Informational cascades in the laboratory: Do they occur for the right reasons?*

Steffen Huck and Jörg Oechssler
Humboldt–University[†]

January 23, 1999

Abstract

Recently, the theory of informational cascades has been tested in an experiment by Anderson and Holt (1997) who report that their data support the theory amazingly well. In this note we report on an experiment designed to find out whether observed cascades are indeed due to rational Bayesian updating. However, we find little support for rational updating. The simple heuristic "follow your own signal" does much better in explaining our data than Bayesian rationality.

1 Introduction

Informational cascades are said to occur when decision makers choose to ignore their own private information in favor of imitating others who faced the same decision earlier on. This, of course, requires that individuals decide in a sequence and that actions (but not signals) are observable. Bikhchandani, Hirshleifer, and Welch (1992) and Banerjee (1992) observed that it can be perfectly rational from an individual point of view to ignore one's private information. Cascades are the unavoidable (and sometimes unfortunate) consequence of a strict application of Bayes' rule. Recently, the theory has been tested in an experiment by Anderson and Holt (1997) who report that

*The first author acknowledges financial support from the Deutsche Forschungsgemeinschaft (DFG). We are grateful to Lisa Anderson and Charles Holt for sharing their unpublished data with us.

[†]Department of Economics, Institute for Economic Theory, Spandauer Straße 1, D–10178 Berlin, Germany, Fax +49 30 2093 5704, email *huck@...* and *oechsler@wiwi.hu-berlin.de*.

their data support the theory amazingly well. They interpret their data as showing that "Individuals generally used information efficiently and followed the decision of others *when it was rational"* (p. 859, our emphasis). Given the subtleties involved in the application of Bayes' rule this seems remarkable.

In this note we report on an experiment designed to find out whether the observed cascades are indeed due to rational Bayesian updating. To do this, we employ a design which differs from the one by Anderson and Holt in some important dimensions. First of all, in our experiment we control for subjects' beliefs. In Anderson and Holt's experiment every decision in a sequence of decisions is taken by a real person. Therefore, it is virtually impossible to know the beliefs of subjects. What does a subject believe about the decision rules of those deciding before him? This cannot be known but it should be known when one wants to analyze whether the theory works for the right reasons. In our experiment we confront subjects with a hypothetical sequence of decisions taken by hypothetical predecessors. Subjects are explicitly told that all previous decisions were taken rationally. To ensure that subjects understand the meaning of this information we used undergraduates in economics and business administration who had just finished an intermediate microeconomics class which covered the principles of Bayesian decision making. The decision task was part of the written final exam in this class. Subjects were rewarded by points for the correct, rational answer. Note that the task is a decision problem and thus there is no ambiguity what the correct answer should be. Hence this design made it possible to control for subjects' beliefs and objectives. Another important difference in the design was that we asked subjects to explain their decision.

The data of our experiment questions the interpretation that subjects decide rationally. Our subjects were in agreement with Bayes' rule only about half of the time. Hardly any subjects were able to explain correctly how to apply Bayes rule, which indicates that subjects who were in agreement with Bayes' rule, were so accidentally. The rule that describes subjects' behavior best is "follow your own signal". As a consequence, fewer cascades happened than predicted by theory. This led us to a closer analysis of the Anderson and Holt data. In their experiment many decisions are rather simple — simple meaning that subjects can obtain the correct decision prescribed by Bayes' rule by applying various other heuristics, such as "follow the majority" or "follow your own signal". Focussing on more difficult decisions in their experiment we find that the theory does not very well in explaining this subset of their data and that, in fact, the instances in which it does well are mainly cases in which Bayes' rule and the simple rule, "follow your own signal", coincide.

The remainder of this note is organized as follows. Section 2 describes our experimental setup and offers a brief theoretical analysis. Section 3 presents the experimental results, Section 4 deals with the data of Anderson and Holt, and Section 5 concludes.

2 Experimental setup

In total 63 subjects participated in our experiment. Subjects were divided into three equal sized groups, each group having to work on a different task. Each subject was confronted with one task, which was embedded into the written final exam in a regular intermediate microeconomics class. The payoff for taking the correct, expected utility maximizing decision was 4 points in the exam which amounted to about 5% of the final grade in the class.[1] All subjects were familiar with the basic notion of Bayesian updating though they were not specifically trained in the logic underlying informational cascades. The frame of the task was neutral (see Appendix A for the instructions) and there was no directly imposed time limit for this particular question.[2] In a second question subjects were asked to explain their decision. This explanation could earn subjects another 4 points in the exam.

All 3 tasks were of the following nature. Each subject was given a series of hypothetical decisions of his hypothetical predecessors. He or she was told the quality of the signals (but not the signal itself) of all predecessors and the predecessors' decisions. Furthermore, he or she was told that all predecessors had observed the decisions of their predecessors and had decided rationally given their information.

Let A and B denote the possible decisions, let a and b be possible signals, and α and β possible states of the world. Payoffs were specified as follows. Action A gave a payoff of 1 if and only if the true state of the world was α, and action B yielded 1 if and only if the true state of the world was β. Otherwise the payoff was 0. Thus, an expected utility maximizer would choose the ex post more likely action regardless of risk aversion.

[1]For experimental incentives like this (which can be very high given the possible monetary costs of not passing) compare, for example, Abbink, Buchta, Sadrieh, and Selten (1997).

[2]The entire exam lasted 120 minutes.

Table 1: The three decision tasks

	Previous decisions	p	Signal	q	Rational action
Task 1	B	0.9	a	0.9	B
Task 2	AA	0.6	b	0.65	A
Task 3	$BBBB$	0.6	a	0.65	A

The *a priori* probability that the true state of the world was α was announced to be $r = 0.49$. Let p be the signal's quality for all predecessors and q the quality of one's own signal. Signal quality is specified as the probability with which the correct signal occurs given the true state of the world, i.e. as $prob(a|\,\alpha) = prob(b|\,\beta)$.

Table 1 summarizes the three decision tasks. The first column shows the sequence of previous decisions. All previous decisions were based on signals of quality p, which is given in the second column. The third column contains the subject's own signal and the fourth this signal's quality. Finally, the fifth column contains the action a Bayesian player would rationally choose.

The derivation of the rational action for Task 1 is straightforward. Given that the first agent chose B he must have received signal b. Since p and q are equally informative, the signal of the first and the second agent cancel out. This leaves the second agent with an ex post probability for α which equals the ex ante probability of 0.49. Hence the rational action is to choose B.

In Task 2 the derivation of the rational action is a bit more subtle. The first agent rationally follows his own signal, that is a. The second agent also must have received a because had he received b, he would have chosen B by a similar argument as above. Thus the third agent is facing two a–signals with quality 0.6 and his own b–signal with quality 0.65. Thus the ex post probability for A being the optimal action is given by Bayes rule as

$$\begin{aligned} prob(\alpha|\,aab) &= \frac{prob(aab|\,\alpha)prob(\alpha)}{prob(aab|\,\alpha)prob(\alpha) + prob(aab|\,\beta)prob(\beta)} \\ &= \frac{(.6)^2(.35)(.49)}{(.6)^2(.35)(.49) + (.4)^2(.65)(.51)} = .5379 > \frac{1}{2}. \end{aligned}$$

Therefore, the rational action is A.

In Task 3 one has to realize that already the second agent is rationally disregarding his own signal (as in Task 1). That is, the decisions of agents

Table 2: Number of correct and incorrect decisions

	Correct	Incorrect
Task 1	6	15
Task 2	7	14
Task 3	16	5

2 through 4 are irrelevant since they are already part of a cascade. The fifth agent therefore has only 2 signals to go on. His own a–signal and the first agent's b–signal. Given that his signal is of better quality a calculation using Baye's rule shows that the rational action is A.

Note that in the first two tasks it is rational to disregard one's own signal, thus joining or starting a cascade. In the third task it is rational to follow one's own signal, thus breaking a cascade.

3 Experimental results

The experimental results are summarized in Table 2 according to whether they agree with rational decision making or not. In Tasks 1 and 2 the majority of subjects chose the wrong decision, while in Task 3 it was the other way round. In all, only 53% of the decisions are in line with Bayesian updating. This is not significantly different from random decisions based on flipping coins. Note further that cascades happen in only 31.0% of cases in which they should happen (namely in Tasks 1 and 2).

The question immediately arises why do subjects do so much better in Task 3 than in Tasks 1 and 2? In which way is Task 3 different from the others? On first sight it seems that Task 3 is even more difficult than the others because to deduce the optimal decisions one has to realize first that the previous sequence of B–decisions is already a cascade. Thus, at least Tasks 1 seems easier. However, in Task 3 rationality leads to a decision which is in line with one's own signal which it is not in Tasks 1 and 2. In other words, in contrast to Tasks 1 and 2, the simple rule "follow your own signal" coincides with Bayesian rationality in Task 3.

In fact, the simple rule "follow your own signal" overall does much better in explaining the data. In 71.4% the decisions agree with this rule which is significantly different from random coin flipping. In fact, given our data no other rule can outperform the rule "follow your own signal".

Even more devastating for rational decision making is the analysis of the second exam question, in which students were asked to explain their decision. Not a single student described the correct Bayesian reasoning in Tasks 1 or 2. The argument in Task 3 was performed correctly by only 3 out of 21 students. In all tasks about a quarter of students tried to apply Bayes' rule but did so incorrectly. About 40% of the students either exclusively considered their own signal or gave their signal excessive weight. The remaining students did not give any explanation.[3]

We interpret our data as giving fairly strong evidence against rational cascade theory. But how does this square with the data of Anderson and Holt, who found support for Bayesian updating. Did their subjects do a better job in applying Bayes' rule? As we will see in the next section this is not necessarily the case. Rather, a decomposition of Anderson and Holt's data suggests that cascades occur in a lab only in very simple settings.

4 Anderson and Holt's data

Considering Anderson and Holt's data in detail it is difficult not to be amazed by the diligence with which University of Virginia undergraduates applied Bayes' rule. However, most of the decisions in their experiment were simpler than the tasks we chose for our experiment. For comparison we focus therefore on Anderson and Holt's asymmetric design (their sessions 7-12). For those sessions Bayes' rule does only slightly better than the rule "follow your own signal" (78.8% vs. 71.8% success rate).

This becomes even more pronounced when we consider cases in which Bayes' rule is inconsistent with a simple counting rule discussed by Anderson and Holt, which prescribes to choose the action which receives the most inferred and observed signals in its favor.[4] It is rather obvious that cases in which Bayesian updating implies a decision inconsistent with the counting rule are cognitively more demanding than others. Therefore, we extract all these cases and analyze the decisions subjects have taken under these non–trivial circumstances. Table 3 gives an overview. It is a crosstable in which the columns distinguish between Bayesian and non–Bayesian decisions and

[3] The second author of this paper is willing to take the blame for failing to teach the students properly about Bayes' rule in his microeconomics class. However, given that most students were able to master the materials of Varian's *Intermediate Microeconomics,* Bayes' rule seems particularly hard to understand and apply.

[4] Only decisions before the onset of a cascade are counted.

Table 3: Difficult decisions in Anderson and Holt

	Bayesian	not
Following own signal	21	31
not	18	9

the rows between decisions following one's own signal and those which do not.

First, it can be seen that only 39 of the altogether 79 more difficult decisions are in line with Bayesian updating while 52 decisions are in line with the simple rule "follow your own signal". These proportions (49.4% versus 65.8%) are nearly identical to those of our experiment. That is, when it came to applying hard–nosed rationality Anderson and Holt's subjects were not better than ours. Moreover, decomposing the number of Bayesian decisions shows that more than half of them might be artifacts as here Bayesian rationality coincides with "follow your own signal".

Finally, consider another heuristic which might also have been used by Anderson and Holt's subjects. In their asymmetric design they had two urns, one with a proportion of 1/7 of *b*–signals and one with a proportion of 2/7. A simple test for finding out which urn was chosen would be to check whether the empirical frequency of relevant (inferred plus observed) *b*–signals is closer to 1/7 than to 2/7. This test yields the same answer as Bayes rule for all possible cases in Anderson and Holt's setup and is, of course, much easier to apply.[5] If some subjects used this rule, this could account for the apparent high predictive success of Bayes' rule.

5 Conclusion

There is no doubt that informational cascades sometimes happen in real life and in the laboratory. The question is, however, whether they happen for the right reason and whether they happen as frequently as they should. The theoretical literature (Banerjee, 1992, and Bikhchandani *et al.*, 1992) derives cascades by an elegant, but cognitively demanding, application of Bayes' rule. Anderson and Holt (1997) found evidence that cascades occur in the laboratory and interpret their data as showing that subjects decided rationally.

[5]Note that in our design such a simple heuristic would not work.

In this note we described an experiment which made it possible to control for subjects' beliefs and to gain insight about their reasoning. The data revealed that hardly any subjects applied Bayes' rule correctly. As a consequence we found much fewer cascades than predicted by theory. In fact, the simple rule "follow your own signal" outperformed Bayes' rule as a predictor of subjects' behavior.

We view our data as showing that subjects, even if they are familiar with Bayes' rule, have trouble applying it correctly, at least in complex situations. In Anderson and Holt's data subjects seem to apply Bayes' rule quite well. However, their decision tasks were mostly much simpler than ours. A closer inspection of more difficult decisions in their data reveals that subjects are in agreement with Bayes' rule less than 50% of the time, which would even be compatible with the hypothesis that none of the subjects applied Bayes' rule correctly as random decisions are also correct 50% of the time.

Finally, there is the issue of whether subjects would ever learn to choose correctly. Clearly, the repetitions that are necessary to test this question were not feasible in our exam–design. In a recent paper Friedman (1998) makes the assertion that all choice anomalies would vanish if subjects were given enough time and feedback to learn. We have considerable doubt whether his assertion is relevant in our case. First of all, subjects in the experiment of Anderson and Holt (1997) had the opportunity to learn as the task was repeated 15 times and feedback was provided. But the data show no significant improvements in the application of Bayes' rule.[6] Second, while Friedman's assertion is almost tautological, the kind of feedback and the number of repetitions required to back it up seem so extreme,[7] that in our view choice anomalies remain an issue to deal with in real world situations.

References

[1] Abbink, Klaus; Buchta, Joachim; Sadrieh, Abdolkarim and Selten, Rein-

[6]We have compared the number of correct Baysian decisions in the first two and the last two rounds of Anderson and Holt's data from all sessions without public draws (Sessions 1,2, 6-12). In the first two rounds 84 of 108 decisions were correct; in the last two 89 of 108.

[7]Friedman provided some subjects with a correct (and an incorrect) theoretical explanation of a task (the famous 3–door problem) plus a comparison of the past performance of the two alternative strategies and yet, only slightly more than 50% of subjects learned to choose correctly.

hard (1998), “How to Play 3x3 Games”, Paper presented at the ESA 1998 Annual Meeting in Mannheim.

[2] Anderson, Lisa and Holt, Charles (1997), “Information Cascades in the Laboratory”, *American Economic Review,* 87, 847-862.

[3] Banerjee, Abhijit (1992), “A Simple Model of Herd Behavior”, *Quarterly Journal of Economics*, 107, 797-817.

[4] Bikhchandani, Sushil; Hirshleifer, David and Welch, Ivo (1992), “A Theory of Fads, Fashion, Custom, and Cultural Changes as Informational Cascades”, *Journal of Political Economy*, 100, 992-1026.

[5] Friedman, Dan (1998), “Monty Hall’s Three Doors: Construction and Deconstruction of a Choice Anomaly”, *American Economic Review,* 88, 933-946.

Appendix: Instructions

You have to choose between actions A and B. If you choose the correct action, your payoff is 1, otherwise 0. *A priori* there is a probability of 49% that A is the correct action.

Before you decide you receive an information (either “a” or “b”) about the correct action. In **65** of 100 cases, in which a particular action is correct, this information is true. That is, if e.g. the correct action is A, then in 65% of cases you receive the information “a”.

Several people before you had to decide the same question. Person 1 had to decide first. Then person 2, and so on. Each person was able to observe the decisions of the people before them (but not their information) knowing that their information was true in **60** of 100 cases. The information of all people was independent of each other. Suppose that all people before you have decided rationally. You decide as the 5th person. The 4 persons before you have decide as follows: B, B, B, B.

a) What should you do? A □ □ B

b) Please explain your decision.