

## Classroom Games

# Information Cascades

Lisa R. Anderson and Charles A. Holt\*

**Abstract:** This paper describes how to set up a classroom exercise in which students see private signals and make public decisions in sequence. A pattern of conforming decisions in this context is called an information cascade. Once a cascade starts, it is rational for students to ignore their private signals and follow the pattern of previous decisions. This exercise provides an interactive framework that facilitates discussions of Bayes' rule, information, and social conformity.

Keywords: information cascades, experimental economics, classroom experiments, Bayes' rule.

## Introduction

Although economic forecasting is typically discussed in the context of a single individual, an interactive exercise is potentially more interesting for students. This exercise is motivated by situations in which people must combine private information with information inferred from the actions of others. For example, an employer might turn down a promising job applicant simply because the applicant is known to have been rejected for several similar jobs in the past. Here, the negative information implied by previous rejections can outweigh the employer's own positive assessment (Stern, 1990). Similarly, "bandwagon effects" in consumer purchases, adoption of production technologies, and politics may be due to inferences about others' private information that people make by observing others' decisions.

This paper describes a sequential prediction game in which students (like the employers in the previous paragraph) receive private sample information (like a job interview), and make

---

\* Support for this research was provided in part by the National Science Foundation (SES93-20617). We wish to thank Sheryl Ball and the editors for helpful suggestions.

predictions about an unknown event (the applicant's ability). The predictions are public and sequential, so students later in the sequence may use observed decisions to make inferences about the private information of students earlier in the sequence. A pattern of conformity can arise if initial predictions coincide and the inferred information dominates the private information of subsequent decision makers. The followers go along with a consensus prediction, even if it would not be the "correct" prediction made only on the basis of their own sample. This is known as an "information cascade" (Bikhchandani, Hirshleifer, and Welch, 1992).

This exercise can be used effectively in conjunction with the material on information, uncertainty, and game theory in introductory or intermediate microeconomics classes. It can also be used in courses that specifically cover Bayes' rule (economic statistics, finance, and managerial economics), or more broadly, in courses that deal with social conformity (political science, psychology, and sociology). Even in an introductory class, there is often a lively discussion about the extent to which consumers and producers are as rational as implied by standard textbook models. We use cascade exercises in principles classes as an example of sophisticated decision making that might initially appear to be irrational. We do not mean to imply that the instructor should dismiss the possibility of an irrational bias towards following the status quo. However, the discussion can be directed toward distinguishing alternative explanations and assessing the importance of the costs of irrationality in specific contexts.

## **Procedures**

The exercise will take from 30 to 45 minutes, depending on the size of the class and the level of the discussion. You will need a deck of cards, three identical marbles of one color, three identical marbles of another color, and three opaque plastic cups, e.g. the ones used for large drinks at college athletic events. Since marbles can vary a lot in size and color, the instructor may prefer to use colored golf balls, ping pong balls, or bath oil beads. For purposes of discussion, suppose that three of the marbles are Dark, and three are Light. Label one cup "urn A" and the other cup "urn B". Urn A will contain two Dark marbles and one Light marble, and urn B will contain one Dark and two Lights. Urn A could correspond to a high-ability worker, and a Dark marble could represent a favorable interview, but it is better to defer specific interpretations until after the decisions have been made.

Students are shown the contents of each urn and are told that each urn is equally likely *ex ante*. It helps to write this information on a chalkboard or transparency. A coin is flipped to select one of the urns; a heads determines urn A, and a tails determines urn B. The contents of the selected urn are emptied into the third, unlabeled cup. Of course, a screen (podium or portable black board) must be used to keep students from observing which urn is being used. The students are approached one at a time in sequence to observe a private draw (with replacement) from the unlabeled cup. Each student makes a public prediction of the urn used (A or B), and subsequent students can use the information contained in this prediction. After the first few draws, students can be reminded that the ball has been returned to the cup, which should be shaken well each time.

Since students make their decisions in sequence, you need a way to determine the order of decisions. A random order is best in a small class when the same six to ten students will be used repeatedly. Playing cards work nicely for this, as each student can be assigned a card. Shuffle the cards and draw them without replacement to determine the random order of individual decisions in each sequence. In a larger class, you can follow the seating pattern, using a different row of six-to-ten students for each new decision sequence.

After each card is drawn, take the student selected to a private location, e.g. behind a screen, and draw a marble from the unmarked cup. Ask this student to predict which urn (A or B) was used, and then announce the prediction (but not the draw) to the class. This is the interactive element of the experiment, since students often gain information about which urn was used by observing the previous decisions of others. To prevent inferences drawn from a student's indecision or hesitation, you should have the student face away from the class, wait 10 seconds, and write the prediction, A or B, which is then announced by recording it on a blackboard or overhead transparency. In particular, the student should not announce their own prediction to the class. This process is continued until all students have seen a private draw from the same cup and have made a public prediction.<sup>1</sup> Then the urn that was actually used is announced to all, so

---

<sup>1</sup> An alternative that we have used with a medium-sized class is to put six-to-ten players into the hall, each with a playing card number. The person who is selected by the card draw is called into the room and observes a draw from the cup, as does the rest of the class in this setup. Then the person's prediction, but not the draw, is written on the chalkboard and announced to those in the hall, before the next person in the sequence is called. It is probably best to have an additional student serve as a hall monitor to keep the others from talking while they are waiting.

students find out whether or not their predictions were correct. Small rewards (perhaps 10 cents or a ball point pen) for correct predictions are not necessary, but they do serve to increase interest and provide more incentive to think carefully.<sup>2</sup>

For classes with less than fifty students, it is nice to have each student make a decision at least once. This may require five sequences with up to ten decision makers in each one. Time should not be a problem, since such sequences will take about five minutes once one or two of them are completed. Even with a small class, go through several sequences with no fewer than six people in a sequence. If time is limited, you may stop after you finish a sequence with a cascade. It is better to defer discussion until after all sequences have been completed, in order to avoid biasing behavior. If decisions are recorded on a chalkboard or transparency as they are made, it will be quicker and easier to refer to specific patterns and events.

### **Information Cascades**

Consider an example using the parameters discussed above, which can be presented on the blackboard as follows:

Urn A: Dark, Dark, Light

Urn B: Dark, Light, Light

If the first decision maker predicts A, the other students can infer that this person saw a Dark marble, since the urns were equally likely to begin with and the sample favors urn A. Note that both urns are equally likely, and hence, before the die is tossed, all six marbles have the same chance of being drawn. Since two of the three Dark marbles are in urn A, the posterior probability of A after seeing a Dark is  $2/3$ . This calculation can be verified with Bayes' rule:

---

<sup>2</sup> Some instructors translate experiment earnings into extra credit points (Williams and Walker, 1993), but this requires parameter designs that meet the test of fairness. In the information cascade experiment, the first two people in the sequence have less information, and in this sense are disadvantaged. This problem can be resolved by repeating the sequential prediction exercise a number of times, with randomly determined orders. We prefer not to use extra points, since this can make the situation rivalistic, where students want to maximize relative payoffs, not absolute payoffs.

$$\begin{aligned}
Pr(A | Dark) &= \frac{(1/2)Pr(Dark | A)}{(1/2)Pr(Dark | A) + (1/2)Pr(Dark | B)} \\
&= \frac{(1/2)(2/3)}{(1/2)(2/3) + (1/2)(1/3)} = 2/3.
\end{aligned}$$

In the previous installment of this column, Anderson and Holt (1996) describe how Bayes' rule can be explained in this context and how the (Bayesian) ball counting heuristic can be generalized for the case of unequal prior probabilities.

If the first person predicts A and the second person sees a Dark, the second person has a sample of two Darks (one inferred and one observed). This person should clearly predict A since the probability that urn A is being used are even higher with two Dark draws.<sup>3</sup> Conversely, if the first person predicts urn A and the second person sees a Light draw, the posterior probability is 1/2, since the urns were equally likely before the coin flip, and the sample of a Dark (inferred) and a Light (observed) does not favor either urn. In such cases, when the second student's private draw differs from the inferred draw of the first student, the second student almost always makes the prediction implied by his or her private draw (Anderson and Holt, 1995).<sup>4</sup> If students follow their private information with a posterior of 1/2, others can infer a Light draw if the second person predicts B and a Dark draw if the second person predicts A. This reasoning implies that the second person's prediction reveals that person's draw, regardless of whether the initial prediction was A or B.

The key insight here is that when the first two people predict A, the two inferred Dark draws outweigh the private information seen by the third person. In particular, when the first two people predict A, the inference of two Dark draws outweighs a Light draw seen by the third

---

<sup>3</sup> After the first Dark draw, the prior for urn A is 2/3. If the second decision maker sees a Dark draw, Bayes' rule can be used to calculate the posterior for A, which is 4/5. This calculation is consistent with the counting heuristic described in the text above: Since urn A is twice as likely as urn B after the first Dark draw, just mentally double the number of marbles in urn A to be 4 Darks and 2 Lights, so that 4 of the 5 Darks are in urn A.

<sup>4</sup> Relying on one's private information when the theoretical posterior is 1/2 can be rational if there is a slight chance that previous decision makers have made one or more errors. Errors will be discussed in more detail below.

person.<sup>5</sup> Hence the third (and subsequent) students should go along with the first two predictions, even if their private draws indicate otherwise. The conformity of subsequent predictions, therefore, reveals no new information. This is known as an "information cascade." These cascades will be very fragile and easily broken by errors. For example, someone who deviates from an established pattern is presumably following his or her private information, which is then signaled to later decision makers. Since the cascade starts with a slight imbalance of information, a single deviation of this nature is sometimes enough to break a cascade.

In controlled laboratory experiments, Anderson and Holt (1997) report that cascades form in about 80 percent of the cases where the possibility arises. As an example, consider a pattern that emerged once when urn B was used: A|Dark, B|Light, B|Light, B|Light, B|Dark, B|Dark. The first subject saw a Dark draw and predicted A. The second person saw a Light draw and predicted B. These opposite predictions balanced each other, and the third person made a prediction of B that was consistent with the observed Light draw. The fourth person saw Light and predicted B, which started a cascade that was followed by the fifth and sixth people in the sequence.

Students may be particularly interested if there is a "reverse cascade" where the first few students are unlucky and see signals that indicate the wrong urn. For example, consider a case where urn A is being used, but the first two draws from urn A were Light, even though the urn only contained one Light marble out of three. As a consequence, the first two people predicted urn B, and the third and fourth people (rationally) followed this pattern with B predictions, despite an observation of a Dark marble in each case. In this pattern, everyone made an incorrect prediction, even though each of them made a rational choice.

The combination of signals needed to make a cascade possible may not occur in all draw

---

<sup>5</sup> If you assume that students randomly predict either urn when the posterior is 1/2, then the calculation is more complicated, but yields the same result: If the first and second decision makers match predictions, all subsequent decision makers should (according to their Bayesian posteriors) follow the pattern, regardless of their private signal. This can be verified formally with Bayes' rule (Anderson and Holt, 1995), but it can be explained intuitively. Suppose the first person predicted urn A. Inferring that the first person saw a Dark, the second person always predicts urn A if a Dark is drawn and randomly chooses between A and B if a Light is drawn. Therefore, a prediction of urn A by the second person means that it is more likely that the second draw was Dark. Since the first person must have seen a Dark and the second person is more likely to have seen a Dark, the information implicit in the two matching predictions dominates any information that the third person might see, whether Dark or Light.

sequences, and mistakes can be made even if the needed combination does occur.<sup>6</sup> Therefore, the sequential prediction exercise should be repeated five to six times, with records kept on the blackboard. If someone deviated from a pattern, it is interesting to find out whether this person made a decision that was consistent with private information. The most likely response is that the deviant based the prediction on the private draw because of the uncertainty associated with inferences drawn from the decisions of others. If students believe that others are making mistakes, it may not be irrational for them to make decisions based solely on private information. In any event, this exercise forces students to make predictions and to think about how others make predictions.

Variations of the basic setup can be used to introduce more economic content. One possibility is to explain the setup as one in which a buyer must decide whether or not to purchase a product of unknown quality at a prespecified price. For example, the value of the urn can depend on the number of Dark marbles it contains, and buyers see private draws and make public purchase decisions in sequence.

### **Further Reading**

This exercise provides useful experience in drawing inferences from the decisions of others. Predictions made in an interactive, group context tend to stimulate class discussion. Examples from naturally occurring situations add to the discussion. In an article entitled "Yes, ten million people can be wrong," *The Economist* considers the possibility that the widespread use of the popular anti-depressant Prozac might be the result of an information cascade. Bikhchandani, Hirshleifer, and Welch (1992) discuss a wide range of situations in which cascades might develop; for example, experimentation with drugs, waves of hostile takeovers in financial markets, and continued rejections of a journal submission that finally receives positive referee reports. Other applications include: bank runs (Diamond and Dybvig, 1983) and investors

---

<sup>6</sup> Anderson (1995) estimates a logistic model of decision errors with data from cascade experiments with financially motivated subjects.

subscriptions to initial public stock offerings (Welch, 1992).<sup>7</sup> For related theoretical models of cascades and herding behavior, see Bannerjee (1992), McKelvey and Ordeshook (1985), and Bikhchandani, Hirshleifer, and Welch (1992). Anderson and Holt (1997) report results of information cascade experiments in controlled settings with financially motivated subjects. Cascades are common in the laboratory, but they are fragile in the sense that they can be broken by public releases of information, as predicted by economic theory.

---

<sup>7</sup> There can be other explanations of these phenomena, of course. It may be rational to try to withdraw cash during a bank run, even if nobody has any private information. Similarly, there can be real value in conformity to a technological standard that others have already adopted.



## References

Anderson, Lisa R., "Information Cascades: A Logistic Error Model of Laboratory Data," working paper, University of Virginia, revised March 1995.

Anderson, Lisa R., and Charles A. Holt, "Information Cascades in the Laboratory," *American Economic Review*, December 1997, 847-862.

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

Bannerjee, A. V., "A Simple Model of Herd Behavior," *Quarterly Journal of Economics*, August 1992, 107(3), 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), 992-1026.

Diamond, Douglas W., and Philip H. Dybvig, "Bank Runs, Deposit Insurance, and Liquidity," *Journal of Political Economy*, June 1983, 91, 401-419.

McKelvey, Richard D., and Peter C. Ordeshook, "Elections with Limited Information: A Fulfilled Expectations Model Using Contemporaneous Poll and Endorsement Data as Information Sources," *Journal of Economic Theory*, June 1985, 36, 55-85.

Stern, Steve, "The Effects of Firm Optimizing Behavior in Firm Matching Models," *Review of Economic Studies*, October 1990, 57, 647-660.

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

Williams, Arlington W., and James M. Walker, "Computerized Laboratory Exercises for Microeconomics Education: Three Applications Motivated by the Methodology of Experimental Economics," *Journal of Economic Education*, Fall 1993, 22, 291-315.

"Yes, Ten Million People Can Be Wrong," *The Economist*, February 18, 1994, vol. 330 no. 7851, 81.