

An Experimental Test of Advice and Social Learning

Boğaçhan Çelen[†]
Columbia University

Shachar Kariv[‡]
UC Berkeley

Andrew Schotter[§]
New York University

February 15, 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, 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., New York, NY, 10003 (Email: 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 naive 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 is 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 ba-

sis 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 naive advice on behavior in experimental games. Schotter (2003, 2005) 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

¹In 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.

best guess as to what the correct action was. Figure 1 summarizes our experimental treatments and procedures.

[Figure 1 here]

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 s_1, s_2, \dots, s_8 that are independently and uniformly distributed over the support $[-10, 10]$. Sequentially, each agent $n \in \{1, \dots, 8\}$ has to make a binary irreversible decision $x_n \in \{A, B\}$ where action A is profitable if and only if $\sum_{i=1}^8 s_i \geq 0$, and action B otherwise.

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

$$x_n = A \text{ if and only if } \mathbf{E} \sum_{i=1}^n x_i \mid x_{n-1} \geq 0$$

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

$$x_n = A \text{ if and only if } \hat{x}_n \geq \mathbf{E} \sum_{i=1}^n x_i \mid x_{n-1}.$$

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

$$x_n = \begin{cases} A & \text{if } \hat{x}_n \geq \hat{c}_n, \\ B & \text{if } \hat{x}_n < \hat{c}_n, \end{cases} \quad (1)$$

where

$$\hat{c}_n(x_{n-1}) = \mathbf{E} \sum_{i=1}^n x_i \mid x_{n-1} \quad (2)$$

is the optimal cutoff which accumulates all the information revealed to agent n from her predecessor's action. Thus, \hat{c}_n is sufficient to characterize agent n 's behavior and the sequence of cutoffs $\{\hat{c}_n\}$ 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{c}_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

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

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

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

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

It follows immediately from (3) that the cutoff rule partitions the signal space into three subsets: $[-1, \bar{c}_n)$, $[\bar{c}_n, \underline{c}_n)$ and $[\underline{c}_n, 1]$. For high-value signals

$\underline{x}_n \in [\underline{x}_n, 1]$ and symmetric low-value signals $\bar{x}_n \in [-1, \bar{x}_n)$ agent n follows her private signal and takes action A or B respectively. As Figure 2 illustrates, in the intermediate subset $[\bar{x}_n, \underline{x}_n)$, which we call an imitation set, private signals are ignored in making a decision and agents imitate their immediate predecessor's action. Furthermore, since $\{\bar{x}_n\}$ and $\{\underline{x}_n\}$ 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.

[Figure 2 here]

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

$$x_n = A \text{ if and only if } x_n \geq -\mathbf{E} \prod_{i=1}^{n-1} x_i | a_{n-1}.$$

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

$$\hat{x}_n(a_{n-1}) = -\mathbf{E} \prod_{i=1}^{n-1} x_i | a_{n-1} \quad (4)$$

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 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 $\hat{c}_n(a_{n-1})$ given by (1) is the same as $\hat{c}_n(x_{n-1})$ 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 $\hat{c}_n(x_{n-1})$ as given by (2), which are identical to those in the truthful equilibrium $\hat{c}_n(a_{n-1})$.

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 $\bar{c}_n(\bar{x}_n)$ instead of $\bar{c}_n(x_n)$. 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 $\{\hat{c}_n\}$.

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 $\hat{p}_1 = 0$ 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 $E[u_2 | p_1, a_1 = A] = 10p - 5$. Thus it is optimal for the second agent to take action A if and only if $\bar{p}_2 \geq 5 - 10p$. Likewise, if the advice received is $a_1 = B$, it is optimal for the second agent to take action A if and only if $\bar{p}_2 \geq 10p - 5$. Thus, after adding noisy advice to the model, the second agent's cutoff rule is

$$\hat{p}_2(p_1, a_1) = \begin{cases} \frac{1}{2} - 5 + 10p & \text{if } a_1 = A, \\ \frac{1}{2} + 5 - 10p & \text{if } a_1 = B. \end{cases}$$

Because $\bar{p}_2 < 0$ and $\underline{p}_2 > 0$ ($\bar{p}_2 = -\underline{p}_2$ 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

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

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

$$\hat{c}_n(a_{n-1}) = -\mathbf{E} \left[\sum_{i=1}^{n-1} x_i \mid x_{n-1}, a_{n-1} \right] \quad (5)$$

is the optimal cuto 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 cuto leading to action A (B) advises her successor A (B) if the realization of her signal is closer to 10 (–10) than to her cuto; 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 *cuto* 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 *cuto* 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) *cuto*, she demonstrates concurrence by selecting a negative (positive) *cuto*, 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) *cuto*, then she disagrees with her predecessor. We say that such decisions are contrary decisions. Finally, neutral decisions are carried out by choosing a zero *cuto*, 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.

Figure 1. The experimental design and procedures.

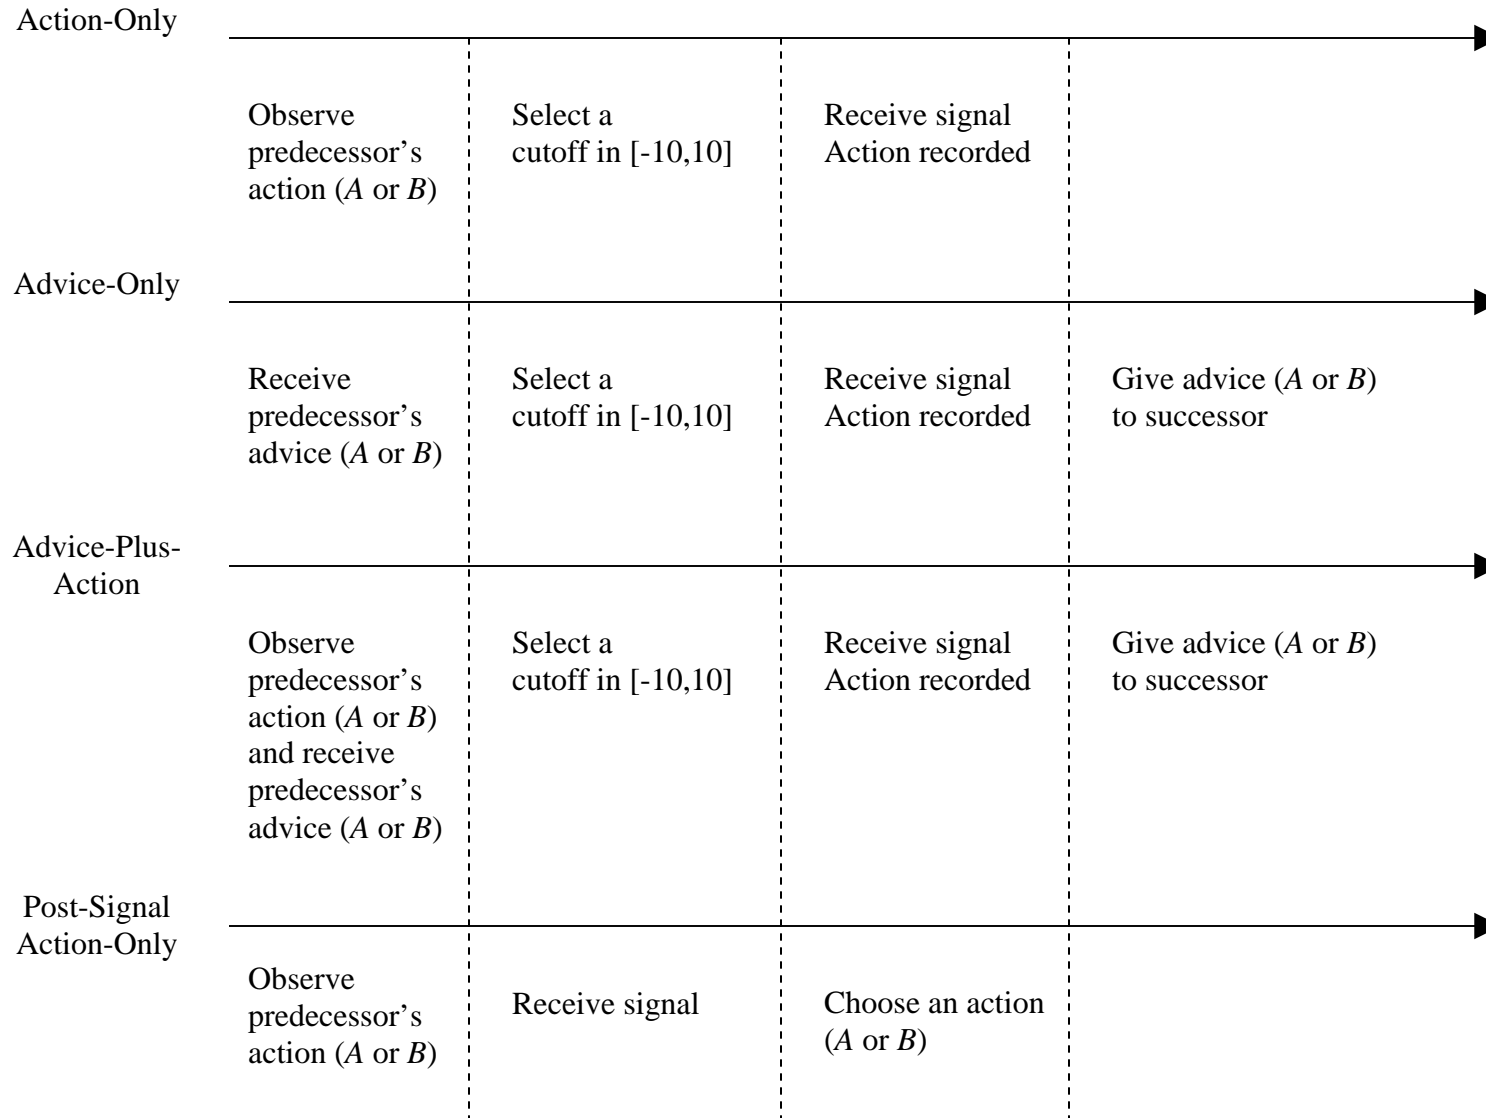

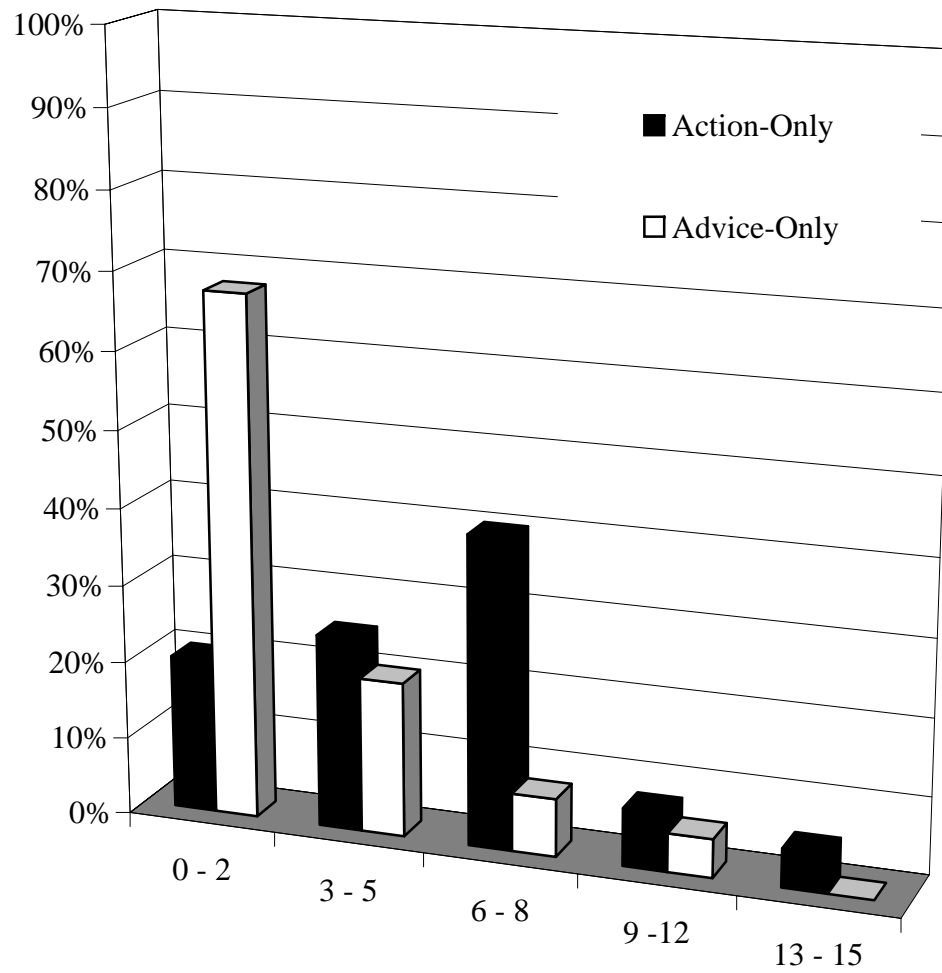


Figure 2: The process of cutoffs and imitation sets




Imitation sets monotonically increase in n regardless of the actual history of actions, and thus, over time, it is more likely that imitation will arise.

Figure 3: The distribution of contrary subjects




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

Figure 4: Mean cutoffs by decision turn in concurring decisions




Conditional means where the conditioning is done on whether the subject's decision was a cooccurring decision.

Figure 5: Mean cutoffs by decision turn in weakly concurring decisions



Conditional means where the conditioning is done on whether the subject's decision was weakly concurring (we include the neutral decisions in our means).

Figure 6: The strength of disagreement



The absolute value of the distance between the cutoff chosen and that which would be set if the subject acted according to the theoretical cutoff rule (1), and between the cutoff chosen and zero (2).

Figure 7: Unconditional mean cutoffs by decision turn





Figure 8: The distribution of contrary subjects



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

Figure 9: Unconditional mean cutoffs by decision turn

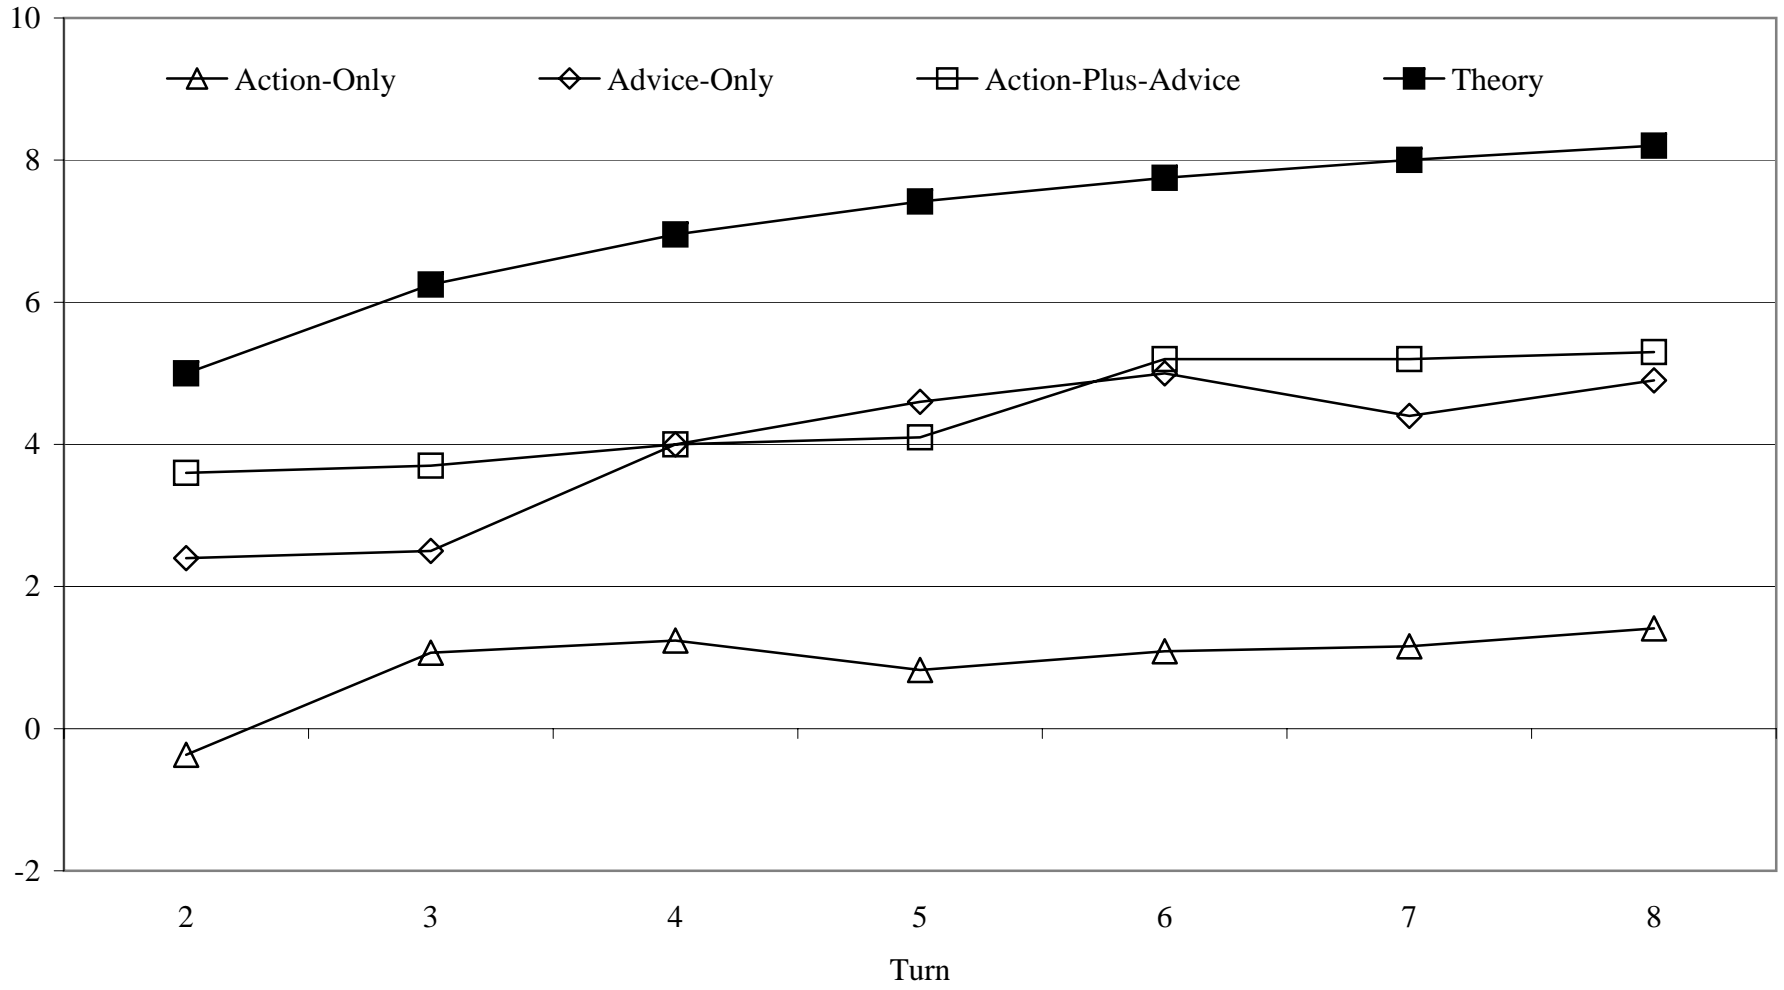


Figure 10: The distribution of advice overturning subjects

