Thus, in the absence of interfering shocks, inefficient cascades/herding on project B persist forever. If, on the other hand, the first individual selects project A , then all uncertainty about A is resolved and the ensuing cascade is always efficient.

If the first individual invests in project B , the resulting cascade is idiosyncratic. The true payoff of project A could be higher than B , whereas if information of many individuals were aggregated optimally, they would converge upon the correct decision with probability 1 .

into words and in interpreting words, and from conflicts of interest between the communicators.

Furthermore, the inefficient cascades outcome is fragile with respect to exogenous shocks. If the model is modified so that there is a release of a public signal of modest precision at a given date, this can break the cascade and affect the behavior of many later individuals. Consider for example a public signal with the same precision as the private signals that individuals receive. The arrival of a public signal realization of H on A can shatter a cascade on B if the next individual also observes a private H signal. If A is preferable, thereafter all individuals choose A instead of B .

In this simple example there can only be an inefficient cascade upon one of the projects, and once the inefficient cascade is broken the correct choice is always made thereafter. In the general model, there is a positive probability of an inefficient cascade on either project. Furthermore, there exist sample paths on which both projects are tried any number of times, yet inefficient cascades persist. Cascades have a positive probability of lasting forever.¹²

3 A Model with Transient Shocks to Payoffs and Noisy Observation of Payoffs

In the example in Section 2, all individuals who adopt the same project receive the same payoff. We now offer a general model that allows for transient or individual-specific payoff shocks, i.e., payoffs that are stochastic conditional on the project state. We will show that under mild distributional assumptions, even though payoffs are observable, cascades always arise, and with positive probability are incorrect. Thus information aggregation is inefficient relative to a situation in which private signals are aggregated perfectly. For notational simplicity, we first model the case in which observation of past payoffs is perfect. We then will show that this analysis generalizes immediately to the case in which signals

¹²When individuals observe a noisy summary statistic of past decisions, and the action space is continuous, a continual accumulation of additional information leads to eventual convergence toward the efficient choice (Vives (1993)). The cascades/herding approach focuses on discrete or bounded action spaces, resulting in information blockages and behavioral fragility. The relation between action space and the formation of cascades/herding is studied in Lee [1993].

about past payoffs are noisy.

3.1 Transient or Individual-Specific Payoff Shocks

3.1.1 The Economic Setting

Let there be a sequence of individuals $i = 1, 2, 3, \dots$. Each must choose between two project alternatives. The state of project n is u_n or d_n (up or down), $n = 1, 2$. Each project has two possible payoffs, $v_{n1} < v_{n2}$, and the project payoff \tilde{v}_n^i is stochastic conditional on the project's state, and its distribution depends on the state but not on the individual. An i superscript is used to denote the payoff of the project if adopted by individual i . The prior probability that $\tilde{v}_n^i = v_{nk}$ is denoted $\mu_k(s_n)$, $s_n = u_n, d_n$. The project payoffs v_n of different projects and the project states s_n are uncorrelated across projects.¹³ Throughout the paper, we assume that the monotone likelihood ratio property holds, that is, a project at the up state is more likely to have higher payoffs and gives out higher signals.

We apply the concept of perfect Bayesian equilibrium.¹⁴ Individuals observe predecessors' actions and payoffs of the projects they have chosen. Individuals differ not only by their positions in the queue but by the signals they privately observe. Each individual i observes a vector of conditionally independent and identically distributed direct signals $\tilde{\sigma}^i = (\tilde{\sigma}_1^i, \tilde{\sigma}_2^i)$, where $\tilde{\sigma}_n^i$ is a signal about the state of project n , with possible values $\sigma_{n1} < \sigma_{n2}$. Let Σ denote the set of observable signals. These signals are independent across projects, and (conditional on state) are independent across individuals. Let $p_q(s_n)$ be the probability that an individual observes signal realization σ_{nq} given that the true state of project n is s_n . We assume that $p_q(s_n) > 0$ for all n, q and s . Let $P_q(s_n)$ be the cumulative distribution

¹³It is straightforward to extend the model to the case in which the number of projects, the number of possible values from each project and the number of signal values are more than two; the proofs of propositions in the appendix allow for these possibilities. However, to avoid notational complexity in the presentation in the main text we examine the binary case.

¹⁴Since an individual's payoffs do not depend on what later individuals do, there is no incentive to make an out-of-equilibrium move to try to influence a later player. Thus, without loss of generality, we assume that if any player is observed to deviate from the equilibrium, subsequent individuals have the same beliefs as though he had chosen his correct action.

of σ_n^i , given the state of the project is s , that is,

$$P_q(s_n) \equiv \Pr(\sigma_n^i \leq \sigma_{nq} | s_n) = \sum_{k=1}^q p_k(s_n).$$

Let J^i be the set of signal realizations that lead individual i to choose action a^i . His decision a^i communicates to others that he observed a signal in the set J^i . If $J^i = \Sigma$, then the individual is in a cascades, and individual i 's action conveys no information about his signal realization.

Let a^i be individual i 's action and let $A^i = (a^1, \dots, a^i)$ represent the history of actions taken by individuals $1, 2, \dots, i$. Let b^i be the *value payoff* from individual i 's action and let $B^i = (b^1, \dots, b^i)$ represent the history of payoffs from the action of individuals $1, 2, \dots, i$. Given history A^{i-1}, B^{i-1} , let $J^i(A^{i-1}, B^{i-1}, a^i)$ be the set of signal realizations that lead individual i to choose action a^i .

Then individual $i+1$'s conditional expectation of adopting project n given his own signal realization $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)$ and the history A^i, B^i is:

$$V_n^{i+1}(A^i, B^i, \sigma) \equiv E[V | \sigma^{i+1} = \sigma, \sigma^k \in J^k(A^{k-1}, B^{k-1}, a^k), \forall k < i+1]$$

Individual $i+1$ adopts project n if the conditional expectation of project n is larger than that of other projects. When there is a tie, we take the convention that the project listed last on the queue is adopted,¹⁵ That is if $V_n^{i+1} \geq V_k^{i+1}, k \neq n$ and $V_n^{i+1} > V_k^{i+1}, k > n$. Therefore, the inference drawn from $i+1$'s action a^{i+1} is that $\sigma^{i+1} \in J^{i+1}(A^i, B^i, a^{i+1})$.

If all information is aggregated efficiently, then by the law of large numbers, the correct decision will be made by everyone except a finite number of early individuals. Thus, we will use the term 'long-run efficient' to refer to a situation where with probability one, all but a finite number of individuals make the correct choice.

We impose a No long-run tie assumption that if individuals learn enough about value by observing predecessors, then they are not indifferent between adopting and rejecting.

Assumption 1. No long-run ties. $\bar{v}_{sn} \neq \bar{v}_{s'n'}$ if $n \neq n'$, where \bar{v}_{sn} denote the expected payoff of project n given state of the project is s .

¹⁵This convention does not affect any of the results. For example we can also assume that individuals will randomize among tied projects.

3.1.2 Implications

In this setting inefficient cascades/herding occur.

Proposition 1 *If Assumption 1 holds, and individuals can observe both actions and payoffs, then as the number of individuals increases, (i) the probability that cascades/herding eventually starts approaches one; (ii) inefficient cascades/herding occur with positive probability; (iii) there is a positive probability that inefficient cascades/herding last forever; and (iv) if payoffs are subject to nondegenerate transient shocks, the probability that cascades/herding last forever can be less than one.*

The intuition is as follows. (i) If individuals try two different projects and thereby observe payoffs many times, by the law of large numbers they will eventually update so much about which alternative is better that someone will start to ‘ignore’ his private signals about project states (i.e., his decision will be independent of his private signals). This starts a cascade.

(ii) Furthermore, an inefficient cascade is possible in this scenario, because at the point at which it is optimal for an individual to conform with the evidence derived from predecessors (through observation of actions and payoffs) despite an opposing private signal, he is still uncertain about which alternative is better. For example, early payoffs may indicate a high expected value for project 1 relative to project 2 even if 2 is better.

(iii) In addition, with positive probability the cascade continues forever. For example, if individuals are cascading on project 1, they do not generate new information about the true mean value of project 2 (given the actual state). This information blockage can prevent people from ever discovering the superiority of project 2. If the state on project 1 is favorable, individuals will tend to continue receiving good news about project 1, in ignorance of the even higher true mean payoff of the alternative project. There is a positive probability that a sequence of low payoffs on project 1 will trigger a temporary or permanent switch to project 2. However, at any point in time there is also a positive probability that no further switch ever occurs.

(iv) It is also possible that the cascade is broken endogenously, because observation of payoffs is noisy. With a long enough sequence of bad payoff outcomes, the cascade can be broken. This holds for both correct and incorrect cascades.

The following example illustrates how inefficient cascades/herding can occur on either project, and how there is a positive probability less than one of an inefficient cascade on either project lasting forever.

Numerical Example of Inefficient Cascades with Heterogeneous Payoffs

Individual 1 observes private signals about A and B . He therefore has 4 possible information outcomes based on high versus low signals about each project. Individual 2 observes the payoff received from the project that individual 1 adopts, as well as his own private signals about A and about B . Therefore, conditional on individual 1's project choice, individual 2 has 2^3 possible observation outcomes.

Projects A and B each have possible payoffs of 0 or 1. Each project can be in the up (u) or down (d) state, which are equally likely. Conditional on the state, the payoff distribution is: $\Pr(v_A = 1|u_A) = 0.8$, $\Pr(v_A = 1|d_A) = 0.2$, $\Pr(v_B = 1|u_B) = 0.7$, $\Pr(v_B = 1|d_B) = 0.3$. Thus, A 's ex ante expected payoff is more sensitive to state than B 's. Prior to choosing on the project, each individual receives signals regarding the state of the two projects. $\Pr(H_A|u_A) = 0.8$, $\Pr(H_A|d_A) = 0.2$, $\Pr(H_B|u_B) = 0.6$, $\Pr(H_B|d_B) = 0.4$. Since the signal on project A is more precise than that of project B , he chooses A if and only if his A signal is High.

Now suppose that the true state is u_A, u_B , so that A is superior. If the first individual receives a Low signal about A , he chooses B . Suppose he obtains a payoff of 1 from B . The second individual infers from the first individual's action that the first signal on project A was Low. So even if the second individual's private signal on project A is High, his conditional expectation of A is 0.5. Thus, the second individual will choose project B even if he receives a Low private signal on project B ; the favorable payoff of the first individual is such a favorable indicator (precision .7) that it outweighs his own private signal (precision .6). Therefore the second individual is in an inefficient cascade— his action is incorrect, and

does not reflect his private information. There is a positive probability that the cascade will be overturned as later individuals receive observation about payoffs from the adopted projects. If a sufficiently long sequence of 0 payoffs from B occurs, cascade will be overturned since posterior expectation of project value will approach 0.3. Nevertheless, since the true mean of B is 0.7, the probability that the posterior mean of B will stay above 0.5 forever is positive.¹⁶ With positive probability, the inefficient cascade lasts forever. \square

As in the basic single-person multi-arm bandit model of Rothschild (1974), it can be socially optimal to desist from experimenting between different choice alternatives even though substantial uncertainty remains about which is more desirable. However, as in the multi-person model of Bolton and Harris (1999), the amount of experimentation is suboptimal, because of the positive informational externality that such experimentation confers upon later observers. In contrast with multi-arm bandit models, here individuals do not just view past payoffs, they also spontaneously receive conditionally independent private signals about value. Unlike the payoff outcomes, individuals receive private signals about each project regardless of what actions were taken. If these signals were reasonably well aggregated, society would converge to accurate decisions. We find that even though both information sources are present (private signals and public information about payoffs), there is still a failure to converge to efficiency.

3.2 No Transient Shocks

We have shown that inefficient cascades/herding can occur and can last forever when individuals can observe actions and payoffs, and that these cascades/herding have a positive probability of being broken. The possibility of cascades/herding breaking spontaneously is somewhat different from the result of Bikhchandani, Hirshleifer and Welch (1992) and Banerjee (1992). In their model, if all individuals are identical and there are no interfering shocks, a cascade once started lasts forever. A similar outcome applies in our model only

¹⁶Beliefs evolve according to the net number of 1 versus 0 profit signals. This number follow a random walk with a positive drift, which need never cross a fixed lower bound— see, e.g., Chung (1974), page 263.

in the special case in which payoffs depend on the state nonstochastically. The reason is that once a cascade starts observing further identical payoffs does not reveal any new information. With no further information arriving, all succeeding individuals join the cascade. This proves the following corollary.

Corollary *If payoffs are non-stochastic functions of state and individuals observe all past actions and project payoffs, then cascades/herding once started last forever.*

When payoffs are non-stochastic functions of state, inefficient cascades/herding are extremely fragile in the sense that a small shock (such as a new public information arrival) can easily shift the long run outcome. If the shock persuades just one individual to try the alternative project, then everyone thereafter shifts to the better project.¹⁷ An individual can be sure that he is in a correct cascade only if both projects have been tried previously or if the payoff of current project is higher than the highest possible value from the untaken projects.

3.3 Noisy Observation of Payoffs

So far we have assumed that payoffs are perfectly observable. Proposition 1 and the Corollary generalize immediately to the case in which individuals observe noisy signals about past project payoffs. We describe the model informally here, and provide formal details and proof in the appendix.

Suppose that each individual observes a set of noisy signals about each of the payoffs of previously-chosen projects. Each signal is equal to the payoff plus independent noise. After observing the past payoff of a project, for two reasons an individual may still not know the true mean payoff of the project (given the true state). First, the realized payoff can be a noisy indicator of the true mean. Second, the individual observes a noisy indicator of the

¹⁷Another shock to the system that could break cascades/herding is the arrival of misfit individuals who have different information, payoffs or beliefs. For example, Bernardo and Welch (1999) show that the presence of overconfident ‘entrepreneurs’ can break cascades/herding and thereby increase the welfare of others. Hirshleifer and Noah (2000) examine the effects of several different kinds of misfit individuals.

realized payoff. Thus, this setting with both stochastic state-conditional payoff and noisy observation of payoff is equivalent to a corresponding setting where there is no observation noise but in which the noise in the state-conditional payoff as an indicator of long-run mean in the no-observation-noise regime in each state is equal to the sum of the noise of the state-conditional payoff and the observation noise in the observation-noise regime. In either scenario, individuals at each stage possess precisely the same information about project states, and therefore behave in an identical fashion. We therefore have:

Proposition 2 *If Assumption 1 holds, individuals can observe past actions, decision payoffs are stochastically heterogeneous and the observation of payoffs is noisy, then as the number of individuals increases, (i) the probability that a cascade eventually starts approaches one; (ii) inefficient cascades occur with positive probability; (iii) there is a positive probability less than one that cascades/herding last forever.*

As pointed out by Banerjee and Fudenberg (1999), observation of a small subsample of past actions causes subsets of observers to have different beliefs. If some individuals by chance observe a subset of actions that includes the less-popular action, they may be inclined to make use of signals about payoffs, improving information aggregation. In a sequential decision setting, an individual who observes a subset of past actions and/or payoffs that is relatively adverse to the currently-most-popular action may break a cascade. Again, this would cause further information to be aggregated. This raises the question of whether information blockages and cascades can persist in the long run when there is random sampling of previous decisionmakers. We do not attempt to model this question here.¹⁸

¹⁸It is clear that if the probability of each predecessor being in the subsample is close to 1, then cascades/herding will still last a long time. More generally, in a setting where individuals only take subsamples among the most recent K predecessors, it is plausible that the system may reach a point where all recent individuals choose the same action, suggesting that cascades may persist forever. We have verified in a special example in which there is random sampling that inefficient cascades still occur and with high probability last forever.

4 Observation of Payoffs Can Reduce Average Welfare and Decision Accuracy: An Example

Long-run efficient decisionmaking is a stringent benchmark. It requires that almost everyone do the right thing almost for sure. One might hope that even if a structure in which payoffs are observable is not long-run efficient, that observing payoffs at least increases welfare relative to not observing payoffs. We will see that this is not always the case.

Consider two projects A and B . The payoff of A is 0 and $2 + \epsilon$ with equal probability, and the payoff of B is $1 + \epsilon, 1 - \epsilon$ with equal probability, where ϵ is a small number as defined later. Each individual observes a signal about A with precision p and no signal about B .¹⁹ We compare the welfare outcomes of two scenarios. In the first, individuals do not observe any payoffs from the adopted projects. In the second scenario, individuals can observe payoffs from the B project if it is adopted but individuals cannot observe payoffs from project A .²⁰ Since there is no observation noise, once the payoff of B is observed it is known precisely thereafter.

Assumption A: We assume that

$$\epsilon < \frac{2p - 1}{2 - p}.$$

Assumption A ensures that if all individuals had to make decisions independently, they would adopt project A without any signals or with an H signal, and adopt B with an L signal.

Result 1 *In the numerical example, allowing individuals to observe payoffs can on average make individuals worse off and decisions less accurate.*

It seems reasonable that allowing an additional means of information transmission, conversation, would improve information flow and lead to improved decisions. Result 1 shows that observing payoffs can potentially lock individuals in an inefficient cascade/herd. While

¹⁹We can allow individuals to observe a very noisy signal of B without affecting any of the results.

²⁰Similarly, we can allow individuals to observe a very noisy signal of the predecessor's realized payoff from adopting project A without affecting any of our results here.

the observation of payoffs provides an individual with valuable information, it can reduce the informational content of the earlier individual's action. Whether later individuals are better off or worse off with payoff observations thus depends on the tradeoff between the observation of payoffs and the possible reduction of information conveyed by previous actions.

Previous papers have pointed out that greater information flow can reduce ex ante welfare. In Hirshleifer (1971) disclosure can reduce welfare by reducing risk-sharing. Teoh (1997) finds that disclosure can reduce ex ante welfare by inhibiting contributions in a public good game. In settings with informational cascades/herding, improved information flows do not necessarily improve overall decisions. For example, Bikhchandani, Hirshleifer and Welch (1992) find that the prospect of a public information disclosure can reduce average welfare. A difference between our model and previous papers is that the amount of additional information generation is endogenous. The ability to observe past payoffs could potentially resolve all uncertainty about which action is superior and thereby bring about the correct action choice in the long run, if only there were sufficient experimentation.

The problem with observation of past payoffs is that this can cause individuals to settle into cascades/herding even earlier than they would have otherwise. This premature cascading can be undesirable. This suggests that if observation of payoffs is delayed, later individuals may sometimes obtain the benefit of observing payoffs without the cost of reduced informativeness of predecessors' actions and payoffs. This suggests that observing payoffs with a delay can make everyone better off. We now modify the example by assuming that the payoffs of project B can be observed with one period delay.

Result 2: *In the numerical example, delaying the observation of payoffs can improve the accuracy of almost all individuals' decision and average welfare.*

According to Alexander Pope in his "Essay on Criticism,"

A little learning is a dangerous thing;
Drink deep or taste not the Pierian spring.
There shallow draughts intoxicate the brain,

Whilst drinking deeply sobers it again.

Taking observation of past actions as a starting point, a deep draught in our setting would be observation of past private signals, which leads to correct decisions in the long run. A shallow draught of learning, observing past payoffs (but not past signals), can be harmful. Observing payoffs only with delay is an even shallower draught of learning. In our setting this shallower draught can be the better one.

5 Conclusion

Convergent behavior often seems highly faddish and error-prone, suggesting that information gets lost when individuals try to assess alternatives by observing the choices of others. An individual who imitates a predecessor despite an opposing private signal makes his own action less informative to later observers than it might otherwise be. Such information blockages have been analyzed in a recent literature on what has been termed informational cascades (see, e.g., Banerjee (1992), Bikhchandani, Hirshleifer, and Welch (1992)). Research on informational cascades/herding has focused on settings where rational individuals spontaneously receive private information about the value of discrete action alternatives, and observe the actions of predecessors.

Another strand of research on word-of-mouth learning has allowed for observation of *payoff outcomes*, but in settings that differ in important respects (such as a continuum of agents, imperfect rationality, small observation samples, or private information only about past payoffs). This approach is consistent with the casual intuition that many decisions, such as choices of restaurants, movies, and repairmen, are influenced by conversation, and the empirical evidence of strong conversational influences in choices of consumer products, business factor inputs, and security investment decisions. Banerjee and Fudenberg (1999) describe fairly mildly assumptions under which a conversational learning system evolves deterministically to the efficient outcome.

We consider here a scenario in which, like previous cascades/herding models, involve potentially path-dependent long-term outcomes, but which, like the word-of-mouth models,

also allows for ‘conversational learning’: individuals can observe the payoffs as well as the actions of predecessors. We find informational cascades/herding still occurs with probability one, and describe conditions under which cascades aggregate information inefficiently and are fragile. In our setting decisions evolve in a stochastic, time-dependent fashion. We show that a cascade once started has a positive probability of lasting forever. We also provide examples to show that the ability to observe/communicate past payoffs can reduce average welfare and the accuracy of action choices, and that delay in observation of payoffs can increase welfare.

The rise of mass media and of interactive communication technology (printing, telegraph, telephone, television, email, the web) have made it easier to observe the payoff outcomes of others (or summary statistics of past payoffs).²¹ The analysis here suggests that the resulting improvements in decisions may not be as great as might have been expected. The recipient of extra information (beyond what he would have obtained by observing past actions) gains directly from receiving more information. However, this benefit may be opposed by the fact that the actions of the recipient may be less informative to later decisionmakers.

²¹These advances have also made it easier to observe past actions and to communicate private information. The discussion here does not try to address the full scope of these changes.

Appendix

Proof of Proposition 1: The proof of Proposition 1 requires the strong law of large numbers (see DeGroot 1970, p. 203).

The Strong Law of Large Numbers. Let Z_1, Z_2, \dots be a sequence of independently and identically distributed random variables with mean λ . Then, with probability one,

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{k=1}^n Z_k = \lambda.$$

We consider the general case of N projects. For each project n , it could be in an up state or a down state represented by $s_n = \{u_n, d_n\}$. Each project has M possible payoffs, $v_{n1} < v_{n2} < \dots < v_{nM}$,²² and the project payoff \tilde{v}_n^i is stochastic conditional on the project's state, and its distribution depends on the state but not on the individual. The project state s_n and project values are uncorrelated across projects. An i superscript is used to denote the payoff of the project if adopted by individual i . The prior probability that $\tilde{v}_n^i = v_{nj}$ given state of project n , s_n , is denoted $\mu_j(s_n)$, $j = 1, \dots, M$. The payoffs of different projects and the indicator variable of states of different projects are uncorrelated.

We apply the concept of perfect Bayesian equilibrium.²³ Individuals observe predecessors' actions and payoffs of the projects they have chosen. Individuals differ not only by their positions in the queue but by the signals they privately observe. Each individual i observes a vector of conditionally independent and identically distributed signals $\tilde{\Sigma}_n^i = (\tilde{\sigma}_1^i, \tilde{\sigma}_2^i, \dots, \tilde{\sigma}_N^i)$, where $\tilde{\sigma}_n^i$ is a signal about project n , with possible values $\sigma_{n1} < \sigma_{n2} < \dots < \sigma_{nM}$. Let $p_q(s_n)$ be the probability that an individual observes signal realization σ_{nq} given that the true state of project n is s_n .

We now prove that with probability one a cascade will start. Suppose that a cascade has not started till individual $i + 1$. By definition, each of the preceding individuals made

²²We can allow the number of payoff values and number of signal realizations to be different but that will increase notational complexity.

²³Since an individual's payoffs do not depend on what later individuals do, there is no incentive to make an out-of-equilibrium move to try to influence a later player. Thus, without loss of generality, we assume that if any player is observed to deviate from the equilibrium, then subsequent individuals have the same beliefs as though he had chosen his correct action.

use of their signals. For each preceding individual, his action provides some information about his signal. Let A^i denote the history of actions, Let B^i denote the history of payoffs from the adopted projects. Let $J^i(A^{i-1}, B^{i-1}, a^i)$ be the set of signals such that individual i will choose project a^i . Let $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_N)$ denotes a realization of the signals. Let Σ denote the spaces composed of all σ , then $J^i(A^{i-1}, B^{i-1}, n)$ forms an ordered N -partition of Σ . The action choice a^i of individual i reveals his signal belongs to $J^i(\cdot, \cdot, a^i)$.

Let Σ_N be the set of ordered N -partitions of Σ , $T = (T_1, T_2, \dots, T_N)$ be an element of Σ_N and $S = (s_1, s_2, \dots, s_N)$ denote the state of the N projects. Let $P_{T_j s_n}$ denote the conditional probability of the signal σ to be in T_j given that the true state of project n is s_n :

$$P_{T_j s_n} \equiv \Pr(\sigma, \sigma \in T_j | s_n).$$

For an arbitrary partition T , let j_T denote the set of individuals such that $J^k(\cdot, \cdot, n) = T_n$ for all n . Let m_T denote the number count in set j_T . Let S^* denote the true state of the world. Let μ_{s_n} denote the prior probability that the state of project n is s_n . Let $p_{\sigma s_n}$ denote the probability of receiving a σ signal when the state of project n is s_n . Let i_n denote the set of individuals who have taken project n and m_n denote the number count of i_n . Let $p_{v_j^i s_n}$ denote the conditional probability that individual i 's payoff on project j is v_j^i given the state on project n is s_n . Then individual $i + 1$'s posterior probability that the state of project n is s_n conditional on the action and payoff history of his predecessors A^i, B^i and on his signal realization $\sigma^{i+1} = \sigma$, is

$$\Pr\{s_n | A^i, B^i, \sigma^{i+1} = \sigma\} = \frac{\mu_{s_n} p_{\sigma s_n} \prod_{T \in \Sigma_N, m_T > 0} \prod_{k \in j_T} P_{T_{a^k s_n}} \prod_{j=1}^N \prod_{i \in i_j} p_{v_j^i s_n}}{\sum_{s_n} \mu_{s_n} p_{\sigma s_n} \prod_{T \in \Sigma_N, m_T > 0} \prod_{k \in j_T} P_{T_{a^k s_n}} \prod_{j=1}^N \prod_{i \in i_j} p_{v_j^i s_n}}$$

Let $L(s_n)$ denote the log of likelihood ratio between states s_n and s_n^* , then

$$L(s_n) = \log \left[\frac{\Pr\{s_n | A^i, B^i, \sigma^{i+1} = \sigma\}}{\Pr\{s_n^* | A^i, B^i, \sigma^{i+1} = \sigma\}} \right] = l_{\mu s_n} + l_{\sigma s_n} + \sum_{T \in \Sigma_N, m_T > 0} \sum_{k \in j_T} l_{T_{a^k s_n}} + \sum_{i \in i_n} l_{v_n^i s_n} \text{ where } l_{\mu s_n} \equiv \log \left[\frac{\mu_{s_n}}{\mu_{s_n^*}} \right]$$

We show that given the true state on project n is s_n^* , as the number of individuals who choose project n goes to infinity, then with probability one,

$$\lim_{m_n \rightarrow \infty} \Pr\{s_n | A^i, B^i, \sigma^{i+1} = \sigma\} = \begin{cases} 1 & \text{if } s_n = s_n^*, \\ 0 & \text{if } s_n \neq s_n^*. \end{cases}$$

Let

$$\lambda_{Ts_n} \equiv E \left[\log \left(\frac{P_{T_{a^i s_n}}}{P_{T_{a^i s_n^*}}} \right) \middle| s_n^* \right].$$

$$\theta_{s_n} \equiv E \left[\log \left(\frac{p_{v_n^i s_n}}{p_{v_n^i s_n^*}} \right) \middle| s_n^* \right]$$

If $s_n \neq s_n^*$, then Jensen's inequality imply that

$$\lambda_{Ts_n} \leq \log \left(E \left[\frac{P_{T_{a^i s_n}}}{P_{T_{a^i s_n^*}}} \middle| s_n^* \right] \right) = \log \left(\sum \frac{P_{T_{a^i s_n}}}{P_{T_{a^i s_n^*}}} P_{T_{a^i s_n^*}} \right) = \log \left(\sum_{a^i} P_{T_{a^i s_n}} \right) = 0.$$

Similarly, if $s_n \neq s_n^*$, then

$$\theta_{s_n} < \log \left(E \left[\frac{p_{v_n^i s_n}}{p_{v_n^i s_n^*}} \middle| s_n^* \right] \right) = 0$$

Thus the strong law of large numbers implies that with probability one

$$\lim_{m_T \rightarrow \infty} \frac{1}{m_T} \sum_{k \in j_T} \log \left(\frac{P_{T_{a^i s_n}}}{P_{T_{a^i s_n^*}}} \right) = \lambda_{Ts_n} \leq 0.$$

$$\lim_{m_n \rightarrow \infty} \frac{1}{m_n} \sum_{k=1}^{m_n} \log \left[\frac{p_{v_n^i s_n}}{p_{v_n^i s_n^*}} \right] = \theta_{s_n}.$$

Hence

$$\frac{1}{m_T} \log \left(\lim_{m_T \rightarrow \infty} \prod_{k=1}^{m_T} \frac{P_{T_{a^i s_n}}}{P_{T_{a^i s_n^*}}} \right) = \lim_{j_T \rightarrow \infty} \frac{1}{j_T} \sum_{k=1}^{j_T} \log \left(\frac{P_{T_{a^i s_n}}}{P_{T_{a^i s_n^*}}} \right) = \lambda_{Ts_n} \leq 0.$$

$$\frac{1}{m_n} \log \lim_{m_n \rightarrow \infty} \prod_{i \in i_n} \frac{p_{v_n^i s_n}}{p_{v_n^i s_n^*}} = \lim_{m_n \rightarrow \infty} \frac{1}{m_n} \sum_{i \in i_n} \log \left[\frac{p_{v_n^i s_n}}{p_{v_n^i s_n^*}} \right] = -\infty.$$

Consequently, if $s_n \neq s_n^*$, then, with probability one

$$\lim_{j_T \rightarrow \infty} \prod_{k=1}^{j_T} \left(\frac{P_{T_{a^k S}}}{P_{T_{a^k S}}} \right) \leq 1$$

$$\lim_{m_n \rightarrow \infty} \prod_{k=1}^{m_n} \frac{p_{v_n^i s_n}}{p_{v_n^i s_n^*}} = 0$$

As a result, we have:

$$\lim_{m_n \rightarrow \infty} L(s_n) = -\infty$$

which implies that

$$\lim_{m_n \rightarrow \infty} \Pr\{s_n^* | A^i, B^i, \sigma^{i+1}\} = 1$$

If $i \rightarrow \infty$, then since $\sum_n m_n = i$, there exists at least one n such that $m_n \rightarrow \infty$. In this case, $\Pr(s_n^* | A^i, B^i, \sigma^{i+1})$ goes to 1. Notice that there can be only one n such that $m_n \rightarrow \infty$. We prove this by contradiction. If there were an $n' \neq n$ and both $m_n \rightarrow \infty$ and $m_{n'} \rightarrow \infty$, then with probability one we would know the payoff of project n and n' .

$$\lim_{m_n \rightarrow \infty} \Pr\{s_n = s_n^* | A^i, B^i, \sigma^{i+1}\} = 1$$

$$\lim_{m_{n'} \rightarrow \infty} \Pr\{s_{n'} = s_{n'}^* | A^i, B^i, \sigma^{i+1}\} = 1$$

Suppose that when the true state is known, project n dominates project n' , i.e., $\bar{v}_{nS^*} > \bar{v}_{n'S^*}$. Conditional on project n being chosen infinitely often, the probability that project n' will be chosen infinitely often is zero as payoff information will reveal that project n' is strictly inferior. Thus the probability that two projects will alternate infinitely often is zero. Consequently all individuals will herd on a single project and will make their decision regardless of their private signals eventually with probability one. If there are ties, then it is possible that individuals will switch back and forth between two projects infinitely often.

(ii) Let u_n denote the up state of project n and d_n denote the low state of project n . Let $\bar{v}_{nu}, \bar{v}_{nd}$ correspond to the up and down states. Suppose the true state is s . Without loss of generality, we assume the first project has the highest possible expected payoff given the state of the economy, that is, $\bar{v}_{1s} > \bar{v}_{2s} > \dots > \bar{v}_{Ns}$. We assume that the conditional expected payoff of the down state in project 1 does not dominate the prior expected payoffs from other projects, that is:

$$\bar{v}_{1d} < \max_{j \neq 1} \{E[v_j | \sigma^M]\}.$$

This assumption implies that individuals will not choose project 1 if it is known that project 1 is in the down state and the best signals of other projects are received. Suppose the first individual chooses a project other than project 1, project k . We show that with positive probability, this can start an inefficient cascade will start. We prove this in several steps.

Definition A run on project k with the highest payoff on project k , v_{kM} observed is called

an M-run on project k . First of all, for any positive number L , an M-run on project k with length L has positive probability,

$$\Pr(a^i = k, i = 1, \dots, L, v_k^i = v_{kM}) > 0$$

This is because with positive probability, individuals will receive the maximum payoff and signal for the adopted project k and worst possible signals for all other projects. This means that with positive probability, a sequence with arbitrary length of consecutive adoption of the same project which gives out the highest possible payoffs will be observed. Let $\mu_k(L) \equiv E[v_k | v_k^i = v_{kM}, i = 1, \dots, L]$, it is easy to verify that,

$$\mu_k(L + 1) \geq \mu_k(L)$$

that is the conditional expectation of project k 's payoff along such a sequence is a submartingale. Similarly, one can show that $\mu_j(L)$ is a supermartingale along such a sequence. This is because individuals believes that project k is more likely to be of the high state when the best possible payoff on project k is observed. Next we show that there exists a number L , such that if a sequence of L consecutive adoptions of project k with payoff v_{kM} (the largest possible payoff on project k) is observed, then a cascade will start.

Lemma 1 *For project $j \neq k$, there exists a number L_j , such that after a run of L_j number of project k with maximum payoff v_{kM} , project j will be dominated by project k , that is the next individual $L_j + 1$ will not choose project j no matter what signal he receives.*

Assertion (ii) holds following lemma 1 if we take $L = \max_{j \neq k} \{L_j\}$. Now we are ready to prove lemma 1. Let $l_j(L)$ denote the likelihood ratio of project j of the down state versus the up state given the public information after an M-run on project k with length L . Then it is easy to verify that $l_k(L)$ is strictly decreasing and goes to zero as L goes to infinity. This is because both the fact that project k is taken and the best payoff on project k is observed provides favorable information about project k which increases the log likelihood ratio $\log(l_k)$ by a finite amount. Now lets check the evolution of the likelihood ratio on project j in an M run on project k ,

$$l_j(L + 1) = l_j(L) \frac{\pi_{uk}(L) \sum_{m=1}^M \pi_{ukm} \hat{\pi}_{djm}(L + 1) + \pi_{lk}(L) \sum_{m=1}^M \pi_{dkm} \hat{\pi}_{djm}(L + 1)}{\pi_{uk}(L) \sum_{m=1}^M \pi_{ukm} \hat{\pi}_{ujm}(L + 1) + \pi_{lk}(L) \sum_{m=1}^M \pi_{dkm} \hat{\pi}_{ujm}(L + 1)} \quad (1)$$

where $\hat{\pi}_{sjm}(L+1)$ denote the probability that individual $L+1$ would choose project k conditional on a signal of σ_{km} is observed on project k and state on project j is s . $\pi_{uk}(L)$ denote the conditional probability that the state of project k is up and $\pi_{lk}(i)$ denote the conditional probability that the state of project k is down.

We can rewrite the expression of the likelihood ratio as:

$$\frac{l_j(L+1)}{l_j(L)} = 1 + \frac{\sum_{m=1}^M (\pi_{ukm} + l_k(L)\pi_{dkm})(\hat{\pi}_{dj1}(L+1) - \hat{\pi}_{uj1}(L+1))}{\pi_{uk}(L) \sum_{m=1}^M \pi_{ukm} \hat{\pi}_{ujm}(L+1) + \pi_{lk}(L) \sum_{m=1}^M \pi_{dkm} \hat{\pi}_{ujm}(L+1)} \quad (2)$$

By the monotone likelihood ratio of the signals, we have:

$$\frac{\hat{\pi}_{dj1}}{\hat{\pi}_{uj1}} \geq 1$$

Suppose that after L individuals have chosen project k and observed the best possible payoff from project k , it is still possible for the next individual to choose project j . Thus if individual $L+1$ receives the worst signal on project k , then it is possible that he may choose a project other than k . We show that individual $L+1$ will choose project j over k only if he receives a signal better than the worst signal on project j .

Let $\mu_k(L), \mu_j(L)$ denote the conditional expectation on project j, k respectively given public information, then clearly μ_k is strictly increasing and μ_k is weakly decreasing in an M-run on project k . Therefore $\mu_k(L) > \mu_k(1) \geq \mu_j(1) \geq \mu_j(L) \geq E[v_j | F_L, \sigma_j^{L+1} = \sigma_{j1}]$. where F_l denote the public information set from observing first L individual's actions and payoffs from their adopted projects. When L is large, π_u is arbitrarily close to one, and thus $E[v_k | F_L, \sigma_k^{L+1} = \sigma_{k1}]$ is arbitrarily close to $\mu_k(L)$ which is strictly larger than $\mu_j(L)$. Consequently, for large L , when the signals individual $L+1$ receive on project j, k are both worst signals, he will choose project k . Therefore $\hat{\pi}_{dj1}(L+1)$ is strictly positive and strictly larger than $\hat{\pi}_{uj1}(L+1)$ by an amount larger than a fixed positive number. As a result, there is exist a number $\delta > 0$ such that:

$$\frac{l_j(L+1)}{l_j(L)} > \delta > 1$$

for $L > L_1$. Thus along an M-run on project k , l_j goes to infinity and eventually individuals believe that it is almost certain the j project is in the down state. Since

$$\lim_{L \rightarrow \infty} \mu_k(L) = \bar{v}_{uk} > \bar{v}_{dj} = \lim_{L \rightarrow \infty} \hat{\mu}_j(L)$$

where $\hat{\mu}_j$ denote the conditional expectation on project j given the best signal on project j , σ_{jM} . Therefore there exist an L_j such that project k will dominate project j after a run of length L_j on project k . The same analysis holds for all other projects. Therefore all projects will be dominated by project k after a long run on project k and an inefficient cascade will eventually occur. Suppose the first individual always choose project 1 regardless of his signal, then with sufficient number of bad payoffs from project 1, individuals are going to switch and then an inefficient cascade can form following the same argument.

(iii) From part (i)(ii), we know that with positive probability, a cascade can start. Let k denote the adopted project, and suppose the true state of k is the up state. This cascade is inefficient if there are other projects with higher expected payoffs when the state of the projects is known. We prove such an inefficient cascade can last forever. At the beginning of cascade i , let μ denote the highest expected value of other projects given the best possible signals

$$\mu = \max_{j \neq k} \{E[v_j | A^i, B^i, \sigma^{i+1} = \sigma^M]\}$$

where σ_M corresponds to the best signals of all projects.

Now clearly we have $\mu_k^i > \mu$.

$$\mu_k^i = \pi_{ku}^i \bar{v}_u + \pi_{kd}^i \bar{v}_d \quad (3)$$

Let l_k^i denote the likelihood ratio of the down state versus the up state of project k . Then the likelihood ratio evolves according to the payoffs received from project k , that is

$$l_k^{i+1} = l_k^i * \frac{p_{v_k d}}{p_{v_k u}}$$

which implies

$$\log(l_k^{i+1}) = \log(l_k^i) + \log \left[\frac{p_{v_k^{i+1} d}}{p_{v_k^{i+1} u}} \right]$$

Let x_k^{i+1} denote the change in log likelihood ratio, then

$$E[x_k^{i+1} | u_k] = E \left[\log \left(\frac{p_{v_k^{i+1} d}}{p_{v_k^{i+1} u}} \right) | u_k \right] < \log \left(E \left[\frac{p_{v_k^{i+1} d}}{p_{v_k^{i+1} u}} | u_k \right] \right) = \log(1) = 0$$

Thus the log of likelihood ratio follows a generalized random walk with a downward drift. This implies that (Chung (1972) page 263), the log likelihood ratio may not come back to the initial point forever with positive probability. Since the expected payoff is a linear function of the likelihood ratio, the expected payoff is larger than μ_l^i forever with positive probability and an inefficient cascade can last forever.

(iv) When payoffs are stochastic and the conditional payoff from the down state of the adopted project k is strictly smaller than the maximum expected payoff of all other projects given the best signals, a long series of bad payoffs can cause individuals believe that project k is of down state and a cascade can break when the best signals of other projects are received. In this case the probability a cascade will last forever is less than one. \parallel .

Proof of Proposition 2: The payoff for individual i is $v_n + \epsilon_{ni}$. The random variable ϵ_{ni} has zero mean, can take values $\epsilon_{n1} < \epsilon_{n2} < \dots < \epsilon_{nK}$. We assume that ϵ is i.i.d across individuals. Furthermore we assume that after individual i has taken his project, a *public* noisy signal, s_n , about the payoffs becomes available, with possible values $s_{n1} < s_{n2} < \dots < s_{nK}$. The probability of observing s_{nq} given that the true payoff of $v_{nl} + \epsilon_{nlm}$ is p_{nqlm}^s . Thus, the probability of observing s_{nq} given the true state is $v_n = v_{nl}$ is:

$$p_{nql}^s = \sum_{m=1}^K p_{nlm}^\epsilon p_{nqlm}^s.$$

We assume that for $l \neq l'$, p_{nql}^s is different from $p_{nql'}^s$ for at least one q . Notice that individuals are risk neutral and their decision rule depends only on the expected payoff of project n which is v_n . The public signal s_{ni} is a noisy signal of the expected payoff v_n . Therefore the model is isomorphic to the model in Proposition 1 and the Proposition 2 follows as a result.

Proof of Result 1: We first consider the case of no observation of payoffs. Consider the subcase in which project A's payoff is $2 + \epsilon$. If the first individual receives an H signal, then he will choose project A. The next individual will choose project A even if he receives an L signal since the prior mean of A is higher than the prior mean of project B. There is an A cascade immediately. If the first two individuals receive two L signals, then the next individual will choose project B independent of his signal and a B cascade forms. Finally,

if the sequence of the signals is LH , the third individual faces the same situation as the first individual. Starting after the first two individuals, the probability of an A cascade is p , the probability of a B cascade is $(1 - p)^2$, and the probability of no cascade is $p - p^2$. Consequently, the ratio of the likelihoods of an A cascade and a B cascade is $p/(1 - p)^2$. Since a cascade will form with probability one, the long-run probability of an A cascade is $\frac{p}{1-p+p^2}$, and the long run probability of an B cascade is $\frac{(1-p)^2}{1-p+p^2}$. Alternatively, when the payoff of project A is zero, the long-run probability of an A cascade is $\frac{1-p}{1-p+p^2}$, and the long-run probability of a B cascade is $\frac{p^2}{1-p+p^2}$.

Thus, the expected payoff of an individual late in the queue is approximately

$$W_0 = 0.5 \left[\frac{p}{1-p+p^2}(2 + \epsilon) + \frac{(1-p)^2}{1-p+p^2} \right] + 0.5 \left[\frac{p^2}{1-p+p^2} \right].$$

In contrast, if individuals can observe the payoffs of project B , a cascade forms immediately after one round if the payoff of project B is $1 + \epsilon$. This is because if the signal is H , the first individual will choose A and there will be no observation of project B thereafter. If the signal is L , B will be adopted and its payoff $1 + \epsilon$ will stop anyone who wants to try A . If the payoff of project B is $1 - \epsilon$, the analysis is the same as the case without payoff observations. Consequently the expected payoff of an individual later in the queue is:

$$\begin{aligned} W_p &= 0.25[p(2 + \epsilon) + (1 - p)(1 + \epsilon)] + 0.25[p(1 + \epsilon)] \\ &\quad + 0.25 \left[\frac{p}{1-p+p^2}(1 + 2\epsilon) + 1 - \epsilon \right] + 0.25 \left[\frac{p^2}{1-p+p^2}(1 - \epsilon) \right] \end{aligned}$$

We have

$$W_0 - W_p = 0.25 \frac{p(1-p)^2}{1-p+p^2} \left[\frac{p}{1-p} - 1 - \epsilon \right] > 0$$

The last inequality follows by the assumed parameter restriction. \parallel

Proof of Result 2: We now assume that the payoffs of project B can be observed with one period delay. Let (v_A, v_B) denote the payoffs of project A and B . There are four possible payoff combinations:

CASE I $(2 + \epsilon, 1 + \epsilon)$ If the first signal is H , there is an A cascade since the payoffs of A is not observable. The probability of such a cascade is p .

If the first signal is L , the first individual adopts project B . The second individual adopts project B if he receives another L signal. A cascade forms since the third individual further observes the payoff of project B and adopts project B regardless of his signal. The probability of this cascade is $(1 - p)^2$.

If the signal sequence is LH , then the third individual infers the signal sequence and he also observes that the payoff of project B is $1 + \epsilon$. The conditional probability of a B cascade is $(1 - p)/(1 - p + p^2)$ and the conditional probability of an A cascade is $p^2/(1 - p + p^2)$. Consequently, the probability of an A cascade is

$$p_A = p + p(1 - p)p^2/(1 - p + p^2) = \frac{p - p^2(1 - p)^2}{1 - p + p^2}$$

Thus, the payoff of an individual far along in the queue is, in the limit,

$$W_I = \frac{p - p^2(1 - p)^2}{1 - p + p^2} + 1 + \epsilon.$$

The expected payoffs of individuals far down the line for the other three cases can be derived similarly:

Case II, $(0, 1 + \epsilon)$:

$$W_{II} = \frac{1 - p - p^2(1 - p)^2}{1 - p + p^2}(-1 - \epsilon) + 1 + \epsilon$$

Case III, $(2 + \epsilon, 1 - \epsilon)$:

$$W_{III} = \frac{p}{1 - p + p^2}(1 + 2\epsilon) + 1 - \epsilon$$

Case IV, $(0, 1 - \epsilon)$:

$$W_{IV} = \frac{p^2}{1 - p + p^2}(1 - \epsilon)$$

Finally, the expected payoff of an individual far down the line is in the limit

$$\begin{aligned} W_d &= 0.25[W_I + W_{II} + W_{III} + W_{IV}] \\ &= 0.25 \left\{ \frac{p - p^2(1 - p)^2}{1 - p + p^2} + 1 + \epsilon + \left[\frac{1 - p - p^2(1 - p)^2}{1 - p + p^2}(-1 - \epsilon) + 1 + \epsilon \right] \right. \\ &\quad \left. + \left[\frac{p}{1 - p + p^2}(1 + 2\epsilon) + 1 - \epsilon \right] + \left[\frac{p^2}{1 - p + p^2}(1 - \epsilon) \right] \right\}. \end{aligned}$$

It is straightforward to show that

$$W_d - W_0 = 0.25 \frac{p^2(1-p)^2}{1-p+p^2} \epsilon > 0.$$

Therefore, we have

$$W_d > W_0 > W_p.$$

||

References

- L. Anderson (1994), Information Cascades: A Logistic Error Model of Laboratory Data, working paper, University of Virginia.
- Anderson, Lisa, and Charles Holt, "Information Cascades in the Laboratory," *American Economic Review*, 1997, 87, 847-862.
- Aoyagi, M. "Mutual Observability and the Convergence of Actions in a Multi-Person Two-Armed Bandit Model," *Journal of Economic Theory*, 82(2), October (1998):405-424.
- Banerjee, Abhijit. "A Simple Model of Herd Behavior." *Quarterly Journal of Economics*, 107, August (1992):797-818.
- Banerjee, Abhijit and Fudenberg, Drew. "Word-of-Mouth Communication and Social Learning." Harvard University Working Paper, (1999).
- Bernardo, Tony and Welch, Ivo. "A Theory of Overconfidence and Entrepreneurship." UCLA Anderson School Working Paper, (1999).
- Bikhchandani, Sushil; Hirshleifer, David; and Welch, Ivo. "A Theory of Fads, Fashion, Custom and Cultural Change as Informational Cascades." *Journal of Political Economy*, 100(5), October (1992):992-1026.
- Bolton, Patrick and Harris, Christopher, "Strategic Experimentation." *Econometrica*, 67(2), March (1999): 349-74.
- Caplin, Andrew and Leahy, John. "Sectoral Shocks, Learning and Aggregate Fluctuations." *Review of Economic Studies*, 60, (1993):777-794.
- A. Caplin and J. Leahy (1994), "Business as Usual, Market Crashes, and Wisdom After the Fact", *American Economic Review*, 84, 548-565.
- Chamley, Cristophe and Gale, Douglas. "Informational Revelation and Strategic Delay in Irreversible Decisions", *Econometrica*, 62(5), September (1994):1065-1085.
- Corb, Howard M. "The Nature of Bank Runs." Graduate School of Business, Stanford University, (1993).

- G. Ellison and D. Fudenberg (1993), "Rules of Thumb for Social Learning," *Journal of Political Economy*, 101, 612–643.
- G. Ellison and D. Fudenberg (1995), "Word-of-Mouth Communication and Social Learning," *Quarterly Journal of Economics*, 110(1), 93-125.
- Gibson, Robert M. and Hoglund, Jacob. "Copying and Sexual Selection." *Trends in Research in Evolution and Ecology*, 7(7), July (1992):229-232.
- Grenadier, Stephen. "Information Revelation Through Option Exercise," *Review of Financial Studies*, 12(1), Spring (1999):95-129.
- Herr, Paul M., Kardes, Frank R., and Kim, John. "Effects of Word-of-Mouth and Product-Attribute Information on Persuasion: An Accessibility-Diagnosticity Perspective," *Journal of Consumer Research*, vol. 17, March (1991):454-62.
- Hirshleifer, David and Noah, Robert, "Misfits and Social Progress." Ohio State University manuscript, (1999).
- Hirshleifer, Jack. "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61(4), September (1971):561-574.
- Kennedy, Robert E., "Strategy Fads and Competitive Convergence: An Empirical Test for Herd Behavior in Prime-time Television Programming." Harvard Business School Working Paper 96-025, August 1999.
- Khanna, Naveen. "Optimal Contracting with Moral Hazard and Cascading." *Review of Financial Studies*, 11(3), (1998):559-96
- Khanna, Naveen and Slezak, Steve L. "The Effect of Organizational Form on Information Aggregation and Project Choice: The Problem of Informational Cascades." *Journal of Economics and Management Strategy*, 9(1), Spring (2000):115-56.
- Kuran, Timur and Cass R. Sunstein, "Availability Cascades and Risk Regulation." *Stanford Law Review* 51-4, April 1999, 683-768.
- Lee, In Ho (1993), On the Convergence of Informational Cascades, *Journal of Economic Theory*, 61, 395–411.

- Lee, In Ho, "Market Crashes and Information Avalanches", *Review of Economic Studies* 65(4), October (1998):741-59.
- Lohmann, Susanne. "Dynamic Informational Cascades: Theory and Experimental Evidence," Political Science Department, UCLA, (1997).
- Lohmann, Susanne, "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-91," *World Politics*, 1994, 47, 42-101.
- Perktold, Josef, "Recurring Informational Cascades," working paper, University of Chicago, 1999.
- Rothschild, Michael. "A Two-Armed Bandit Theory of Market Pricing," *Journal of Economic Theory*, 9 October (1974):185-202).
- Schlag, K., 1998 "Why Imitate and If So How? A Boundedly Rational Approach to Multi-Armed Bandits," *Journal of Economic Theory*, 78, 130-56.
- Shiller, Robert J., *Irrational Exuberance*, Princeton University Press, 2000.
- "Conversation, Information, and Herd Behavior," *American Economic Review* (1995), 85(2): 181-85.
- Shiller, Robert J. and Pound, John. "Survey Evidence on the Diffusion of Interest and Information Among Investors," *Journal of Economic Behavior and Organization*, 12 (1989):47-66.
- Smith, Lones and Peter Sorenson, "Pathological Outcomes of Observational Learning," *Econometrica*, 3, forthcoming,(2000).
- Taylor, Curtis R., "Time on the Market as a Sign of Quality," *Review of Economic Studies*, 555-578.
- Teoh, Siew Hong, "Information Disclosure and Voluntary Contributions to Public Goods." *Rand Journal of Economics*, 28(3), Autumn (1997):385-406.
- Vives, Xavier, How Fast Do Rational Agents Learn, *Review of Economic Studies*, 60, (1993):329–347.
- Welch, Ivo. "Sequential Sales, Learning and Cascades." *Journal of Finance* 47 (1992).
- Zhang, Jianbo, "Strategic Delay and the Onset of Investment Cascades," *Rand Journal of Economics*, 1997, 28, 188-205.