An Experimental Test of Advice and Social Learning*

Boğaçhan Çelen† Shachar Kariv‡ Andrew Schotter§

Columbia University UC Berkeley New York University

January 26, 2007

Abstract

Social learning is the process of individuals learning by observing the actions of others. In the real world, however, although people learn by observing the actions of others, they also learn from advice. This paper introduces advice giving into a standard social-learning problem. The experiment is designed so that both pieces of information – actions and advice – are equally informative (in fact, identical) in equilibrium. Despite the informational equivalence of advice and actions, in the laboratory, subjects are more willing to follow the advice given to them by their predecessors than to copy their actions. In addition, when advice is given subject behavior is more consistent with the prediction of the theory. Consequently, advice is both more informative and welfare improving.

JEL Classification Numbers: C91,C92, D8.

Keywords: Advice, social learning, experiment.

*This research was supported by the Center for Experimental Social Sciences (C.E.S.S.) and the C. V. Starr Center for Applied Economics at New York University, and the Experimental Social Science Laboratory (X-Lab) at the University of California, Berkeley. We acknowledge helpful discussions of Colin Camerer, Gary Charness, Dan Friedman, Jacob Goeree, Teck Ho, Charles Holt, Tom Palfrey, Barry Sopher and Georg Weizsäcker. This paper has also benefited from suggestions by the participants of the 2003 International ESA meeting in Pittsburgh, SITE 2004 Summer Workshop and seminars in several universities. An earlier version of the paper was titled "The Advice Puzzle: An Experimental Study of Social Learning Where Words Speak Louder than Actions."

†Columbia Business School, 3022 Broadway, 602 Uris Hall, New York, NY 10027. E-mail: bc2132@columbia.edu, url: http://celen.gsb.columbia.edu/.

‡Department of Economics, University of California, Berkeley, 549 Evans Hall # 3880, Berkeley, CA 94720. E-mail: kariv@berkeley.edu, url: http://ist-socrates.berkeley.edu/~kariv/.

§Department of Economics, New York University, 269 Mercer Street, 7th Floor, New York, NY 10003. E-mail: as7@nyu.edu, url: http://homepages.nyu.edu/~as7/.

1 Introduction

Individuals typically know only a small fraction of the information distributed throughout society as a whole. Consequently, they have a very strong incentive to try to benefit from the knowledge of others before making a decision. In social settings, when the information needed is not available from public sources and individuals can observe one another's actions, it is rational for individuals to try to learn from one another. This process is called social learning. The literature on social learning contains numerous examples of social phenomena that can be explained in this way. In particular, it argues that the striking uniformity of social behavior is one implication of social learning. At the same time, the standard social-learning model has several special features that are quite restrictive.

Perhaps most importantly, one strange aspect of the social learning literature is that it is not very social. In the real world, individuals learn by observing the actions of others, but they also learn from the advice people offer them. For example, we choose restaurants not only by viewing which of them is popular but also by being advised to do so; we choose doctors not only by viewing how crowded their waiting rooms are but also by asking advice about who to see; and so on. Furthermore, individuals make their decisions in many situations by relying only on the so-called naïve advice of non-experts: friends, neighbors, co-workers, and so on. Thus, social learning tends to be far more social than we, economists, describe it. For greater realism, we introduce advice giving into the standard social-learning problem in this paper.

In our experimental design, a sequence of subjects draw private signals from a uniform distribution over $[-10, 10]$. The decision problem is predicting whether the sum of all subjects' signals is positive or negative, and choosing an appropriate action, A or B. A is the profitable action when the sum is positive, and B if the sum is negative. But rather than choosing action A or B directly (after being informed only about the decision, action or advice (or both), of the preceding subject and before observing their own private signals) subjects are asked to select a cutoff such that they will choose action A if the signal received is greater than the cutoff and action B otherwise. After a subject reports her cutoff, she informed of her private signal and her action is recorded accordingly. In the experiments containing advice, after choosing a cutoff and observing a signal, the subject is prompted to give a binary piece of advice (choose A or choose B) to her immediate successor.

We use two treatments containing advice in our experimental design: the Advice-Only experiment, in which each subject only receives her immediate predecessor's advice as to which action to take, A or B; and the Action-Plus-Advice experiment, in which each subject observes the action chosen by her immediate predecessor and receives her advice. In both treatments the subjects' payoffs is also a function of the payoffs achieved by both themselves and their successors, so all subjects have an incentive to offer sincere advice. For comparison purposes, we will present our new results along with the results of Çelen and Kariv (2005) (hereafter, ÇK), which deal with the case in which each subject can observe only her immediate predecessor's action. We thus call ÇK the Action-Only experiment. Aside from the information structure, this new experimental design is identical to the one employed in ÇK. That is, all of the experiments use the same procedures, but the information structure is different.

Most importantly, we design the experiment so that both pieces of information - actions and advice - should, in equilibrium, be equally informative. In fact, the advice offered should be identical to the action taken by a subject after her signal has been observed. Despite this informational equivalence, we find that in a laboratory subjects are far more willing to follow the advice given to them by their predecessor than to copy their actions. As a consequence, subject behavior is much more consistent with the predictions of the theory in the presence of advice since advice appears to be more persuasive and welfare improving.

A sound concern that the reader may raise about our experimental design is that the willingness to follow advice is an artifact of the belief, on the part of subjects, that advice is more informative since in the Action-Only experiment subjects state a cutoff, which then determines their action, before observing their signal, while in the experiments containing advice they advise their successor after observing their signal. In response, we conducted a Post-Signal Action-Only experiment in which subjects observe their private signal and their predecessor's action before taking an action, A or B, directly. This experiment closely resembles the social learning experimental paradigm of Anderson and Holt (1997), and is informationally equivalent to the Action-Only and Advice-Only experiments. Comparing the behavior in this experiment with that in the Action-Only and Advice-Only experiments reinforces that subjects appear to be more willing to follow the advice given to them by their predecessor than to copy their action, and the presence of advice increases subjects' welfare. This establishes that the impact of advice is not an artifact of the experimental design.

Our paper contributes to a large literature on social learning. Banerjee (1992) and Bikhchandani, Hirshleifer and Welch (1992) introduced the basic concepts, and their work was extended by Smith and Sørensen (2000). Ellison and Fudenberg (1993, 1995), and Banerjee and Fudenberg (2004), combine certain features of the social-learning and word-of-mouth learning literatures. Anderson and Holt (1997) experimentally investigate the social-learning model of Bikhchandani, Hirshleifer and Welch (1992) and replicate informational cascades in the laboratory. Hung and Plott (2001), Kübler and Weizsäcker (2003), Çelen and Kariv (2004a), and Goeree, McKelvey, Palfrey and Rogers (2005) among others extend Anderson and Holt (1997) to investigate other possible explanations for informational cascades. The paper also contributes to the large and growing body of work on the influence of naïve advice on behavior in experimental games. Schotter (2003) provides a recent review of the experimental work which clearly demonstrates that subjects tend to give good advice and to follow the advice of others to a remarkable extent.

The rest of this paper is organized as follows. Section 2 formulates our research questions. Section 3 illustrates the underlying theory. Section 4 describes the experimental design and procedures. Section 5 summarizes the results, and Section 6 concludes.

2 Research Questions

In this section we spell out two fundamental questions that we attempt to answer in the remainder of the paper. We first ask whether subjects tend to follow advice more often than actions when each is observed under identical circumstances? That is, consider two subjects, one performing our Action-Only experiment (observing the predecessor's action) and the other performing our Advice-Only experiment (receiving the predecessor's advice). If the Action-Only subject observes her predecessor taking action A while the Advice-Only subject is told to choose action A by her predecessor, is the conditional probability of choosing A greater in the Advice-Only experiment?

Question 1. Do subjects tend to follow advice more often than action when each is observed under identical circumstances?

In the Action-Plus-Advice experiment, subjects both receive advice and observe the action taken by their predecessor. A natural question is whether this changes their behavior from what it was in the Advice-Only or Action-Only experiment. In fact, the Action-Plus-Advice experiment can give us some insight into whether subjects actually value advice more than action, because in some cases subjects actually give advice that differs from the action they take. In those cases, the predecessor may be saying, "do as I say, not as I have done," and the question is which datum is more informative and why.

Question 2. Which information—advice or action—is more valued by the subjects? And, under what circumstances do subjects offer advice that differs from their action?

3 Experimental Design

Our data come from experiments we conducted at the Center for Experimental Social Science (C.E.S.S.) at NYU and at the Experimental Social Science Laboratory (X-Lab) at UC Berkeley, and an earlier experiment of ÇK which also are discussed here for comparison purposes. We will designate the new experiments as the Advice-Only, Action-Plus-Advice and Post-Signal Action-Only experiments, and the earlier experiment as the Action-Only experiment. All experiments used the same basic procedures but they differed according to the information received by subjects. We will explain these informational regimes shortly.

In each of the experiments, we have observations from 40 subjects (in one case, 48 subjects) who had no previous experience in advice or social learning experiments. Each subject participated in only one experimental session, and eight subjects were recruited for each session. The treatment was held constant throughout a given session. After subjects read the instructions (reproduced in Appendix I), they were also read aloud by an experimental administrator. The experiment lasted for about one and one-half hours. Participation fee and subsequent earnings for correct decisions were paid in private at the end of the session. Throughout the experiment, we assured anonymity and an effective isolation of subjects in order to minimize any interpersonal factors that might have caused a tendency towards uniform behavior.

Each experimental session entailed 15 independent rounds, divided into eight decision-turns. In each round, all eight subjects made decisions sequentially, in random order. A round began with the computer drawing eight numbers from a uniform distribution over $[-10, 10]$. The numbers drawn in each round were independent of each other and of the numbers in any of the other rounds. Each subject was informed only of the number corresponding to her turn to move. The value of this number was a private signal. In practice, the subjects observed their signals up to two decimal points.

In the Action-Only, Advice-Only and Action-Plus-Advice experiments, upon being called to participate and before being informed of the private signal, the subject first received some information relevant to decision-making (either the action of the previous subject, the subject's advice, or both depending on the treatment). After receiving this information, each subject was asked to select a number between -10 and 10 (a cutoff), for which the subject would take action A if the signal was above the cutoff and action B if not. Action A was profitable if and only if the sum of the eight numbers was positive. Only after submitting a decision would the computer inform the subject of the value of her private signal. Then, the computer recorded her decision as A if the signal was higher than the cutoff she selected; otherwise, the computer recorded B.

The Action-Only, Advice-Only and Action-Plus-Advice experiments use the same procedures, but the information structure is different. In the Action-Only experiment (ÇK), subjects were able to observe only the action taken by their immediate predecessor. For example, the fifth to choose was informed only what action the fourth has taken. In the Advice-Only experiment, when subjects were called upon to make their decision they were not able to observe any of the actions taken by their predecessors. Rather, they received advice from their immediate predecessor as to what the correct action to take was. In the Action-Plus-Advice experiment, subjects were not only able to receive advice from their immediate predecessor but also observed her action. In both cases, subjects gave advice after the computer recorded their action according to their cutoff and after they observed their private signal. Thus, as a benchmark, we also conducted a Post-Signal Action-Only experiment in which each subject knew her own private signal and the action of the immediate predecessor before taking an action, A or B, directly. This experiment does not contain advice.[1]

After all subjects had made their decisions, the computer informed everyone what the sum of the eight numbers actually was. Everyone whose decision determined their action as A earned $2 if the sum of the subjects' private signals was positive (or zero), and nothing otherwise. On the other hand, everyone whose decision determined their action as B earned $2 if the sum was negative, and nothing otherwise. In addition, in the experiments containing advice, everyone earned $1 if her successor took the correct action. This was paid to insure that the advice subjects give would be their best guess as to what the correct action was. Figure 1 summarizes our experimental treatments and procedures.

Figure 1: The timeline of the experiments.

Experiment	Timeline			
Action-Only	Observe predecessor's action (A or B)	Select a cutoff from $[-10, 10]$	Receive signal, action recorded	
Advice-Only	Receive predecessor's advice (A or B)	Select a cutoff from $[-10, 10]$	Receive signal, action recorded	
Action-Plus Advice	Observe predecessor's action (A or B) Receive predecessor's advice (A or B)	Select a cutoff from $[-10, 10]$	Receive signal, action recorded	Give advice (A or B) to successor
Post-Signal Action-Only	Observe predecessor's action (A or B)	Receive signal	Choose an action, (A or B)	Give advice (A or B) to successor

[1] In this setup, unlike Anderson and Holt (1997), while there are two events which, ex ante, are equally likely to occur, there is a continuous signal space. Çelen and Kariv (2004a) discuss the importance of this difference.

4 Some Theory

In this section, we discuss the theoretical implications of the model tested in the laboratory. Çelen and Kariv (2004b) provides an extensive analysis of a general version of the Action-Only case. The main goal of this section is to demonstrate that, in the Advice-Only case, it is always optimal for a decision-maker to offer advice equal to her chosen action no matter what signal she or she observes. As a result, substituting advice for actions in our experiment cannot convey more information. This implies that the environment in the Advice-Only experiment is not informationally richer than the environment in the Action-Only experiment.

4.1 Preliminaries

Suppose that the eight agents receive private signals $\theta_1, \theta_2, \cdots, \theta_8$ that are independently and uniformly distributed over the support $[-10, 10]$. Sequentially, each agent $n \in \{1, \cdots, 8\}$ has to make a binary irreversible decision $x_n \in \{A, B\}$ where action A is profitable if and only if $\sum_{i=1}^{8} \theta_i \geq 0$, and action B otherwise.

It is immediate that $\sum_{i=1}^{8} \theta_i$ defines the set of the states of the world which are partitioned into two decision-relevant events, $\sum_{i=1}^{8} \theta_i \geq 0$ and $\sum_{i=1}^{8} \theta_i < 0$. The decision problem involves incomplete and asymmetric information: agents are uncertain about the underlying decision-relevant event, $\sum_{i=1}^{8} \theta_i \geq 0$ or $\sum_{i=1}^{8} \theta_i < 0$, and the information about it is shared asymmetrically among them.

In what follows, we will first discuss the theory behind the Action-Only case that constitutes the backbone of all three experiments. Then, we will discuss the Advice-Only and Action-Plus-Advice cases in order to demonstrate their connection.

4.2 Action-Only

The Decision Problem In the Action-Only case, except for the first agent, everyone observes only her immediate predecessor's decision. In such a situation, conditional on the information available to her, agent n's optimal decision rule is

$$x_n = A \text{ if and only if } \mathbb{E}\left[\sum_{i=1}^{8} \theta_i \mid \theta_n, x_{n-1}\right] \geq 0$$

and because agents do not know any of their successors' actions,

$$x_n = A \text{ if and only if } \theta_n \geq -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid x_{n-1}\right].$$

It follows that the optimal decision takes the form of this cutoff strategy,

$$x_n = \begin{cases} A & \text{if} \quad \theta_n \geq \hat{\theta}_n, \\ B & \text{if} \quad \theta_n < \hat{\theta}_n, \end{cases} \tag{1}$$

where

$$\hat{\theta}_n(x_{n-1}) = -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid x_{n-1}\right] \tag{2}$$

is the optimal cutoff which accumulates all the information revealed to agent n from her predecessor's action. Thus, $\hat{\theta}_n$ is sufficient to characterize agent n's behavior and the sequence of cutoffs $\{\hat{\theta}_n\}$ characterizes the social behavior. That is why we take the cutoff equilibrium (an equilibrium in which all follow the cutoff strategy (1) and (2)) as the primitives of the experimental design and of our analysis.

THE CUTOFF PROCESS Clearly, the cutoff rule, $\hat{\theta}_n$, of any agent n can take two different values conditional on whether agent $(n-1)$ took action A or action B which we denote by

$$\hat{\theta}_n(x_{n-1}) = \begin{cases} \overline{\theta}_n \text{ if } x_{n-1} = A \\ \underline{\theta}_n \text{ if } x_{n-1} = B \end{cases}$$

Çelen and Kariv (2004b) show that by using symmetry, $\overline{\theta}_n = -\underline{\theta}_n$, the dynamics of the cutoff rule $\hat{\theta}_n$ are described recursively in a closed-form solution as follows:

$$\hat{\theta}_n(x_{n-1}) = \begin{cases} -5 - \frac{\hat{\theta}_{n-1}^2}{200} & \text{if} \quad x_{n-1} = A, \\ 5 + \frac{\hat{\theta}_{n-1}^2}{200} & \text{if} \quad x_{n-1} = B, \end{cases} \tag{3}$$

where $\hat{\theta}_1 = 0$.

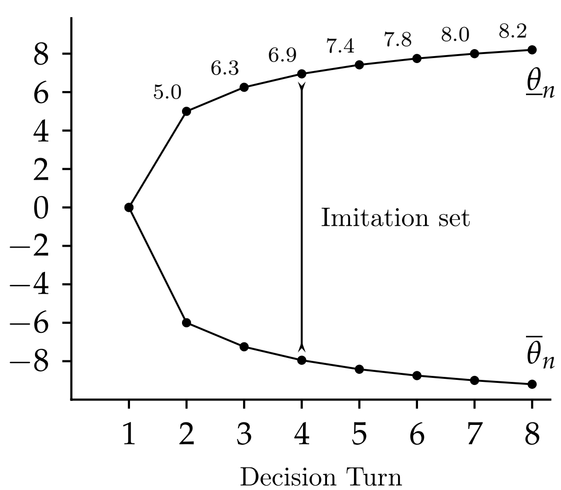

Figure 2: The process of cutoffs.

It follows immediately from (3) that the cutoff rule partitions the signal space into three subsets: $[-1, \overline{\theta}_n), [\overline{\theta}_n, \underline{\theta}_n)$ and $[\underline{\theta}_n, 1]$. For high-value signals $\theta_n \in [\underline{\theta}_n, 1]$ and symmetric low-value signals $\theta_n \in [-1, \overline{\theta}_n)$ agent n follows her private signal and takes action A or B respectively. As Figure 2 illustrates, in the intermediate subset $[\overline{\theta}_n, \underline{\theta}_n)$, which we call an imitation set, private signals are ignored in making a decision and agents imitate their immediate predecessor's action. Furthermore, since $\{\overline{\theta}_n\}$ and $\{\underline{\theta}_n\}$ are decreasing and increasing sequences respectively, imitation sets monotonically increase in n regardless of the actual history of actions. Hence, over time, agents tend to rely more on the information revealed by the predecessor's action rather than their private signal.

4.3 Advice-Only

Next, we investigate the differences between the decision problem underlying our Action-Only and Advice-Only experiments. Recall that in the games played with advice, advice is profitable if and only if the successor takes the correct action. Our purpose in this section is to demonstrate that in the Advice-Only case, advice cannot convey more information than observation of a predecessor's action. This is because in the only relevant equilibrium of the Advice-Only case, it is optimal to send advice equal to the action taken.

The Decision Problem In the Advice-Only case, everyone except for the first agent receives binary advice, denoted by $a_n \in \{A, B\}$, from her immediate predecessor. In this case, conditional on the information available to her, agent n's optimal decision rule is

$$x_n = A \text{ if and only if } \theta_n \geq -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid a_{n-1}\right].$$

It follows that the optimal decision will take the form of the cutoff strategy given by (1) where

$$\hat{\theta}_n(a_{n-1}) = -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid a_{n-1}\right] \tag{4}$$

is the optimal cutoff which includes all of the information revealed to agent n from her predecessor's advice. There are three equilibria in the the Advice-Only case: the truthful, mirror, and babbling equilibria. Here we explain and characterize these equilibria, and then demonstrate that they exhaust the set of equilibria for the Advice-Only case.

The Truthful Equilibrium When all agents believe that the advice given to them by their predecessor is identical to her action, ($a_n = x_n$, so that the beliefs are consistent) then the unique equilibrium in the Action-Only case is also an equilibrium in the Advice-Only case. We call this the truthful equilibrium. That is, with a consistent belief system, agent n's optimal cutoff $\hat{\theta}_n(a_{n-1})$ given by (1) is the same as $\hat{\theta}_n(x_{n-1})$ given by (2), and the optimal advice rule is to give advice equal to her chosen action, $a_n = x_n$. Throughout the paper, whenever we refer to the theoretical sequence of cutoffs, we consider these to be the unique equilibrium cutoffs in the Action-Only case $\hat{\theta}_n(x_{n-1})$ as given by (2), which are identical to those in the truthful equilibrium $\hat{\theta}_n(a_{n-1})$.

The Mirror and Babbling Equilibria The truthful equilibrium is not unique in the Advice-Only case, but it is easy enough to verify that there are only two other equilibria:

the mirror equilibrium and the babbling equilibrium. Here we discuss their properties and show that there are no other equilibria in the Advice-Only case.

In the mirror equilibrium, agents advise their successors to take the opposite action to theirs, $a_n \neq x_n$; successors believe that the advice given to them by their predecessor is opposite to her action; and they set their cutoffs optimally according to (4), given their beliefs (this equilibrium is the mirror image of the truthful equilibrium). In such an equilibrium, everyone who is advised by her predecessor to take action A (B) believes that the action she actually took was B (A) and thus sets her cutoff optimally at $\underline{\theta}_n$ ($\overline{\theta}_n$) instead of $\overline{\theta}_n$ ($\underline{\theta}_n$). Then, everyone advises her successor to take action A (B) if the action she herself took was B (A). Clearly, this equilibrium and the truthful equilibrium define the same process of cutoffs $\{\hat{\theta}_n\}$.

In the babbling equilibrium, agents give noisy advice, in the sense that it is uncorrelated with their action and thus independent of the available information (for example, agents randomly advise A or B); they believe that the advice given to them by their predecessor is also noisy; and they ignore advice and set their cutoffs optimally at zero, given their beliefs. Hence, in the babbling equilibrium the advice does not reveal any information to successors, no information is accumulated, and agents make decisions solely on the basis of private information simply by setting cutoffs optimally at zero.

NO OTHER EQUILIBRIA In this section we show that there are no other equilibria in the Advice-Only case. If any other equilibria existed, they would take the form of agent n advising her successor to take the same action as she did, $a_n = x_n$, with some probability $0 < p_n < 1$, and the opposite action, $a_n \neq x_n$, with probability $1 - p_n$. In the truthful equilibrium, $p_n = 1$, while in the mirror and babbling equilibria, $p_n = 0$ and $p_n = 1/2$, respectively. With a consistent belief system (agent $n+1$ believes that the advice given to her by agent n is indeed the same as the chosen action with probability p_n), it is obvious that it is optimal for agent n to always advise others to make the same decision she did, $a_n = x_n$, if $p_n > 1/2$ and the opposite action, $a_n \neq x_n$, if $p_n < 1/2$. Remember that the agent only gets a positive payoff if the successor takes the correct action.

We can prove this result by contradiction. Suppose there is an equilibrium in which the first agent sets her optimal cutoff $\hat{\theta}_1 = 0$ but advises the second agent to take the same action that she did with some probability $1/2 < p_1 < 1$ and the opposite action with probability $1 - p_1$. With a consistent belief system, the second agent conditions her decision on p_1 and on whether the advice received is $a_1 = A$ or $a_1 = B$. If the advice received is $a_1 = A$, then a simple calculation shows that $\mathbb{E}\left[\theta_1 \,|p_1, a_1 = A\right] = 10p - 5$. Thus it is optimal for the second agent to take action A if and only if $\theta_2 \geq 5 - 10p$. Likewise, if the advice received is $a_1 = B$, it is optimal for the second agent to take action A if and only if $\theta_2 \geq 10p - 5$.

Thus, after adding noisy advice to the model, the second agent's cutoff rule is

$$\hat{\theta}_2(p_1, a_1) = \begin{cases} -5 + 10p & \text{if} \quad a_1 = A, \\ 5 - 10p & \text{if} \quad a_1 = B. \end{cases}$$

Because $\overline{\theta}_2 < 0$ and $\underline{\theta}_2 > 0$ ($\overline{\theta}_2 = -\underline{\theta}_2$ as in the Action-Only case), the second agent may still follow the advice given to her, even though she would have made a contrary decision had she based her decision solely on her own signal. But then the first agent is better off by never offering advice which differs from her action. An analogous argument also applies if $0 < p_1 < 1/2$. A contradiction.

4.4 Action-Plus-Advice

We analyze Action-Plus-Advice to see whether advice has a greater impact on subjects than do actions when both are observable, not to test for finding the more sophisticated equilibria which are possible in this case. In fact, our data suggests that these equilibria are behaviorally and empirically dismissed.

The Decision Problem In the Action-Plus-Advice case, agents are not only able to receive advice from their immediate predecessor, but also to observe her action, which opens up signaling possibilities. In such a situation, conditional on the information available to her, agent n's optimal decision rule is

$$x_n = A \text{ if and only if } \mathbb{E}\left[\sum_{i=1}^{8} \theta_i \mid \theta_n, x_{n-1}, a_{n-1}\right] \geq 0.$$

Again it readily follows that the optimal decision takes the form of the cutoff strategy given by (1) where

$$\hat{\theta}_n(a_{n-1}) = -\mathbb{E}\left[\sum_{i=1}^{n-1} \theta_i \mid x_{n-1}, a_{n-1}\right] \tag{5}$$

is the optimal cutoff that accumulates all of the information revealed to agent n from her predecessor's action and advice.

Observing action and advice enables agents to engage in a more sophisticated, and hence informationally richer, strategy; they can combine all four available action-advice pairs (x_{n-1}, a_{n-1}) to partition their signal space into four regions and thus convey more information to their successors. Hence, the informational pipeline in this case is less constrained and there exists more informationally rich equilibria, which we call signalling equilibria, than in the Action-Only and Advice-Only cases.

Signaling Equilibria To illustrate what these signaling equilibria might look like, consider an equilibrium in which everyone with a cutoff leading to action A (B) advises her successor A (B) if the realization of her signal is closer to 10 (-10) than to her cutoff; she advises B (A) otherwise. Assuming consistent beliefs, such a strategy is clearly more informative than the equilibria we discussed in the Advice-Only or Action-Only cases, because agents use a finer signalling partition here to convey information about their signals. Alternatively, consider an equilibrium where all agents, except the first one, give advice equal to their predecessor's action, $a_n = x_{n-1}$. This is equivalent to the case in which agents observe the two most recent actions, in that their successor gets to see their action directly, and their predecessor's action is learned via advice.

4.5 Summary

To conclude, the observed action taken or advice given generate a binary partition of the signal space in the Action-Only and Advice-Only cases; in the Action-Plus-Advice case each element of the binary partition generated by the action can be partitioned again by the advice. So, in the Action-Plus-Advice case, the barriers to communication clog the information pipeline less tightly and agents can make better decisions. However, the truthful equilibrium in the Advice-Only case, in which agents simply advise their successor to do as they did $a_n = x_n$, is also an equilibrium in the Action-Plus-Advice case. In particular, when a convention exists such that agents ignore conflicting advice and make decisions solely on the basis of the action observed, then the resulting equilibrium is, of course, the truthful equilibrium.

5 Results

In this section we first present the results of the Action-Only, Advice-Only and Action-Plus-Advice experiments and use them to answer the two questions listed before. The Post-Signal Action-Only experiment, which does not contain advice and cutoff elicitation, provides a useful benchmark for our subsequent analysis. We will compare the behavior in the Post-Signal Action-Only experiment with that in each of the other experiments at the end of this section.

5.1 Question 1

Do subjects tend to follow advice more often than action when each is observed under identical circumstances?

In short, the answer to Question 1 is "yes." We define decisions made by subjects as concurring decisions if the sign of their cutoff agrees with the action taken (advice received). For instance, when a subject observes that her predecessor took action A (B) (or gave advice A (B)) and adopts a negative (positive) cutoff, she demonstrates concurrence by selecting a negative (positive) cutoff, she adopts a higher probability of taking action A (B). Similarly, if a subject observes action (receives advice) A (B) and selects a positive (negative) cutoff, then she disagrees with her predecessor. We say that such decisions are contrary decisions. Finally, neutral decisions are carried out by choosing a zero cutoff, which neither agrees nor disagrees with the predecessor's action (advice). Table 1 presents the percentages of concurring, contrary, and neutral decisions in the Action-Only and Advice-Only experiments.

Table 1: Agreement and contrariness in Action-Only and Advice-Only experiments.

	Concurring	Neutral	Contrary
Action-Only	44.2%	16.6%	39.2%
Advice-Only	74.1%	9.1%	16.8%

The most notable pattern in Table 1 is that advice is followed far more often than action. Over all decision-turns except the first, subjects tend to set a cutoff consistent with the advice they receive 74.1 percent of the time in the Advice-Only experiment, but only 44.2 percent of the time in the Action-Only experiment. Together with the neutral cutoffs, subjects tend to weakly agree (set a concurring or neutral cutoff) with advice 83.2 percent of the time in the Advice-Only experiment but only 60.8 percent of the time in the Action-Only experiment. These two distributions are significantly different using a Kolmogorov-Smirnov test (p-value 0.000).

The decision-level data in Table 1 potentially obscure the presence of individual effects. Thus, while Table 1 presents data on the number of decisions that were concurring, neutral, or contrary, Figure 3 shows the distribution of concurring, neutral, or contrary decisions aggregated to the subject level. The horizontal axis measures the number of disagreements (disagreed with the observed action in less than two rounds, three to five rounds, and so on) and the vertical axis measures the percentage of subjects corresponding to each interval. When subjects in the Advice-Only experiment disagreed with the advice they were given, they did so less than twice; 67.5 percent of the time. Subjects in the Action-Only experiment, tended to disagree far more often; only 20 percent of the subjects disagreed two times or less and 40 percent of the subjects disagreed six to eight times. That percentage was only 7.5 in the Advice-Only experiment. These two distributions are also significantly different using a Kolmogorov-Smirnov test (p-value 0.000).

Figure 3: The distribution of contrary subjects.

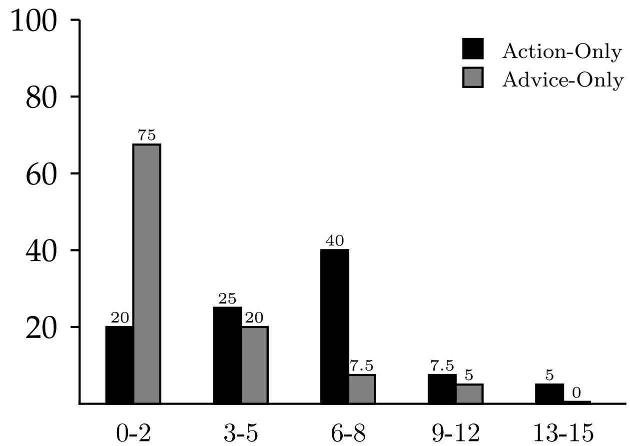

The percentage of the subjects who disagreed with the observed action/advice in less than two rounds, three to five rounds and so on.

The signs of the cutoffs as indications of agreement or disagreement tells only part of the story because they ignore the strength of the agreement or disagreement, which can be measured by the magnitude of the cutoff set. For example, if a subject observes action (receives advice) A and sets a cutoff close to -10, then not only does she agree with the action (advice) she observed, but she also does so very strongly since she will almost surely take action A. In contrast, selection of a negative cutoff that is closer to zero clearly indicates a much weaker agreement.

Because the cutoff strategy is symmetric around zero, in the sense that the strength of agreement or disagreement is independent of the actual action observed, we proceed by transforming the data generated by our subjects in the following way: take the absolute value of cutoffs in concurring decision points, and negative of the absolute value of cutoffs at contrary decision points. For instance, if a subject observes action (receives advice) A and selects a cutoff of -5, we take it as 5, since she acts in a concurring manner. On the other hand, if she places a cutoff of 5, we take it as -5, since she acts in a contrary manner. In the remainder of the paper we will refer to this as mirror image transformation.[2] Figure 4 presents the theoretical cutoffs in the truthful equilibrium and the mean cutoff (after mirror image transformation) of concurring decisions, turn by turn, for our Advice-Only and Action-Only experiments.

It is evident from Figure 4 that there is little difference in the magnitude of the cutoffs set by subjects when they strictly agreed with either the advice offered or the action observed by

[2] To make sure that there was no bias towards any of the actions A or B, we ran $y = \alpha D_A + \beta D_B + \varepsilon$, where y is the vector of reported cutoffs and D_x is the dummy variable which takes value of 1 when the action observed (advice received) is x. We fail to reject the hypothesis $\alpha + \beta = 0$.

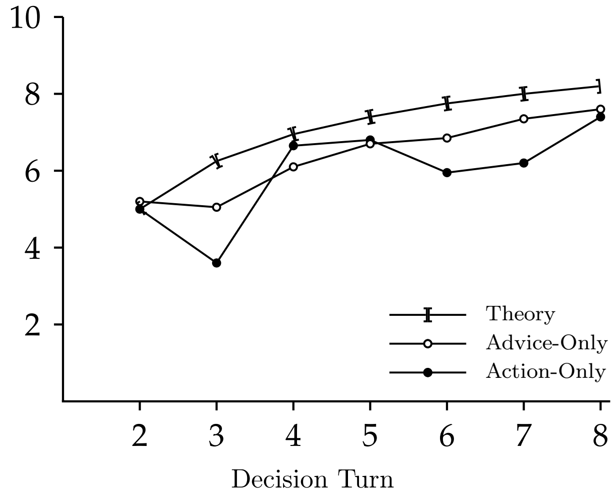

Figure 4: Mean cutoffs in concurring decisions.

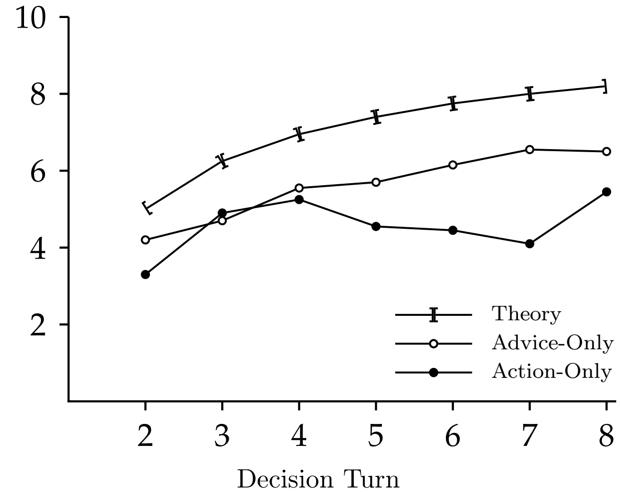

Figure 5: Mean cutoffs in weakly concurring decisions.

their predecessor. In other words, once a subject has decided to follow or imitate the advice offered or action taken, she does so with equal intensity.[3] Note too that there is a substantial degree of conformity with the theory in the magnitude of the cutoffs chosen by subjects when they agree with the action observed (advice received). However, Figure 5 shows that the situation is reversed in the Action-Only experiment, particularly in late decision-turns, when we include neutral decisions in our sample.[4]

So far, we have focused on concurring decisions. However, there is a complementary subset of contrary decisions. Once a subject decides not to follow her predecessor's action (advice), the intensity of her disagreement can be measured in several ways. Figure 6 presents the intensity of disagreement in two ways. First, we use the absolute value of the distance between the cutoff actually chosen and the one which would be selected if the subject acted according to the theoretical cutoff rule given by (2). Second, we use the absolute value of the distance of the chosen cutoff from zero. The strength of disagreement is rather severe, because when subjects disagree with their predecessor, they tend to do so in quite an extreme way.[5]

All of the results presented above condition our data on whether decisions are concurring or contrary. Figure 7 shows that if we do not condition the data on agreement and disagree-

[3]A set of two-sample Wilcoxon rank-sum (Mann-Whitney) tests run decision turn by decision turn detect no significant difference between the strongly agreeing cutoffs set in the Advice-Only and Action-Only experiments for any decision turn.

[4]A set of Wilcoxon tests run decision turn by turn detects a significant difference between the weakly agreeing cutoffs set in the Advice-Only and Action-Only experiments only in the sixth and seventh decision-turns (p -value 0.009 and 0.002, respectively).

[5]Wilcoxon tests detect a significant difference only in the fourth and eighth decision-turns (p-value 0.095 and 0.046, respectively).

Figure 6: The strength of disagreement.

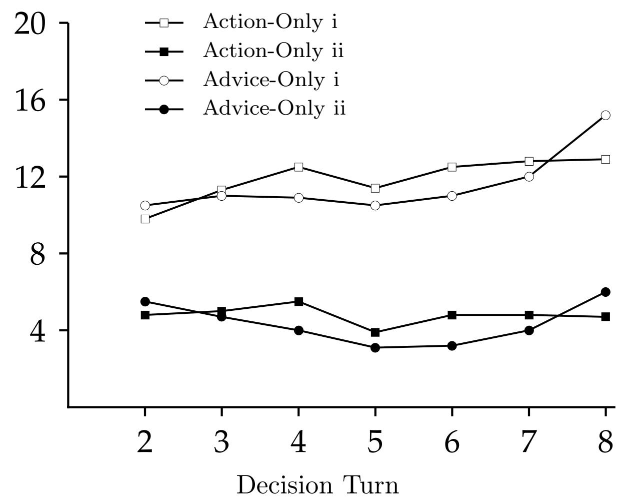

The mean of the absolute value of the distance between the subjects' cutoff and the theoretical cutoff (i); and the mean of the absolute value of the distance between the subjects' cutoff and zero (ii).

ment, it appears overall that there is a significant difference between the Action-Only and Advice-Only cutoffs. Most interestingly, this difference in fact is compositional, representing the distribution of decisions over our concurring and contrary categories and not differences in how persuasive predecessors' actions and advice are once they are followed. Put differently, the difference in behavior is the result of the fact that subjects follow advice much more frequently. When advice is followed, or action imitated, it is done so with the same intensity.

Figure 7: Unconditional mean cutoffs by decision turn.

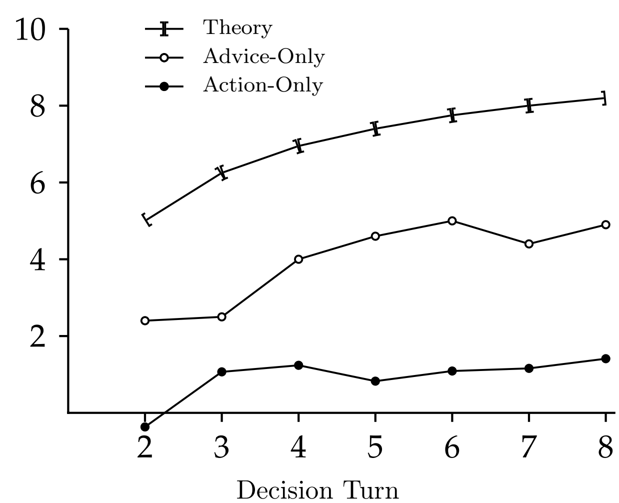

The regression analysis presented in Table 2 summarizes our discussion so far. We regress the mirror image transformation of the cutoff set by subjects on the decision-turn at which this cutoff was set, as well as a dummy variable which takes a value of 1 if the experiment containing the observation was Advice-Only. Referring to Table 2, the experimental dummy is highly significant and positive, indicating that the cutoffs set by a subject for any decision-turn in the Advice-Only experiment can be expected to be 3.05 units higher than the cutoff set under identical circumstances in the Action-Only experiment. This indicates more confidence in advice than action. The significance of the dummy variable in the regression clearly indicates that the process of setting cutoffs is different in the Advice-Only and Action-Only experiments and that this difference is consistent with our observations above that subjects are more persuaded by advice offered than by action observed.[6].

Table 2: Cutoff behavior in Advice-Only and Action-Only experiments.

	Coefficient	Std. Err.	**t**	*p*-value
Constant	-0.51	0.475	-1.064	0.288
D(Advice-Only)	3.05	0.336	9.093	0.000
Turn 3	0.74	0.628	1.183	0.237
Turn 4	1.62	0.628	2.573	0.010
Turn 5	1.67	0.628	2.661	0.008
Turn 6	2.02	0.628	3.218	0.001
Turn 7	1.75	0.628	2.835	0.005
Turn 8	2.12	0.628	3.387	0.001

Obs.$=$ 1050, $R^2 = .088$

5.2 Question 2

Which information—advice or action—is more valued by the subjects? And, under what circumstances do subjects offer advice, different from their action?

The Action-Plus-Advice experiment provides us with an extremely good opportunity to try to separate the impact of advice versus action on behavior. If the predecessor's action and advice differ, then two cases can be observed. The predecessor chooses A and advises B or the predecessor chooses B and advises A. Based on either of these occurring, the successor subject could choose to set either a negative cutoff (concurring with the action observed, as action A is more likely to be chosen) or a positive one (concurring with the advice received,

[6]GLS random-effects estimators and robust variance estimators for independent data and clustered data yield similar results.

as action B is more likely to be chosen). This defines four contingencies depicted in Table 3. Most interestingly, when the advice and action of one's predecessor differ, the successors are far more likely to choose an action consistent with the received advice than the observed action. In fact, in 60.2 percent of the cases where the advice offered differs from the action, subjects chose to follow the advice they received rather than to imitate their predecessor's action; only 24.1 percent of the time did they imitate the action taken, and 15.7 percent of the time they were neutral and chose a cutoff of zero.

Table 3: Advice taking in the Action-Plus-Advice experiment.

Successor / Predecessor	Choose **A** Cutoff (-)	Choose **B** Cutoff (+)	Cutoff = 0
Action **A**/Advice **B**	13 (15.66%)	33 (39.76%)	6 (7.23%)
Action **B**/Advice **A**	17 (20.48%)	7 (8.43%)	7 (8.43%)

Table 3 looks at behavior when the advice received is different from the action observed. But we might also ask whether getting advice that is consistent with the action taken makes a subject more likely to follow it, by setting more extreme cutoffs indicating stronger agreement. A priori we would expect this to be the case because we should expect subjects to view advice as more compelling when it agrees with a predecessors' action. Our results are mixed. First, as illustrated in Table 4, it is true that subjects are more likely to follow advice (as indicated by the sign of their cutoff) when it is backed up by action. If a subject is told to follow an action by a predecessor who took that action herself, the recommendation is followed 84.2 percent of the time, while the advice would be followed only 74.1 percent of the time in the Advice-Only experiment. When only the action is observed, it is imitated 44.2 percent of the time. So, it should be clear that a predecessor who does as she says is seen as being more believable than one whose advice cannot be backed up by action.

Table 4: Decision conformity with advice and action.

	Action Taken		
	Concurring	Neutral	Contrary
Action-Only	44.2%	16.6%	39.2%
Advice-Only	74.1%	9.1%	16.8%
Action-Plus-Advice	84.2%	7.0%	8.8%

In addition, Figure 8 shows that when subjects are in an Action-Plus-Advice experiment and receive advice consistent with the action they observe, they will only act in a contrary manner less often than they would in either the Action-Only or Advice-Only experiments. In fact, as we see in Figure 8, when they act in a contrary manner 80 percent of the time,

Figure 8: The distribution of contrary subjects.

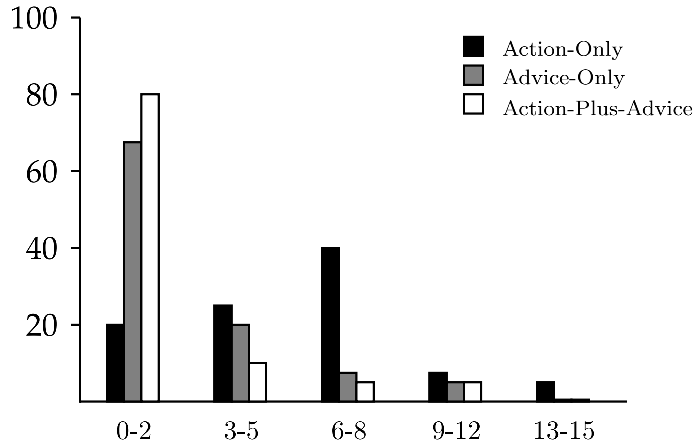

The percentage of the subjects who disagreed with the observed action/advice in less than two rounds, three to five rounds and so on.

they make two or less contrary decisions. This indicates that most subjects follow advice backed up by action most of the time.

On the other hand, as Figure 9 indicates, when subjects set a cutoff conditional on receiving advice that is consistent with the action they observe, the magnitude of the cutoff set does not differ very much from the cutoff set in the Advice-Only experiment (subjects there were not able to see whether the advice offered was consistent with the actions taken by their predecessors).[7] Again, the impact of advice, this time backed up by action, is to increase the fraction of time that decisions are followed. But once they are followed, the strength of commitment to the decision is practically identical.

Recall that our initial hypothesis was that, in any decision-turn, the cutoff set by a subject would be greatest (after mirror image transformation) in the Action-Plus-Advice experiment after receiving advice consistent with the action observed. The second highest cutoff was after being advised to choose an action in the Advice-Only experiment. The third highest was after observing an action in the Action-Only experiment; and the smallest was when one's predecessor chooses one action but advises another, in the Action-Plus-Advice experiment. To test this hypothesis we ran the following two regressions: we pooled our data and separated it into two sets. In one set, we pooled the data from the Action-Only experiment and the data from the Action-Plus-Advice experiment; in the other set, we did the same thing for the Advice-Only and Action-Plus-Advice experiments.

In the first case, we regressed the cutoff set by our subject on three dummy variables depicting whether the observation came from the Action-Only experiment, the Action-Plus-

[7] A set Wilcoxon tests detect no significant difference between the distributions of the cutoffs set in the Advice-Only and Action-Plus-Advice experiments for any turn.

Figure 9: Unconditional mean cutoffs by decision turn.

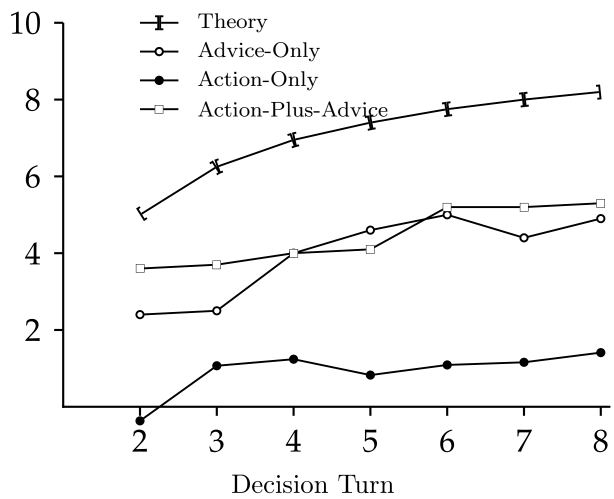

Advice experiment where advice was consistent with action, or the Action-Plus-Advice experiment with inconsistent action and advice. Using the Action-Only dummy as the baseline, only two dummies were coded. In the second case, we ran the same regression using the other data set and including the Advice-Only experiment as the baseline. The other right hand variables were the decision-turn dummies.

The results of these regressions confirm our hypotheses. In brief, at any decision-turn, subjects tend to set the highest cutoff when they are in the Action-Plus-Advice experiment and they receive advice consistent with the action they just observed. Their second highest cutoff is when they receive advice in the Advice-Only experiment; the third highest is when they view an action in the Action-Only experiment. Finally, they tend to set the lowest cutoff when they get conflicting advice from the action they observe. These results are presented in Table 5A and Table 5B.[8]

In both tables, the dummy variable D_1 depicts an observation where the advice-action pair is either AA or BB, D_2 depicts an observation where the advice-action pair is either AB or BA. In Table 5A, the constant term is associated with an observation from the Action-Only experiment while in Table 5B the observation comes from the Advice-Only experiment. Tables 5A and 5B provide support for our hypothesis. Looking at Table 5A, for example, a subject's cutoff increases when she receives advice that is consistent with the action observed; it decreases if the action and the advice disagree. However, the cutoff does not increase as the decision-turn increases, because the coefficient on each decision-turn is not significant.

[8] GLS random-effects estimators and robust variance estimators for independent data and clustered data yield similar results.

The results are slightly different in the regression presented in Table 5B that uses Advice-Only experiment as the baseline. Most notably, the Advice-Only subjects are so confident about their cutoff that observing an action consistent with that advice in the Action-Plus-Advice experiment has no significant impact on their cutoff (D_1 is not significantly different from zero). Seeing mixed Advice-Only lowers their cutoff, though. In addition, the decision-turn has a significant impact on the cutoff level from the fourth decision-turn onward. Finally, looking at absolute levels, the regressions imply that cutoffs are highest in situations where advice and action are the same in Action-Plus-Advice; second highest in Advice-Only; third highest in Action-Only; and lowest where advice and action are not the same, in Action-Plus-Advice.

Table 5A: Cutoff behavior in Action-Only and Action-Plus-Advice experiments.

	Coefficient	Std. Err.	**t**	p-value
Constant (Action)	0.042	0.048	0.870	0.384
$\mathbf{D}_1$(**AA** or **BB**)	0.347	0.042	8.157	0.000
$\mathbf{D}_2$(**AB** or **BA**)	-0.370	0.065	-5.611	0.000
Turn 3	0.045	0.058	0.773	0.440
Turn 4	0.023	0.059	0.391	0.696
Turn 5	0.058	0.059	0.994	0.320
Turn 6	0.100	0.059	1.707	0.088
Turn 7	0.118	0.059	1.996	0.046
Turn 8	0.033	0.059	0.561	0.575

Obs. = 1050, R^2 = .151

Table 5B: Cutoff behavior in Advice-Only and Action-Plus-Advice experiments.

	Coefficient	Std. Err.	**t**	p-value
Constant (Advice)	0.257	0.056	4.526	0.000
$\mathbf{D}_1$(**AA** or **BB**)	0.045	0.067	0.670	0.503
$\mathbf{D}_2$(**AB** or **BA**)	-0.214	0.079	-2.707	0.007
Turn 3	0.060	0.050	1.234	0.217
Turn 4	0.097	0.049	1.965	0.049
Turn 5	0.161	0.049	3.293	0.001
Turn 6	0.218	0.049	4.481	0.000
Turn 7	0.240	0.049	4.860	0.000
Turn 8	0.228	0.049	4.653	0.000

Obs. = 1050, R^2 = .030

Next we turn our attention the question of when subjects offer advice that differs from their action? We conjecture that subjects would overturn their action (offer advice that differs from the action taken) when they set extreme cutoffs relatively close to -10 or 10 and the signal they observe is consistent with their beliefs but very close to their cutoff.

Say a subject is relatively certain that the sum of the signals is negative, so that B is the correct decision. Under those assumptions, she sets a cutoff of 8.5. This means that for any signal below 8.5, she would like to choose action B. Such a strong cutoff clearly indicates a strong belief that the sum of all signals is negative and that B is the profitable action. If the signal she receives is below 8.5 and very negative, say -5, she feels pretty safe in her belief and is happy to have B chosen for her because her signal is such a strong confirmation of her prior belief. For the same reason, she also would be happy to offer B as advice. We will call a signal that is below a positive cutoff, or above a negative cutoff, a consistent signal because it confirms the subject's belief about the true state of the world. That is, it gives evidence that the state she already believes in is more likely to occur.

What if her signal is instead 8.49? Here, the signal is still consistent with her belief, and she will still choose action B, but its magnitude shakes her faith in her prior. It is exactly under these circumstances that we expect that subjects will offer advice opposite to the action they took. If their cutoff was not extreme, then we do not expect overturns, because non-extreme cutoffs indicate lack of strength in a subject's belief about the state of nature. Hence, any realization of the signal is not likely to cause the subject to overturn it in her advice giving. To summarize, we expect that subject overturns occur when the signal observed is consistent but only marginally so. If the signal were marginally inconsistent, for example 8.53 in our previous example, then we expect subjects would tend not to overturn the A action chosen.

Figure 10: The distribution of advice overturning subjects.

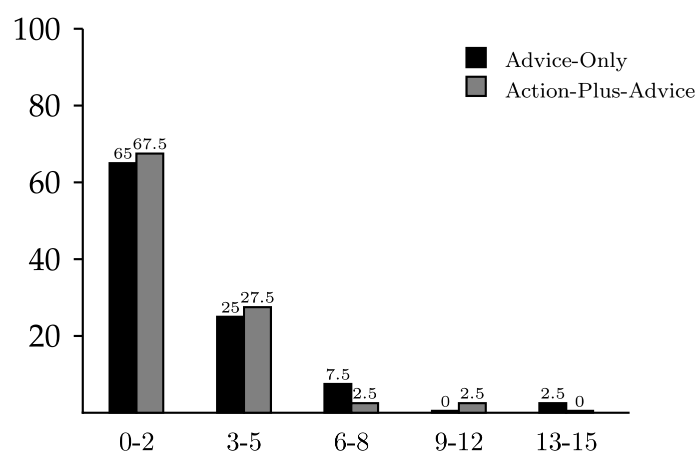

Our data largely support these conjectures with one noticeable exception. First, overturns are relatively rare, occurring only 17.5 percent of the time in the Advice-Only experiment and 15.8 percent of the time in the Action-Plus-Advice experiment. In addition, as Figure 10 illustrates, it is rare that any subject will overturn many outcomes even if she overturns some. For example, in the Advice-Only and Action-Plus-Advice experiments, 65 and 67.5 percent of our subjects ,respectively, offered advice that overturned two or fewer of the actions they took out of 15 trials, if they ever overturned. Only 10 percent and 5 percent overturned six or more outcomes in the Advice-Only or Action-Plus-Advice experiments respectively. In short, overturns are rare and infrequent for any given subject.[9]

To investigate this conjecture, we specified a Logit model in which the left-hand variable was a binary variable which took a value of 1 when there was an overturn and 0 otherwise. The right-hand variables consisted of the decision-turn, `Turn`; the absolute value of the cutoff set by the subject, `Abs`; the distance between the absolute value of the cutoff and the mirror of the signal received, `Dst`. We also entered an interaction term, `Product`, consisting of the distance, and a dummy that takes a value of 1 if the signal and the cutoff have the same sign and the signal is below (above) a negative (positive) cutoff. We also entered a dummy, `AcAd`, for the experiment that takes the value 1 for the Action-Plus-Advice experiment and 0 for the Advice-Only experiment. The results are presented in Table 6.

The model substantiates our conjecture. The probability of an overturn is increasing according to how extreme is the cutoff set by the subject (`Abs`), and decreasing in the distance between the cutoff set and the signal received (`Dst`). Also, there is no experiment effect, in the sense that the overturn behavior does not seem to be affected by whether we look at the Advice-Only or the Action-Plus-Advice data (`AcAd`). Finally, the interaction (`Product`) term is insignificant. Still, overturns occur when the cutoff is high and close to signal.

Table 6: Overturning behavior.

	Odds Ratio	Std. Err.	**Z**	*p*-value
`AcAd`	0.973	0.169	-0.156	0.876
`Turn`	0.990	0.044	-0.216	0.829
`Abs`	1.093	0.034	2.912	0.004
`Dst`	0.834	0.019	-7.957	0.000
`Product`	1.070	0.049	1.487	0.137

Log likelihood = -429.34084, Pseudo R^2= 0.0925

[9] In addition, the number of overturns is insensitive to the decision turn in that, when a regression of the number of overturns on decision turn is run in both the Action-Plus-Advice and Advice-Only experiments, the slope coefficients are not significantly different from zero.

5.3 Post-Signal Action-Only

One possible explanation why subjects tend to follow their predecessor's advice more in the Advice-Only experiment than they imitate their action in the Action-Only experiment is that advice was offered after the predecessor observed her signal while her action was determined by her cutoff which was stated before her signal was observed. Under this explanation advice is followed more because it is perceived as being more reasoned. To investigate this claim we conducted the Post-Signal Action-Only experiment, which provides a useful benchmark for our preceding analysis. In this experiment, subjects observe their private signal and their predecessor's action before taking their own action. Thus, if both advice and action are equally persuasive then subjects will not be more willing to follow advice given to them by their predecessor in the Advice-Only experiment than to copy their action in the Post-Signal Action-Only.

In the Post-Signal Action-Only experiment, if the sign of the private signal agrees with the action taken by the predecessor (the predecessor took action A (B) and the signal is positive (negative)) then the subject should trivially follow the predecessor's action. The inference problem for each subject becomes more interesting if her signal disagree with her predecessor's action. Figure 11 presents, by turn, the fraction of time that the predecessor's action was mimicked in the subset of decisions where the signal disagreed with the predecessor's action, and compares it to the theoretical probability. The number of observations at each decision turn is given in parenthesis. It is evident from Figure 11 that subjects did not mimic the action of the predecessor as frequent as the theory predicts. Comparing the data presented in Figure 11 with our analysis of concurring and contrary decision in the Action-Only and Advice-Only experiments above clearly indicates that subjects have more confidence in advice than action. Over all decision-turns except the first, in the Post-Signal Action-Only experiment, when their signal disagreed with the predecessor's action, subjects chose the same action as their predecessor only 46.5 percent of the time. In comparison, subjects stated a cutoff consistent with the advice they receive 74.1 percent of the time in the Advice-Only experiment, but only 44.2 percent of the time in the Action-Only experiment.

Informational efficiency is an important question for economists and the efficiency of information aggregation is one of the main concerns of the theoretical and experimental literatures on social learning. A central result of the literature is that conformity might result in most agents choosing the wrong action (where the right action is defined relative to the information available in the economy). This outcome is both informationally inefficient and Pareto inefficient. Thus, we compare the likelihoods of correct actions in the various treatments. Table 7 below presents the results (notice that uninformed random actions will be correct half the time). Most interestingly, the likelihood of correct actions is significantly higher in the experiments containing advice, Advice-Only and Advice-Plus-Action, and there

Figure 11: Unconditional mean cutoffs by decision turn.

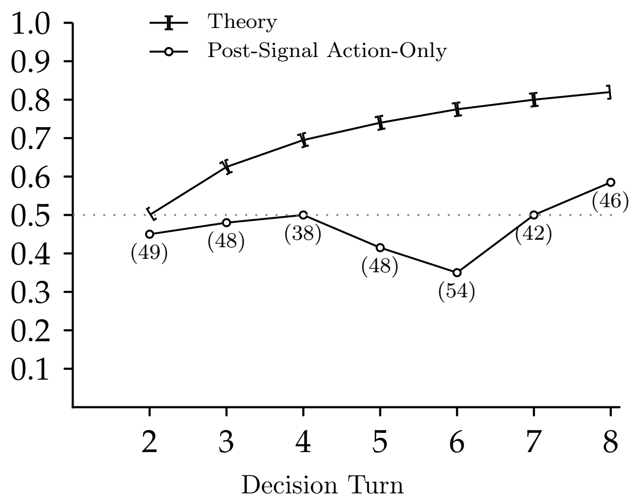

are no significant differences between the likelihoods of correct actions in the Post-Signal Action-Only and Action-Only experiments.

Table 7: The likelihood of correct actions by experiment.

	Prop.	Std. Err.	[95% Conf. Interval]
Post-Signal Action-Only	0.632	0.018	[0.597, 0.667]
Action-Only	0.625	0.020	[0.586, 0.663]
Advice-Only	0.720	0.018	[0.684, 0.756]
Action-Plus-Advice	0.768	0.017	[0.734, 0.802]

6 Conclusion

Social learning is a process in which economic agents learn by observing the actions of others. In the real world, though, people also learn from advice. In this paper, we introduced advice giving into a standard social learning problem. We designed the experiment so that both pieces of information should, in equilibrium, be equally informative (in fact, identical). The basic regularities of our experiment may be summarized as follows: first, subjects are more willing to follow the advice given to them by their predecessor than to copy her action. Second, subjects' behavior with advice is closer to the prediction of the theory than without it. Finally, and perhaps most importantly, the presence of advice increases subjects' welfare. Our findings therefore suggest that models of social learning must be modified in order to account for the observed behavior. To determine which factors are important in explaining decision making with advice in a variety of settings, it will be necessary to investigate a larger class of social learning situations in the laboratory. This is perhaps one of the most important topics for future research. Progress in this area requires both new theory and new experimental data.

References

[1] Anderson, L. and C. Holt (1997) "Information Cascades in the Laboratory." American Economic Review, 87(5), pp. 847-62.

[2] Banerjee, A. (1992) "A Simple Model of Herd Behavior." Quarterly Journal of Economics, 107(3), pp. 797-817.

[3] Banerjee, A. and D. Fudenberg (2004) "Word-of-Mouth Learning." Games and Economic Behavior, 46, pp. 1-22.

[4] Bikhchandani, S., D. Hirshleifer and I. Welch (1992) "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascade." Journal of Political Economy, 100(5), pp. 992-1026.

[5] Çelen, B. and S. Kariv (2004a) "Distinguishing Informational Cascades from Herd Behavior in the Laboratory." American Economic Review, 94(3), pp. 484-497.

[6] Çelen, B. and S. Kariv (2004b) "Observational Learning Under Imperfect Information." Games and Economic Behavior, 47(1), pp. 72-86.

[7] Çelen, B. and S. Kariv (2005) "An Experimental Test of Observational Learning under Imperfect Information." Economic Theory, 26(3), pp. 677-699.

[8] Ellison, G. and D. Fudenberg (1993) "Rules of Thumb for Social Learning." Journal of Political Economy, 101, pp. 612-643.

[9] Ellison, G. and D. Fudenberg (1995) "Word of Mouth Communication and Social Learning." Quarterly Journal of Economics, 110, pp. 93-126.

[10] Goeree, J., R. McKelvey, T. Palfrey, and B. Rogers (2005) "Self-Correcting Information Cascades." Review of Economic Studies, forthcoming.

[11] Hung, A. and C. Plott (2001) "Information Cascades: Replication and an Extension to Majority Rule and Conformity-Rewarding Institutions." American Economic Review, 91(5), pp. 1508-1520.

[12] Kübler, D. and G. Weizsäcker (2003) "Limited depth of reasoning and failure of cascade formation in the laboratory." Review of Economic Studies, 71, pp. 425-441.

[13] Schotter, A. (2003) "Decision Making in the Face of Naive Advice." American Economic Review Papers & Proceedings, 93(2), pp. 196-201.

[14] Smith, L. and P. Sørensen (2000) "Pathological Outcomes of Observational Learning." Econometrica, 68(2), pp. 371-398.