Information Cascade Experiments

Lisa R. Anderson and Charles A. Holt*

I. Cascades

The theoretical literature on "herding" pertains to situations where people with private, incomplete information make public decisions in sequence. Hence, the first few decision makers reveal their information, and subsequent decision makers may follow an established pattern even when their private information suggests that they should deviate. This type of "information cascade" can occur with perfectly rational individuals, when the information implied by early decisions outweighs any one person's private information. These theories have been used to explain fads, investment patterns, etc. (Bannerjee, 1992; Bikhchandani, Hirshleifer, and Welch, 1992). For example, a waiting line for a movie or restaurant may be enough to lure additional customers, even if they have heard mixed reviews from other sources. Economists are particularly interested in market implications of herding behavior, e.g., the possibility that investment booms and busts are analogous to information cascades. This paper surveys the results of experiments designed to evaluate cascade behavior with human subjects, both in simple "ball-and-urn" settings and in more complex, asset-market environments.

The concept of an information cascade can be explained in the context of a specific numerical example that was used in initial laboratory experiments (Anderson and Holt, 1997). In this example, there are two states of nature, A and B, which are equally likely, *ex ante*. Each decision maker obtains an independent, private signal, "a" or "b," which has a two-thirds chance of indicating the correct state, i.e. Pr(a|A) = Pr(b|B) = 2/3. The decision makers are selected in sequence and asked to predict the state, with a monetary reward paid for a correct prediction. The predictions are publicly announced as they are made, but individuals are not able to observe others' private signals. The first person in the sequence must predict only on the basis of private information. This person should predict the state indicated by the signal because, with a prior probability of 1/2 for each state, the posterior probability is 2/3 for the state that is signaled:

$$= \frac{Pr(a|A)\ Pr(A)}{Pr(a|A)\ Pr(A)\ +\ Pr(a|B)\ Pr(B)} = \frac{(2/3)\ (1/2)}{(2/3)\ (1/2)\ +\ (1/3)\ (1/2)} \qquad (1)$$

Therefore, the first decision reveals that person's signal, even though subsequent individuals see only the prediction, not the private information on which the prediction was based.

Without loss of generality, suppose that the first person sees an "a" signal and predicts A. If the second decision maker in the sequence also sees an "a" signal, then it is obvious that the optimal prediction is A. If the second person sees a "b" signal, then the inferred and observed signals cancel each other, and the posterior probability is exactly 1/2, as can be verified by Bayes' rule. If we assume that people make a prediction based on their own private information when the posterior is 1/2, then the second decision also reveals the associated private signal, regardless of whether or not it conforms to the first prediction. Therefore, two initial predictions of A reveal two "a" signals, which loosely speaking, are more informative than the private signal seen by the third person in the sequence, even if this is a "b". Whenever the first two predictions match, the third person should follow.[1] This is how a cascade develops; the third person's decision does not reveal his or her private draw in this case, and the fourth person makes a decision based on the same prior information as the third. Thus a string of matching decisions can create a false degree of certainty, since all are driven by the first two predictions when they match.

Anderson and Holt (1997) implemented this setup in an experiment by putting balls labeled "a" or "b" in urns labeled A and B, with three labeled balls in each:[2]

$$Urn\ A:\ \ a,\ a,\ b \qquad Urn\ B:\ \ b,\ b,\ a$$

The urns were equally likely to be chosen by the throw of a six-sided die. A throw of 1, 2, or 3 determined that urn A would be used for the draws, and a throw of 4, 5, or 6 determined that urn B would be used. Hence, each of the 6 balls were, *ex ante*, equally likely to be drawn. Since 2 of the 3 balls labeled "a" were in urn A, the posterior probability of event A given signal "a" is 2/3. Similarly, the posterior probability of event A given signal "b" is 1/3.[3]

Subjects were chosen in a random order to see a private signal and make a public prediction about which urn was used. Once each subject made a prediction, a monitor announced which urn was actually used. Everyone who predicted correctly earned $2; others earned nothing. This process was repeated fifteen times for each group of six subjects with a new die throw to select the urn at the beginning of each repetition. New subjects were recruited for six different sessions of the experiment using the parameters described above.[4]

Sample results from one of these sessions are presented in figure 1. In period 1, the first two subjects in the sequence made conflicting predictions, based on their own private signals. Reading across the first row of the figure, the next three subjects saw "a" signals and predicted A, and the final subject saw a "b" signal and followed the others with an A prediction. This was an incorrect cascade since urn B was actually being used for draws, as shown in the far right column of the figure. Decisions made in an incorrect cascade are colored blue in the figure. The second period begins with an error, as indicated by the red shading of the incorrect A prediction that followed a "b" signal. Cascades, indicated by green shading, formed in many of the subsequent periods. The longest cascade was the incorrect cascade in period 9, shown in blue. This session is atypical in the sense that cascades were possible in most of the periods. Over all sessions, the sequence of draws made cascades possible in about sixty percent of the periods, and cascades actually formed in seventy percent of the periods in which they were possible.

Despite the overall consistency of the data with predicted behavior, mistakes are not uncommon. The first subject to make a decision in period 2 saw a "b" signal and predicted A. This type of error, which is inconsistent with Bayes' rule and private information, is probably the result of confusion or carelessness. Another type of error was committed by the third decision maker in period 8. This person saw a "b" signal and made a B prediction, consistent with the private information but inconsistent with the Bayesian posterior determined by the two previous A decisions. Perhaps it is not surprising that this particular subject relied on private information, since this person, who was the final decision maker in period 1, went against private information to follow the others in a series of predictions that turned out to be incorrect. Overall, this type of error occurred in about one-fourth of the cases where the optimal Bayesian decision differed from the decision implied by private information.

Anderson (1994) estimates error rates in these experiments using a logistic error model, which explains a number of interesting patterns in the data. For example, when there is some chance that the first person in the sequence will make an error (a prediction that is inconsistent with the primally observed signal), the second person should rely on private information when this information is inconsistent with the first decision. In this case, the logistic error model predicts that the second person will follow private information with a probability of .96. In fact, this reliance on private information occurred in ninety-five percent of the cases where the second person's signal differed from the first person's prediction.

Another prediction of the logistic error model is that increases in the payoff associated with a correct prediction will reduce the incidence of errors. Anderson (1996) replicates the basic cascade design with three different payoffs for a correct decision: $0 (no payoff), $2 (payoff), and $4 (double payoff).[5] Increasing payoffs from $0 to $2 resulted in a decrease in the number of errors, but the increase from $2 to $4 had no significant effect on errors.

Cascades can arise in other contexts, where decisions, but not private signals, are observed by other decision makers. In particular, Allsop and Hey (1997) report the results of an experiment based on the Banerjee (1992) model in which only one of a finite number of assets will have a positive payoff. Each subject receives a signal with probability α. A signal will reveal the correct asset with probability β. If two or more people have selected a particular asset that is different from the one indicated by a person's private signal, then it is optimal for that person to choose the most commonly selected asset, independent of the values of α and β. Cascades can form in this setup, as it is optimal for people to ignore their private information and follow others. Allsop and Hey report that the incidence of observed cascades is lower in their experiments than would be predicted by the Banerjee model.[6] Moreover, subjects' behavior is affected by the α and β parameters, despite the fact that theoretical predictions are independent of these parameters. The most common deviation from predicted behavior was the tendency for individuals to select the asset indicated by their own signal, even when it is irrational to do so. All of the analysis is based on the assumption that others do not make mistakes, and we conjecture that the anomalous behavior patterns may be explained by incorporating the possibility of decision error. When others may make errors, the option of following one's own information becomes more attractive. Moreover, the α and β parameters affect the relative costs of not following the herd, and therefore, these parameters affect behavior in a theoretical model with decision error.

II. MARKET APPLICATIONS

Much of the cascade literature is motivated by an interest in financial market applications, where following the crowd may result in price bubbles or crashes. The connection between investment decisions and inferences drawn from other's investment decisions was noted by Keynes (1965), who compared investment decisions with a guessing game in which participants have to predict which contestant will receive the most votes in a beauty contest. This process is complicated if each person tries to think about what the others will find attractive, and what the others will think about what others find attractive, etc. Similarly, investors in financial markets will try to guess which stocks will become popular, even in the short term, perhaps by looking at others' purchases as revealed sequentially on a ticker tape. Note that the sequence of investment decisions usually is not exogenously specified, as was the case for the cascade experiments considered in the previous section. Here we review several papers in which the order of decisions is determined by the subjects in experiments, based on their own information.

Bounmy, Verganud, Willinger, and Ziegelmeyer (1997) conducted an experiment in which paired subjects each received a signal that pertained to the value of an asset. Moreover, the quality of the signal was apparent when the signal was received. The signal was either positive, indicating that it is better to buy, or negative, indicating that it is better to sell. (Uninformative, zero, signals were also possible.) The magnitude of the signal indicated its quality, e.g. a large positive signal indicated that it is more likely that the correct decision is to buy. At each decision point, subjects could buy, sell, or incur a small cost by waiting. The prediction of the theoretical model is that subjects with less informative signals should wait and then imitate the decision made by the other person if that person decides earlier. These predictions tended to describe observed behavior.

Camerer and Weigelt (1991) report that imitation of earlier decisions may occur even if initial trades are not based on superior information. For example, randomness in initial decisions by uninformed traders may create a price movement that seems to indicate a conformity of inside information. Then other traders may imitate these decisions, often resulting in an incorrect cascade.

Plott, Wit, and Yang (1997) present results from experiments that implement a type of parimutuel betting. The setup is analogous to a horse race where a cash prize is divided among those who bet on the winning horse in proportion to the amounts that they bet. In the experiment, there were six assets, and only the asset that corresponds to the true state had value to investors. Participants received private and imperfect information about the true state, and then decided how to allocate their endowment between purchases of each asset. Purchases were revealed as they occurred, so individuals could see others' purchases and make inferences about others' information. Information aggregation occurred to a large extent. In most cases, the asset corresponding to the true state was most heavily purchased. In some cases, however, heavy purchases of an asset that did not correspond to the true state induced others to imitate, which created a herding pattern, indicating an incorrect cascade.

This literature, which builds on simplified models of inference in sequential decision making, seems to be progressing toward more interesting applications, like the parimutuel betting example. Even though these applications are motivated by naturally occurring institutions, the usefulness of field data is limited by the fact that the private information of traders and/or betters typically cannot be observed. Laboratory experiments are particularly useful in examining herding behavior because private information is observed by the experimenter and the flow is information can be precisely controlled.

REFERENCES

Allsopp, Louise and John D. Hey "An Experiment to Test a Model of Herd Behavior," unpublished working paper, University of York, 1997.

Anderson, Lisa R. Information Cascades, unpublished doctoral dissertation, University of Virginia, 1994.

Anderson Lisa R. "Payoff Effects in Information Cascade Experiments," working paper, 1996.

Anderson, Lisa R. and Charles A. Holt "Information Cascades in the Laboratory," American Economic Review, December, 1997.

Anderson, Lisa R. and Charles A. Holt "Classroom Games: Information Cascades," Journal of Economic Perspectives, Fall 1996, 10(4), pp. 187-193.

Bannerjee, A. V. "A Simple Model of Herd Behavior," Quarterly Journal of Economics, August 1992, 107(3), pp. 797-817.

Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," Journal of Political Economy, October 1992, 100 (5), pp. 992-1026.

Bounmy, Kene, Jean-Christophe Vergnaud, Marc Willinger, Anthony Ziegelmeyer "Information Externalities and Mimetic Behavior with Endogeneous Timing of Decisions: Theoretical Predictions and Experimental Evidence," working paper, University Louis Pasteur, Strasbourg, 1997.

Camerer, Colin F. and Keith Weigelt "Information Mirages in Experimental Asset Markets," Journal of Business, October 1991, 64(4), pp. 463-93.

Holt, Charles A. and Lisa R. Anderson "Classroom Games: Understanding Bayes' Rule," Journal of Economic Perspectives, Spring 1996, 10(2), pp. 179-187.

Keynes, John Maynard The General Theory of Employment, Interest, and Money. New York: Harcourt, Brace & World, 1965.

Plott, Charles R. and Shyam Sunder "Efficiency of Experimental Security Markets with Insider Information: An Application of Rational-Expectations Models," Journal of Political Economy, August 1982, 90, pp. 663-698.

Plott, Charles R. and Shyam Sunder "Rational Expectations and the Aggregation of Diverse Information in Laboratory Security Markets," Econometrica, 1988, 56, pp. 1085-1118.

Plott, Charles R., J. Wit, and W. C. Yang "Parimutuel Betting Markets as Information Aggregation Devices: Experimental Results," working paper, Caltech, January 1997.

Welch, Ivo. "Sequential Sales, Learning, and Cascades," Journal of Finance, June 1992, 47 (2), pp. 695-732.

Endnotes

*. Lisa R. Anderson is Assistant Professor of Economics at The College of William and Mary, Williamsburg, VA. Charles A. Holt is Professor of Economics at the University of Virginia, Charlottesville, VA. This research was supported in part by grants from the National Science Foundation (SES93-20617).

1. The argument given in the text is based on the assumption that, when the posterior is 1/2, the person bases the prediction on private information. This assumption is strongly supported by the laboratory evidence in Anderson and Holt (1997), who also note that relying on private information in this case is rational when the initial decision makers may have made a mistake. The cascade formation described in the text is unchanged if we make the alternative assumption that the prediction is equally likely to be an A or a B when the posterior is 1/2. Then the second prediction is also informative: a B prediction reveals a "b" signal, and an A prediction reveals that an "a" signal was more likely, since the "a" signal always results in an A prediction and the "b" signal only yields an A prediction half of the time. If the first two predictions match, then the third decision maker should reason: the first A prediction reveals an "a" signal, the second A prediction reveals that an "a" signal was more likely, and the total information content of these two observations is more favorable for state A than my own signal, even if it is a "b" signal. In this manner, the optimal decision of the third person in a sequence is to follow the first two predictions when they match, regardless of the signal observed.

2. In the experiment, the balls were actually identified by a "light" or "dark" color, instead of being labeled by letters.

3. In fact, this counting method of determining the posterior probability can be generalized to provide

a natural and intuitive way of teaching Bayes' rule to students in a classroom setting. For example, suppose that the prior probability of urn A is 2/3 instead of 1/2. To reflect the fact that Urn A is twice as likely in this case, just double the number of balls listed for Urn A, keeping the proportions unchanged: Urn A: a, a, b, a, a, b Urn B: b, b, a. Now four of the five "a" balls are listed in Urn A, so the draw of an "a" ball results in a posterior of 4/5 for Urn A, as can be verified by Bayes' rule. Holt and Anderson (1996) show how this ball counting heuristic is related to the algebra of Bayes' rule, and how the relationship can be used in the teaching of Bayes' rule.

4. Anderson and Holt (1997) report six additional sessions using an asymmetric design in which urns A and B contained different proportions of "a" and "b" balls.

5. In all three designs, subjects were paid $5 for participation in the experiment. In addition, subjects in the no payoff treatment were paid a fixed amount, $20, independent of their decisions.

6. Banerjee's result is based on a number of tie-breaking assumptions that are not listed here.