

An Experimental Test of Advice and Social Learning*

Boğaçhan Çelen[†]
Columbia University

Shachar Kariv[‡]
UC Berkeley

Andrew Schotter[§]
New York University

June 17, 2010

Abstract

Social learning describes any situation in which individuals learn by observing the behavior of others. In the real world, however, individuals learn not just by observing the actions of others, but also learn from advice. This paper introduces advice giving into the standard social-learning experiment of Çelen and Kariv (2005). The experiments are designed so that both pieces of information—action and advice—are equally informative (in fact, identical) in equilibrium. Despite the informational equivalence of advice and actions, we find that subjects in a laboratory social-learning situation appear to be more willing to follow the advice given to them by their predecessor than to copy their action, and that the presence of advice increases subjects' welfare.

*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 (Xlab) at the University of California, Berkeley. We acknowledge helpful discussions with Colin Camerer, Gary Charness, Dan Friedman, Jacob Goeree, Teck Ho, Charles Holt, Tom Palfrey, Matthew Rabin, Barry Sopher, and Georg Weizsäcker. This paper has also benefited from suggestions by the participants of 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, 508-1 Evans Hall # 3880, Berkeley, CA 94720 (E-mail: kariv@berkeley.edu, URL: <http://emlab.berkeley.edu/~kariv/>).

[§]Department of Economics, New York University, 19 W. 4th Street, New York, NY, 10012 (Email: as7@nyu.edu, URL: <http://homepages.nyu.edu/~as7/>).

1 Introduction

Whether choosing a restaurant, adopting a new technology, or investing in a portfolio, an individual's decisions can reveal useful private information. So in social settings, where agents can observe one another's actions, it is rational for them to try to learn from one another. This interaction 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 has been argued that the striking uniformity of social behavior is one implication of social learning.¹

At the same time, the standard social-learning model has several special features that are quite restrictive and deserve further examination. Perhaps most importantly, the odd aspect of the social-learning literature is that it does not really accurately reflect social behavior. In the real world, while people learn by observing the actions of others, they also learn from their advice. For example, people choose restaurants not only by noting which of them are popular, but also by receiving advice about them. Similarly, people choose doctors not only by seeing how crowded their waiting rooms are, but also by asking recommendations about which physician is more qualified. Furthermore, people make their decisions in many situations by relying only on the so-called *naïve* advice of nonexperts such as friends, neighbors, and coworkers. Thus, social learning tends to be far more social than economists describe it.

In this paper, we introduce advice giving into a standard social-learning situation of the type that has been already investigated theoretically by Çelen and Kariv (2004a) and experimentally by Çelen and Kariv (2004b, 2005). In our experimental design, a sequence of subjects draw private signals from a uniform distribution over $[-10, 10]$. The decision problem is to predict whether the sum of all subjects' signals is positive or negative and to choose an appropriate action, A or B . A is the profitable action when this sum is positive and B is the appropriate action when this sum is negative. However, instead of choosing action A or B directly, after observing their immediate predecessor's action (A or B) or receiving advice from their immediate predecessor about which action (A or B) to choose or both and before receiving their own private signal, subjects are asked to select a cutoff that would result in action A being chosen if the signal they receive is greater than the cutoff and in action B being chosen if the signal they receive is less than the cutoff. Subjects are informed of their private signal

¹For surveys see: Gale (1996), Bikhchandani et al. (1998), and Chamley (2004).

only after they report their cutoff; their action is then recorded accordingly. In addition, in the treatments containing advice, after the action is recorded the subject is prompted to give a binary piece of advice, A or B , to her immediate successor.

We use two treatments containing advice in our experimental design. In the **Advice-Only** treatment, each subject only receives her immediate predecessor's advice as to which action to take, A or B . In the **Action-Plus-Advice** treatment, each subject observes the action chosen by her immediate predecessor and also receives her advice. In both treatments the subjects' payoffs are a function of the payoffs achieved both by themselves and by their successor, 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), which deal with the case in which each subject can observe only her immediate predecessor's action. We thus call the experiment of Çelen and Kariv (2005) the **Action-Only** treatment. Aside from the information structure, the treatments containing advice are identical to the experiment of Çelen and Kariv (2005). That is, all the treatments use the same procedures, but the information structure is different.

Most importantly, we design the **Action-Only** and **Advice-Only** treatments 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 action has been recorded. Despite this informational equivalence, we find that subjects are far more willing to follow the advice given to them by their predecessor than to copy their action. As a consequence, in the presence of advice, subject behavior is much more consistent with the predictions of the theory, and the presence of advice increases subjects' welfare.

A possible concern about the 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** treatment subjects first state a cutoff which determines their action and only then receive their signal, whereas in the **Advice-Only** and **Action-Plus-Advice** treatments subjects first observe their signal and then advise their successor. We therefore conducted a **Post-Signal Action-Only** treatment in which subjects observe their private signal and their predecessor's action before taking an action, A or B , directly. This treatment is informationally equivalent to the **Action-Only** and **Advice-Only** treatments. The results of the **Post-Signal Action-Only** treatment reinforces the finding that subjects appear to be more willing to follow the advice given to them by their predecessor

than to copy her action, and that the presence of advice increases subjects' welfare. This establishes that the impact of advice is not an artifact of the experimental design.

The rest of this paper is organized as follows. Section 2 discusses some situations in which advice plays an important role. In Section 3 we formulate the research questions that are subsequently answered. Section 4 summarizes the experimental design and procedures. Section 5 describes the theoretical model that lies behind the experimental design. The results are contained in Section 6. Some concluding remarks and important topics for further research are contained in Section 7.

2 Applications

We will not attempt to review the vast body of work on advice. Advice is an aspect of many disciplines, including economics, business, and even psychology. One important area in which advice is prevalent and salient is *professional advice*. The term professional advice refers to services rendered by experts.² Two inherent features of markets for professional advice—concerns about reputation, and competition—potentially alter the informational content of advice. It is well documented that “[p]rofessional advisers are often concerned with their reputation for being well informed, rather than with the decisions made on the basis of their recommendations[,]” as noted by Ottoviani and Sørensen (2006a). Similarly, competition among experts can distort the information transmitted through interaction. As argued in Ottoviani and Sørensen (2006b), markets for professional advice often take the form of contests (such as the semiannual *Wall Street Journal* Forecasting Survey, *Wall Street Journal* annual list of All-Star Analysts, etc.) where experts are evaluated based on the relative performance of their opinion.

Sometimes people learn from experts. At other times they need information that is not available from these professional sources and then they must try to find the information by seeking advice in their local environment. This type of advice is often referred to as *naïve advice*. Many studies have demonstrated that naïve advice has welfare-improving effects in many environments. For example, Chaudhuri et al. (2006) show that in a public-good experiment, “advice generates a process of social learning that leads to high contributions and less free-riding,” and Steinel et al. (2007) show that the

²For examples see, security analysts: Graham (1999), Hong et al. (2000), and Welch (2000); mutual fund managers: Chevalier and Ellison (1999); economic forecasters: Lamont (2002).

existence of naïve advice significantly improves effectiveness and efficiency in negotiations.

In the real world, naïve advice is essential in *viral marketing*, or viral advertising, i.e. marketing techniques that use preexisting social networks to create and spread viral messages. In this way, information percolates through the social network over time and people revise their decisions as new information arrives. The classic example of viral marketing is Hotmail.com, one of the first free Web-based e-mail services. Hotmail advertised itself in outgoing mail from their users by attaching a simple tag, “Get your private, free e-mail at www.hotmail.com,” at the bottom of every e-mail sent out. Once the message is sent by users who have a very large number of contacts, it spreads rapidly throughout the entire population.

Advice is also crucial for the functioning of consumer-generated advertising, which recently became an important channel in which marketers and consumers interact in value cocreation processes. Consumer-generated advertising refers to advertising that extensively uses consumer-generated media such as blogging, podcasting, video, and wikis. It seems reasonable to generalize that consumer-generated advertising is analogous to transmitting information in the form of advice. Perhaps the function of advice becomes more transparent in the way online stores such as Amazon.com or iTunes.com use consumer-generated content. These online stores almost always provide a section where user reviews of the products are submitted for the consideration of other buyers. In other words, in addition to a product’s popularity (consumers’ action), the reviews (consumers’ advice) also provide information for the buyers.

In this experiment we restrict our attention to naïve advice and ignore the complications of strategic behavior and reputation motivations in professional advice in order to focus on behavior motivated by purely social learning. Nonstrategic behavior is simpler to analyze and is also adequate for comparing several prominent models of social learning. However, our results so far do suggest a number of possible extensions for applying a similar methodology to professional advice in order to examine the impact of strategic behavior and reputation on the informational content of advice. The paper thus mainly contributes to the large and growing body of work on the influence of naïve advice on behavior in experimental games but also makes some contributions regarding the influence of professional advice. Schotter (2003, 2005) provide comprehensive, though now somewhat dated, reviews of the experimental work which clearly demonstrate that subjects tend to give good advice and to follow the advice of others to a remarkable extent.

Finally, the paper also contributes to a large literature on social learning. Banerjee (1992) and Bikhchandani et al. (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) investigate the social-learning model of Bikhchandani et al. (1992) experimentally and replicate informational cascades in the laboratory. Among others, Hung and Plott (2001), Kübler and Weizsäcker (2003), Çelen and Kariv (2004b, 2005), and Goeree et al. (2007) extend Anderson and Holt (1997) to further investigate possible explanations for informational cascades.

3 Research Questions

In this section we ask two questions that can be explored using the experimental data. 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** treatment (observing the predecessor’s action) and the other performing our **Advice-Only** treatment (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 the subject choosing A greater in the **Advice-Only** treatment?

Question 1. *Do subjects tend to follow advice more often than action when they observe each under identical circumstances?*

In the **Action-Plus-Advice** treatment, subjects both receive advice and observe the action taken by their predecessor. A natural question is whether this results in different behavior than that observed in the **Advice-Only** or **Action-Only** treatments. In fact, the **Action-Plus-Advice** treatment can give us some insight into whether subjects actually value advice more than action, because in some cases subjects gave advice that differed from the action they took. In those cases, the question is which datum the subject thinks more informative and why.

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

4 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 (Xlab) at UC Berkeley, as well as from the earlier experiment of Çelen and Kariv (2005). We will designate the new treatments as the **Advice-Only**, **Action-Plus-Advice**, and **Post-Signal Action-Only** treatments, and the earlier experiment as the **Action-Only** treatment. All treatments used the same basic procedures, but they differed with regard to the information received by subjects. We will explain these informational regimes shortly.

In each of the treatments 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, they were also read aloud by an experimental administrator.³ Participation fees 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 toward uniform behavior.⁴

Each experimental session entailed fifteen independent rounds, each 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 (each with two decimal points) 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 the **Action-Only**, **Advice-Only**, and **Action-Plus-Advice** treatments, upon being called to participate and before being informed of her private signal, the subject first received some information relevant to decision making (either the predecessor's action, the predecessor's advice, or both, depending on the treatment). As in Çelen and Kariv (2004b, 2005), after receiving this information, each subject was asked to select a number

³Sample instructions, including the computer program dialog windows, are available at http://emlab.berkeley.edu/~kariv/CKS_I_A1.pdf.

⁴Participants' working stations were isolated by cubicles, making it impossible for participants to observe each other's screens or to communicate. We also made sure that all participants remained silent throughout the session.

between -10 and 10 (a cutoff) that would result in the subject taking action A if the signal was above the cutoff and action B if the signal was below the cutoff. Action A was profitable if and only if the sum of the eight numbers was positive and action B otherwise. The subject was informed of the value of her private signal only after she submitted her decision. Then the computer recorded her decision as A if the signal was higher than the cutoff she selected. Otherwise, the computer recorded her action as B .

The **Action-Only**, **Advice-Only**, and **Action-Plus-Advice** treatments use the same procedures, but the information structure is different. In the **Action-Only** treatment, subjects were able to observe only the action, A or B , taken by their immediate predecessor. In the **Advice-Only** treatment, when subjects were called upon to make their decision they were not able to observe their predecessor's action. Rather, they received advice from their immediate predecessor as to what the correct action (A or B) to take was. In the **Action-Plus-Advice** treatment, subjects were able not only to receive advice from their immediate predecessor but also to observe her action. In both cases, subjects gave their 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** treatment 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 treatment did not contain advice, and subjects were not asked to select a cutoff that determined the choice of action A or B .

After all subjects had made their decisions, the computer informed everyone what the sum of the eight numbers actually was. Everyone whose decision determined that their action would be 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 that determined their action would be B earned \$2 if the sum was negative, and nothing otherwise. In addition, in the treatments containing advice, everyone earned \$1 if their successor took the correct action. This was paid to induce subjects to give advice that was their best guess as to what the correct action was. At the end of a session, subjects were paid in private. Figure 1 summarizes our experimental treatments and procedures.

[Figure 1 here]

5 Theory and Predictions

In this section, we discuss the theoretical implications of the model tested using the four treatments in the laboratory. Çelen and Kariv (2004a) provide 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 optimal to offer advice equal to the action chosen. As a result, substituting advice for actions in our experiment cannot convey more information. This implies that the environment in the **Advice-Only** treatment is not informationally richer than the environment in the **Action-Only** treatment.

5.1 Preliminaries

Suppose that the eight agents receive private signals $\theta_1, \theta_2, \dots, \theta_8$ 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 $\sum_{i=1}^8 \theta_i \geq 0$, and action B is profitable otherwise.

It follows immediately 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 whether the realization of the underlying decision-relevant event will be $\sum_{i=1}^8 \theta_i \geq 0$ 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 treatments.

5.2 Action-Only

The Decision Problem In the **Action-Only** case, except for the first agent, everyone observes only her immediate predecessor’s action. Since agents do not know any of their successors’ actions, agent n ’s optimal decision rule—conditional on the information available to her—is

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

Çelen and Kariv (2004a) show that the optimal decision takes the form of this *cutoff strategy*:

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

where

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

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 agents follow the cutoff strategy (1) and (2)) as the primitive of the experimental design and of our analysis.

The Cutoff Process Clearly, agent n 's cutoff rule, $\hat{\theta}_n$, can take two different values, conditional on whether agent $(n-1)$ took action A or action B , which we denote by

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

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

$$\hat{\theta}_n(x_{n-1}) = \begin{cases} -5 - \frac{\hat{\theta}_{n-1}^2}{20} & \text{if } x_{n-1} = A, \\ 5 + \frac{\hat{\theta}_{n-1}^2}{20} & \text{if } x_{n-1} = B, \end{cases} \quad (3)$$

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

It follows immediately from (3) that the cutoff rule partitions the signal space into three subsets: $[-10, \bar{\theta}_n)$, $[\bar{\theta}_n, \underline{\theta}_n)$, and $[\underline{\theta}_n, 10]$. For high-value signals $\theta_n \in [\underline{\theta}_n, 10]$ and symmetric low-value signals $\theta_n \in [-10, \bar{\theta}_n)$, agent n follows her private signal and takes action A or B respectively. In the intermediate subset $[\bar{\theta}_n, \underline{\theta}_n)$, which we call an *imitation set*, agent n ignores her private signal when she makes a decision, and all agents imitate their immediate predecessor's action. are ignored in making a decision and agents imitate their immediate predecessor's action. Furthermore, as Figure 2 illustrates, since $\{\bar{\theta}_n\}$ and $\{\underline{\theta}_n\}$ are decreasing and increasing sequences respectively, the imitation sets $[\bar{\theta}_n, \underline{\theta}_n)$ 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 on their private signal.⁵

⁵Çelen and Kariv (2004a) show that this has an important implication: beliefs and

[Figure 2 here]

5.3 Advice-Only

Next, we investigate the differences between the decision problems underlying our **Action-Only** and **Advice-Only** treatments. 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 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

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

is the optimal cutoff which includes all of the information revealed to agent n from her predecessor's advice. There are only 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 there are no other equilibria in 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 $a_n = x_n$, 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 (4) is the same as $\hat{\theta}_n(x_{n-1})$ given by (2), and the optimal advice rule is to give advice

actions are not convergent but cycle forever. Despite this instability, over time private information is increasingly ignored and decision makers become increasingly likely to imitate their predecessors. Consequently, behavior is typified by longer and longer periods of uniform behavior, punctuated by (increasingly rare) switches.

equal to her chosen action, $a_n = x_n$. Throughout the paper, whenever we refer to the theoretical sequence of cutoffs, we mean the unique equilibrium cutoffs in the **Action-Only** case $\hat{\theta}_n(x_{n-1})$ given by (2), which are identical to those in the truthful equilibrium $\hat{\theta}_n(a_{n-1})$ given by (4). Figure ?? also depicts the sequence of cutoffs in the truthful equilibrium in the **Action-Only** case.

The Mirror and Babbling Equilibria The truthful equilibrium is not the only equilibrium 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.

In the mirror equilibrium, agents advise their successor to take the opposite action to theirs, $a_n \neq x_n$; the successor believes that the advice given to her by the predecessor is opposite to her predecessor’s action; and she sets her cutoffs optimally according to (4), given her 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 her predecessor actually took was B (A) and thus sets her cutoff optimally at $\underline{\theta}_n$ ($\bar{\theta}_n$) instead of $\bar{\theta}_n$ ($\underline{\theta}_n$). 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 $\{\hat{\theta}_n\}$ given by (2) and depicted in Figure ??.

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 the successor, 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 Next, 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 her successor to take the same action she took, $a_n = x_n$, if $p_n > 1/2$ and always to advise her to take the opposite action, $a_n \neq x_n$, if $p_n < 1/2$.

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 or B . If the advice received is $a_1 = A$, then a simple calculation shows that $\mathbb{E}[\theta_1 | p_1, a_1 = A] = 10p - 5$. 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

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

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

5.4 Action-Plus-Advice

In the **Action-Plus-Advice** case, agents are able not only to receive advice from their immediate predecessor, but also to observe her action; this opens up *signaling* possibilities. In such a situation, conditional on the information available to her, agent n 's optimal decision takes the form of the cutoff strategy given by (1), where

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

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 more sophisticated, and hence informationally richer, strategies. These strategies combine all four available action-advice pairs (x_{n-1}, a_{n-1}) to partition their signal

space into four subsets and thus convey more information to an agent’s successor. Hence, the informational pipeline in this case is less constrained and there exists more informationally rich equilibria, which we call *signaling equilibria*, than in the **Action-Only** and **Advice-Only** cases. 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.⁶

6 Experimental Results

The data from the **Action-Only**, **Advice-Only**, and **Action-Plus-Advice** treatments provide answers to the two research questions posted above. The **Post-Signal Action-Only** treatment, which does not contain advice or cutoff elicitation, provides a useful benchmark for our **Action-Only**, **Advice-Only**, and **Action-Plus-Advice** treatments. We will compare the behavior in the **Post-Signal Action-Only** treatment with the behavior in each of the other treatments at the end of this section.

6.1 Question 1

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

We focus on the data from the data from **Action-Only**, and **Advice-Only** treatments. At any turn n , the data generated by the choice are the cutoff $\tilde{\theta}_n$, the action taken $x_n = \{A, B\}$, and the advice given $a_n = \{A, B\}$ in the **Advice-Only** treatment. To organize these data, following Çelen and Kariv (2005), we first define decisions made by subjects as *concurring decisions* if the sign of their cutoff agrees with the action observed or advice received. For example, when a subject observes that her predecessor took action or gave advice A (resp. B) and adopts a negative (resp. positive)

⁶To illustrate signaling equilibria, consider an equilibrium in which everyone with a cutoff leading to action A (B) advises her successor to take action A (B) if the realization of her signal is closer to 10 (-10) than to her cutoff; she advises her successor to take action 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 signaling partition here to convey information about their signals.

cutoff, she demonstrates concurrence, since by selecting a negative (resp. positive) cutoff she adopts a higher probability of taking action A (resp. B). Similarly, if a subject observes action or receives advice A (resp. B) and selects a positive (resp. 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 or advice but simply entails a choice based on private information.

Table 1: Concurring, contrary, and neutral decisions in the **Action-Only** and **Advice-Only** treatments

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

Table 1 presents the percentages of concurring, contrary, and neutral decisions in the **Action-Only** and **Advice-Only** treatments. 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% of the time in the **Advice-Only** treatment, but only 44.2% of the time in the **Action-Only** treatment. Together with the neutral cutoffs, subjects tend to weakly agree (set a concurring or neutral cutoff) with advice 83.2% of the time in the **Advice-Only** treatment but only 60.8% of the time in the **Action-Only** treatment. These distributions of the concurring, contrary, and neutral decisions in the **Action-Only** and **Advice-Only** treatments are significantly different according to 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, the histograms in Figure ?? shows the distribution of concurring, neutral, or contrary decisions aggregated to the subject level. The horizontal axis measures the number of contrary decisions (those that 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. In the **Advice-Only** treatment, 67.5% of the subjects disagreed with the advice they received in one or two rounds. In the **Action-Only** treatment, subjects tended to disagree far more often: only 20.0% of the subjects disagreed in one or two rounds, and 40.0% of the subjects disagreed in six to eight rounds. The distributions presented in Figure 3 are significantly different according to Kolmogorov-Smirnov test (p -value 0.000).

[Figure 3 here]

The sign of the cutoffs as an indication of agreement or disagreement tells only part of the story because it ignores 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 or receives advice A and sets a cutoff close to -10 , then not only does she agree with the action observed or advice received, but she also does so very strongly since she will then almost surely take action A . In contrast, selection of a negative cutoff that is closer to zero clearly indicates a much weaker agreement. Since the cutoff strategy is symmetric around zero, the strength of agreement or disagreement is independent of the actual action observed (A or B) or advice received (to choose A or B). We therefore proceed as per Çelen and Kariv (2005) and in the **Action-Only** treatment transform the cutoffs in any turn $n > 1$ using the *mirror image transformation*

$$\tilde{\theta}_n = \begin{cases} |\hat{\theta}_n| & \text{if } x_{n-1} = A \text{ and } \hat{\theta}_n \geq 0 \text{ or } x_{n-1} = B \text{ and } \hat{\theta}_n < 0, \\ -|\hat{\theta}_n| & \text{otherwise.} \end{cases}$$

Analogously, we define the mirror image transformation in the **Advice-Only** treatment by replacing x_{n-1} with a_{n-1} . That is, we take the absolute value of the cutoffs in concurring decision points, and the negative of the absolute value of the cutoffs at contrary decision points. For example, if a subject observes action A or receives advice to choose 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 sets 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.

Figure 4A depicts, turn by turn, the theoretical cutoffs $\underline{\theta}_n$ in the truthful equilibrium given by (3) and the mean cutoff after mirror image transformation $\tilde{\theta}_n$ in the subset of concurring decisions in the **Advice-Only** and **Action-Only** treatments. It is evident from Figure 4A 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 their predecessor. In other words, once a subject has decided to follow the advice offered or imitate the action taken, she does so with equal intensity. Also note 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 4B shows that the situation is reversed in the **Action-Only** treatment, particularly in late decision turns, when we include neutral decisions in our sample.

[Figure 4 here]

Next, we focus on the complementary subset of contrary decisions. Once a subject decides not to follow her predecessor's action or advice, the intensity of her disagreement can be measured in several ways. Figure 5 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 truthful equilibrium cutoff rule given by (3). Second, we use the absolute value of the distance of the chosen cutoff from zero. As Figure 5 shows, the strength of disagreement is rather severe, because when subjects disagree with their predecessor, they tend to do so in quite an extreme way.

[Figure 5 here]

All of the results presented above condition our data on whether decisions are concurring or contrary. Figure 6 shows that if we consider the data regardless of agreement or disagreement, it appears overall that there is a significant difference between the mirror image of the cutoffs set in the **Action-Only** and the **Advice-Only** treatments. 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 the predecessors' actions and advice that have been followed are. Put differently, the difference in behavior is the result of the fact that subjects follow advice much more frequently than they imitate action. However, action is imitated with the same intensity that advice is followed.⁷

[Figure 6 here]

We next turn to regression analyses that examine the patterns in the data more systematically. Let $\tilde{\theta}_n$ be the data generated by the choices of subjects at decision turn $n > 1$, and let χ and χ_n be indicator variables for the

⁷Turn by turn, Wilcoxon rank-sum (Mann-Whitney) tests detect statistically significant differences between the cutoffs $\tilde{\theta}_n$ in the **Advice-Only** and **Action-Only** treatments in all decision turns. However, none of these differences are significant if we focus on the subset of concurring or contrary decisions. We note that the Wilcoxon test requires independence. The outcomes of games in which the same subjects appear are not independent. This biases the standard errors downward, increasing the likelihood of finding a significant treatment effect. We use the null of independence and recognize that there is no simple adjustment that will take care of the possible dependence problem.

Advice-Only treatment and decision turns, respectively. Our econometric specification has the form:

$$\tilde{\theta}_n = [\alpha_n + \beta_n \chi] \chi_n + \epsilon_n$$

where ϵ_n is assumed to be distributed normally with mean zero and variance σ_n^2 . We generate estimates of the α and β coefficients using a Tobit model that accounts for the censored distribution (the cutoffs are bounded between -10 and 10), and use robust standard errors that allow for clustering at the level of the individual subject.

Table 2 reports the estimation results. The last column presents the results for the full sample. There are marked differences between the α and β estimates. Except for the third decision turn, all the β estimates are significantly positive. This implies that the later a subject's turn is, the more she relies on the information revealed from advice, and that subjects are more likely to follow their predecessor's advice than to imitate her action. The magnitudes are very large, implying an overall shift in cutoffs $\tilde{\theta}_n$ from approximately 1.18 to 3.57. This is roughly consistent with the effect implied by the summary statistics presented above.

[Table 2 here]

6.2 Question 2

*Which information—advice or action—is more valued by the subjects?
Under what circumstances do subjects offer advice that is different from their action?*

We next turn to the **Action-Plus-Advice** treatment, which is identical to the **Action-Only** and **Advice-Only** treatments with the exception that each subject observes the action chosen by her immediate predecessor and also receives her advice. The **Action-Plus-Advice** treatment allows us to separate the impact of advice from the impact of action on behavior. If the predecessor's action and advice differ—for example, if the predecessor chooses A and advises B —then 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, as action B is more likely to be chosen). Overall, such overturns $x_{n-1} \neq a_{n-1}$ are relatively rare, accounting for only 17.5% of the data from the **Action-Plus-Advice** treatment, which is only marginally higher than the 15.8% from the **Advice-Only** treatment. As will be discussed below,

the fact that overturns are infrequent implies that subjects do not use the advice they give to hedge out the risk they face from their action.

To organize the data from the **Action-Plus-Advice** treatment, we adopt the convention that decisions made by subjects are defined as concurring or contrary decisions with respect to *advice*. That is, decisions are defined by whether the sign of the cutoff agrees or disagrees with advice received. A neutral decision is again defined as choosing cutoff zero, which does not favor any action, *A* or *B*. Table 3 presents the percentages of concurring, contrary, and neutral decisions in the **Action-Plus-Advice** treatment and compares them with the analogous percentages in the **Action-Only** and **Advice-Only** treatments reported in Table 1. In the **Action-Plus-Advice** treatment, subjects set a cutoff consistent with the advice they receive 84.2% of the time when $x_{n-1} = a_{n-1}$, but only 60.2 percent of the time when $x_{n-1} \neq a_{n-1}$. Hence, advice is more likely to be followed when backed by an action. The distributions of the concurring, contrary, and neutral decisions in the **Action-Plus-Advice** treatments are significantly different according to Kolmogorov-Smirnov test (p -value 0.000).

Table 3: Concurring, contrary, and neutral decisions in the **Action-Plus-Advice** treatment

	Concurring	Neutral	Contrary
Action-Plus-Advice ($x_{n-1} = a_{n-1}$)	84.2%	7.0%	8.8%
Action-Plus-Advice ($x_{n-1} \neq a_{n-1}$)	60.2%	15.7%	24.1%
Action-Only	44.2%	16.6%	39.2%
Advice-Only	74.1%	9.1%	16.8%

Figure 7 shows the distribution of contrary decisions in the **Action-Plus-Advice** treatment aggregated at the subject level and compares them with the corresponding distributions in the **Action-Only** and **Advice-Only** treatments depicted in Figure 3. We present the distribution for *all* decisions, as well as the distribution for the subset of decisions where the advice received was consistent with the action observed, $x_{n-1} = a_{n-1}$. The horizontal axis measures the number of concurring decisions and the vertical axis measures the percentage of subjects corresponding to each interval. For the subset of consistent decisions $x_{n-1} = a_{n-1}$, 80.0% of the subjects disagreed with the advice they received only in one or two rounds. This distribution is significantly different from the analogous distributions in the **Action-Only** and **Advice-Only** treatments using Kolmogorov-Smirnov tests (p -values 0.000). For the full sample, only 62.5% of the subjects disagreed with the advice they received in one or two rounds, and 35.0% disagreed in three to five rounds.

[Figure 7 here]

Figure 8 depicts, turn by turn, the mean cutoff after the mirror image transformation ($\tilde{\theta}_n$) in the **Action-Plus-Advice** treatment. We present the mean cutoffs for all decisions, as well as the mean cutoffs for the subset of decisions where the advice received was consistent with the action observed ($x_{n-1} = a_{n-1}$), and compare the cutoffs to those in the **Action-Only** and **Advice-Only** treatments. Figure 8 shows that the magnitude of the cutoffs set in the **Action-Plus-Advice** treatment does not differ much from the magnitude of the cutoffs set in the **Advice-Only** treatment. Hence, when backed up by action, the impact of advice is to increase the number of times that predecessor's decision is followed. But once it is followed, the strength of commitment to the predecessor's decision is practically identical.

[Figure 8 here]

We now turn to regression analyses that examine the patterns in the data from the **Action-Plus-Advice** treatment more systematically. Let χ^c and χ^i be indicator variables for the **Action-Plus-Advice** treatment when the predecessor's action is consistent or inconsistent with her advice, respectively. This generates the following econometric specification:

$$\tilde{\theta}_n = [\alpha_n + \beta_n \chi^c + \delta_n \chi^i] \chi_n + \epsilon_n$$

where ϵ_n is assumed to be distributed normally with mean zero and variance σ_n^2 . We again generate estimates using a Tobit model that accounts for the censored distribution, and use robust standard errors that allow for clustering at the level of the individual subject.

Table 4 reports the estimation results. In Table 4A, we focus on the data from the **Advice-Only** and **Action-Plus-Advice** treatments. In Table 4B, we repeat the estimations using the data from the **Action-Only** and **Action-Plus-Advice** treatments. The last column presents the results for the full sample. The most apparent cross-sectional feature of the beta series is that for each decision turn, the estimated δ coefficients in Table 4A and the estimated α coefficients in Table 4B are monotonic and significantly positive. The other coefficients in Tables 4A and 4B are not significantly different from zero in most turns. This implies that for a fixed decision turn n , the cutoffs $\tilde{\theta}_n$ in the **Action-Plus-Advice** treatment when the advice received is consistent with the action observed $x_{n-1} = a_{n-1}$ are the same as the cutoffs in the **Advice-Only** treatment and higher than the cutoffs in the **Action-Only** treatment. For the full sample, the estimates show

that the cutoffs $\tilde{\theta}_n$ are lowest in the **Action-Only** treatment, higher in the **Action-Plus-Advice** treatment when the advice received is inconsistent with the action observed $x_{n-1} \neq a_{n-1}$, and highest in the **Advice-Only** and **Action-Plus-Advice** treatment when the advice received is consistent with the action observed.

[Table 4 here]

Next we turn our attention to the question of when subjects *offer* advice that differs from their action. Recall that overturns are relatively rare, accounting for only 17.5% and 15.8% of the decisions in the **Advice-Only** treatment and the **Action-Plus-Advice** treatment, respectively. In addition, as Figure 9 illustrates, overturns are rare and infrequent for any given subject—in the **Advice-Only** and **Action-Plus-Advice** treatments, 65.0% and 67.5% of our subjects, respectively, offered advice that overturned two or fewer (if any) of the actions they took (out of 15 rounds), and only 10% and 5%, respectively, offered advice that overturned six or more actions.

[Figure 9 here]

Table 5 reports the results of a Logit regression that examines the extent to which overturning ($x_n \neq a_n$) is correlated with the decision turn n , the absolute value of the cutoff set by the subject $|\hat{\theta}|$, and the distance between the cutoff and the signal received $|\hat{\theta}_n - \theta_n|$. We also include an indicator variable for the **Action-Plus-Advice** treatment. We again use robust standard errors that allow for clustering at the level of the individual subject. The results show that the probability of an overturn increases in the absolute value of a subject’s cutoff, and decreases in the distance between a subject’s cutoff and her signal. The correlations with the indicator variable and the decision turn are insignificant. Thus, we conclude that subjects are more likely to overturn their action $x_n = a_n$ when they set a relatively extreme cutoff and the signal they observe is close to their cutoff. In the **Action-Plus-Advice** treatment, this outcome is consistent with a signaling equilibrium.

Table 5: Overturning behavior

	Odds Ratio	Std. Err.	Z	p -value
Action-Plus-Advice	0.965	0.296	-0.12	0,906
n	0.988	0.032	-0.38	0.703
$ \hat{\theta}_n $	1.119	0.027	4.62	0.000
$ \hat{\theta}_n - \theta_n $	0.830	0.026	-5.99	0.000

Log likelihood = -429.389, Pseudo R^2 = 0.0924.

6.3 Post-Signal Action-Only

One possible explanation of why subjects tend to follow their predecessor’s advice more in the **Advice-Only** treatment than they imitate their predecessor’s action in the **Action-Only** treatment is that the advice was offered *after* the predecessor observed her signal, while her action was determined by her cutoff which was set *before* she observed her signal. To investigate this claim we conducted the **Post-Signal Action-Only** treatment, which provides a useful benchmark for our preceding analysis. In this treatment, subjects observe both their private signal and their predecessor’s action before taking their own action. Thus, in this treatment the action taken by one’s predecessor *includes* information about the signal observed, just as the advice does in the **Advice-Only** treatment. If advice is followed more often in the **Advice-Only** treatment than actions are copied in the **Post-Signal Action-Only** treatment, then we can conclude that advice is more persuasive than actions even in those situations where they are based on identical information.

To make this comparison, however, we must be sure to compare the right situations. In the **Post-Signal Action-Only** treatment, if the sign of the signal agrees with the action taken by the predecessor—i.e., if the predecessor took action A (resp. B) and the signal is positive (resp. negative)—then the subject should obviously follow the predecessor’s action. The decision problem is more interesting if the sign of the signal disagrees with the predecessor’s action. Over all the decision turns except the first, in the **Post-Signal Action-Only** treatment when their signal disagreed with the predecessor’s action, subjects chose the same action as their predecessor only 46.5% of the time. In comparison, in the treatments involving advice we say that a cutoff agrees with the advice received or action taken when a subject observes that her predecessor gave advice or took action A (resp. B) and adopts a negative (resp. positive) cutoff. In the **Action-Only** treatment, the cutoffs agree with the action observed in only 44.2% of the decisions,

which is significantly lower than 74.1% in the **Advice-Only** treatment but very similar to the analogous frequency in the **Post-Signal Action-Only** treatment.

Next, we also compare the likelihoods of correct actions (defined relative to the information available) in the various treatments. Table 6 presents the results (note that uninformed random actions will be correct half the time). Most interestingly, the likelihood of correct actions is significantly higher in the treatments involving advice—the **Advice-Only** and the **Action-Plus-Advice** treatments—and there are no significant differences between the likelihoods of correct actions being taken in the **Post-Signal Action-Only** and the **Action-Only** treatments.

	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]

7 Conclusion

There are innumerable social and economic situations in which agents are influenced by the decisions of others. From the point of view of rational choice theory, however, the important question is why rational, maximizing agents should behave in this way. Several economic theories explain the existence of uniform social behavior. These include the benefits from conformity for its own sake, sanctions imposed on deviants, strategic complementarities, and social learning, which describes any situation in which agents learn by observing the actions of others. Among these theories, only social learning explains why a society settles on a single pattern of behavior. This is an important result and it helps us understand the basis for uniformity of social behavior.

At the same time, the odd aspect of the social learning literature is that it does not really accurately reflect social behavior. In the real world, while people learn by observing the actions of others, they also learn from their advice. In this paper, we introduce advice giving (which has not previously been explored in experimental studies) into the standard sequential social-learning problem. The games that make up the various treatments in our experiment differ only with respect to their information structure, but we

designed the experiment so that both pieces of information—action and advice—should, in equilibrium, be equally informative (in fact, identical).

The experiment generates sharp and suggestive results:

- At the aggregate level, what we find is the truly puzzling result that subjects in laboratory social-learning situations are more willing to follow the advice given to them by their predecessor than to copy her action. As a result, we find that imitation is much more frequent when subjects give and receive advice, but that it is less frequent than the theory predicts.
- We also focus on the data at the individual level. We use our cutoff elicitation technique to elicit subjects' beliefs. The signs of the cutoffs indicate agreement or disagreement. This enables us to look at the distribution of subjects in terms of the frequency with which they either agreed or disagreed with their predecessor's action or advice (or both). We find that subjects tended to disagree much more often with the action observed than the advice received.
- We find that once a subject has decided to imitate her predecessor's action or to follow her advice, she does so consistent with the theoretical predictions. More precisely, among the subjects who agree with their predecessor's action or advice there is a good degree of conformity with the theory, which we fail to observe in the aggregate data. Hence, the difference between the treatments with and without advice is a compositional difference reflecting the fractions of agreements and disagreements and not a difference in how persuasive a predecessor's action or advice that has been followed is.
- Also, we find that the presence of advice improves the accuracy of decisions and subjects' payoffs. The increase in the payoffs in the treatments with advice, relative to those without advice, is mainly attributable to the increasing number of subjects who follow their predecessor's advice even if it conflicts with their private information. That is, over time the information revealed by advice is relied upon more, and subjects become increasingly likely to imitate their predecessors; however, subjects do not tend to rely more over time on the information revealed by the predecessor's action.

We can account for the differences between the results of the treatments with and without advice and for discrepancies between the data and the

predictions of the theory, by introducing noisy individuals in the model. Çelen and Kariv (2005) test a structural model that describes subjects' behavior in the **Action-Only** treatment as a form of generalized Bayesian behavior that incorporates limits on the rationality of others. They conclude that in the **Action-Only** treatment, overall the follow-own-signal heuristic outperforms Bayes' rule as a predictor.

We repeat the exercise of Çelen and Kariv (2005) for the **Advice-Only** treatment.⁸ We find that subjects gradually increase their confidence in the information revealed by their predecessor's advice. As a result, the cutoffs process exhibits an upward trend showing that over time subjects tend to adhere more closely to Bayesian updating. This model generates results that are very similar to those produced by the linear Tobit model we employ in this paper. For that reason and to economize on space, we favor the reduced form approach, which provides a good fit and offers flexibility, tractability, and a straightforward interpretation.

Our conclusions that subjects follow naïve advice and that the subjects' behavior in the experiments with the advice are closer to the prediction of the theory than the behavior in the experiments without the advice are consistent with previous experimental evidence. These findings lack a proper theoretical explanation. As noted by Schotter (2003), “[d]espite the prevalence of reliance on advice, economic theory has relatively little to say about it.” One plausible explanation of why the presence of advice increases rationality is that the process of giving advice requires subjects to rethink the problem. This may be due to reputation concerns, responsibility, or some other psychological reasons. No matter what the reason is, if the subjects know that their predecessor *thought twice* before giving the advice, they are (rationally) more willing to follow their predecessor's advice than to imitate her action.

Although we are ultimately forced to leave our results as puzzles for theorists to ponder, they are sufficiently dramatic that they suggest that there might be similar results for more general cases. Our findings suggest that models of social learning need to 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.

⁸For the detailed analysis of Çelen and Kariv (2005), see the NYU C.E.S.S. working paper with the same title.

Our results reinforce the effectiveness of user-generated content as a marketing tool. In recent years, companies increasingly utilize user-generated content to increase brand awareness or to achieve other marketing objectives. The massive outpouring of viral marketing and viral advertising takes place via pre-existing social networks in which advertisers aim to get consumers to talk among themselves. In addition, numerous user reviews of products being sold by online sellers are submitted by regular visitors to their sites. Recent evidence suggests that consumers rely more on reviews of other customers (naïve advice) than on sale rankings (action). However, these observational studies are subject to identification problems because many crucial parameters and variables are unobserved. In the laboratory, by contrast, we can, in principle, observe all the relevant parameters and variables. This provides an opportunity to learn more about the empirical properties of advice and thus of the word-of-mouth marketing that is widely used in the real world.

References

- [1] Anderson, L., C. Holt. 1997. Information cascades in the laboratory. *American Economic Review* **87**(5) 847-62.
- [2] Banerjee, A. 1992. A simple model of herd behavior. *Quarterly Journal of Economics* **107**(3) 797-817.
- [3] Banerjee, A., D. Fudenberg. 2004. Word of mouth learning. *Games and Economic Behavior* **46**(1) 1-22.
- [4] Bikhchandani, S., D. Hirshleifer, I. Welch. 1992. A theory of fads, fashion, custom, and cultural change as informational cascade. *Journal of Political Economy* **100**(5) 992-1026.
- [5] Bikhchandani, S., D. Hirshleifer, I. Welch. 1998. Learning from the behavior of others: Conformity, fads, and informational cascades. *Journal of Economic Perspectives* **12**(3) 151-170.
- [6] Çelen, B., S. Kariv. 2004a. Observational learning under imperfect information. *Games and Economic Behavior* **47**(1) 72-86.
- [7] Çelen, B., S. Kariv. 2004b. Distinguishing informational cascades from herd behavior in the laboratory. *American Economic Review* **94**(3) 484-497.

- [8] Çelen, B., S. Kariv. 2005. An experimental test of observational learning under imperfect information. *Economic Theory* **26**(3) 677-699.
- [9] Chamley, C. 2004. *Rational herds economic models of social learning*. Cambridge University Press.
- [10] Chaudhuri, A., S. Grazino, P. Maitra. 2006. Social learning and norms in a public goods experiment with inter-generational advice. *Review of Economic Studies* **73**(2) 357-380.
- [11] Chevalier, J., G. Ellison. 1999. Career concerns of mutual fund managers. *Quarterly Journal of Economics* **114**(2) 389-432.
- [12] Ellison, G., D. Fudenberg. 1993. Rules of thumb for social learning. *Journal of Political Economy* **101**(4) 612-643.
- [13] Ellison, G., D. Fudenberg. 1995. Word of mouth communication and social learning. *Quarterly Journal of Economics* **110**(1) 93-126.
- [14] Gale, D. 1996. What have we learned from social learning? *European Economic Review* **40**(3-5) 617-628.
- [15] Goeree, J., R. McKelvey, T. Palfrey, B. Rogers. (2005) Self-correcting information cascades. *Review of Economic Studies* **74**(3) 733-762.
- [16] Graham, J. 1999. Herding among investment newsletters: Theory and evidence. *Journal of Finance* **54**(1) 231-268.
- [17] Hong, H., J.D. Kubik, A. Solomon. 2000. Security analysts' career concerns and herding of earning forecasts. *Rand Journal of Economics* **31**(1) 121-144.
- [18] Hung, A., C. Plott. 2001. Information cascades: Replication and an extension to majority rule and conformity-rewarding institutions. *American Economic Review* **91**(5) 1508-1520.
- [19] Kübler, D., G. Weizsäcker. 2003. Limited depth of reasoning and failure of cascade formation in the laboratory. *Review of Economic Studies* **71**(2) 425-441.
- [20] Lamont, O. 2002. Macroeconomic forecasts and microeconomic forecasters. *Journal of Economic Behavior and Organization* **48**(3) 265-280.

- [21] Ottaviani, M., P. N. Sørensen. 2006a. The strategy of professional forecasting. *Journal of Financial Economics* **81** (2) 441-466.
- [22] Ottaviani, M., P. N. Sørensen. 2006b. Professional advice, *Journal of Economic Theory* **126**(1) 120-142.
- [23] Schotter, A. 2003. Decision making in the face of naïve advice. *American Economic Review Papers & Proceedings* **93** (2) 196-201.
- [24] Schotter, A. 2005. Decision making with naïve advice. In *Experimental Business Research* Vol. 2, A. Rapoport and R. Zwick, editors, Springer, Dordrecht, The Netherlands.
- [25] Smith, L., P. N. Sørensen. 2000. Pathological outcomes of observational learning. *Econometrica* **68**(2) 371-398.
- [26] Steinel, W., A. E. Abele, C. K. W. De Dreu. 2007. Effects of experience and advice on process and performance in negotiations. *Group Processes & Intergroup Relations* **10**(4) 533-550.
- [27] Welch, I. 2000. Herding among security analysts. *Journal of Financial Economics* **58**(3) 369-396.

FIGURE 1: Experimental design and procedures

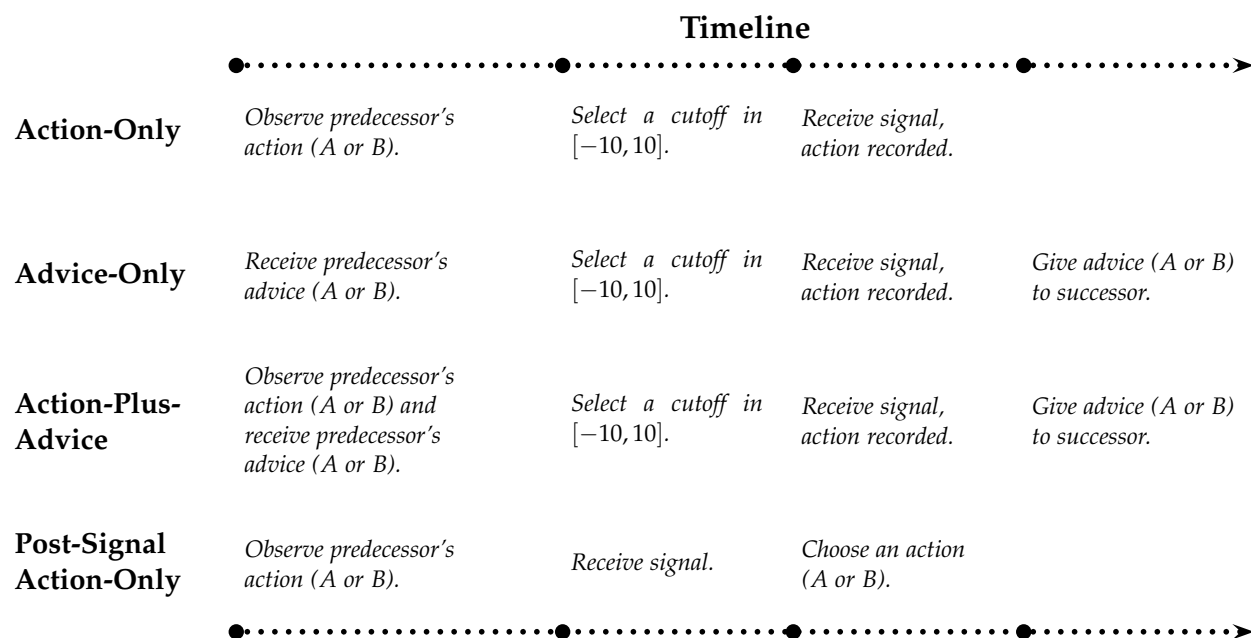
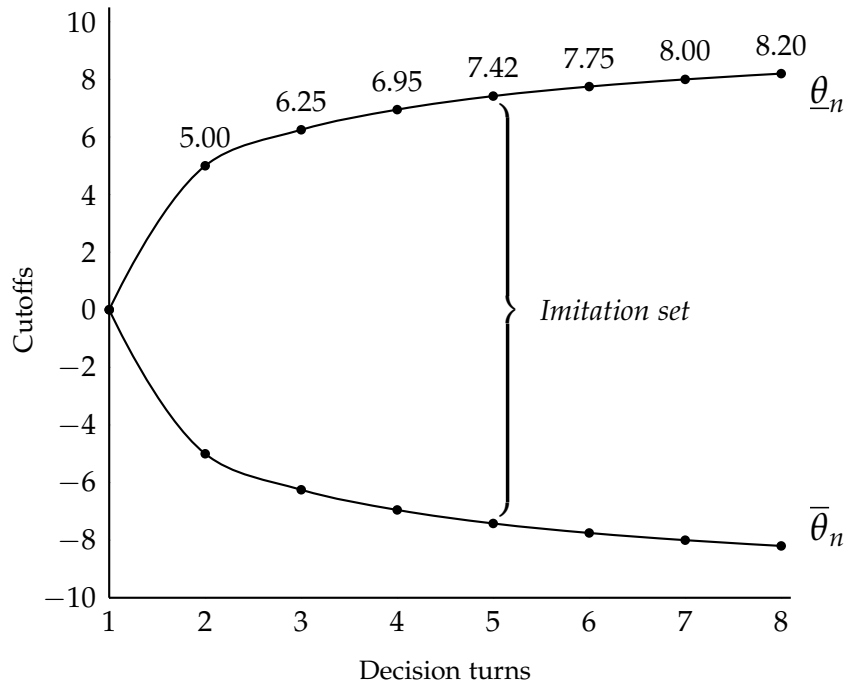
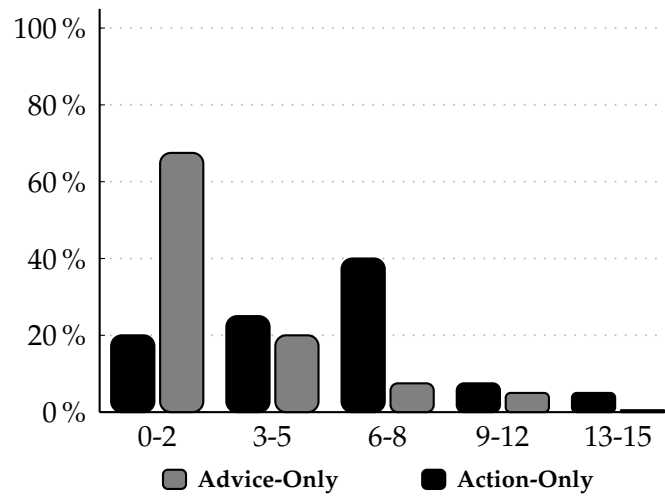


FIGURE 2: The process of cutoffs and imitation sets




Imitation sets monotonically increase in n regardless of the actual history of actions. Thus, overtime, it is more likely that imitation will arise.

FIGURE 3: The distribution of contrary subjects

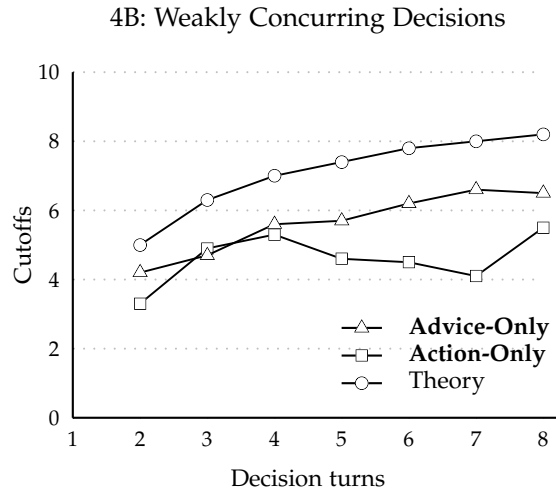


This graph shows the percent of subjects who disagreed with the observed action (advice) in fewer than two rounds, three to five rounds, and so on.

FIGURE 4: Mean cutoffs

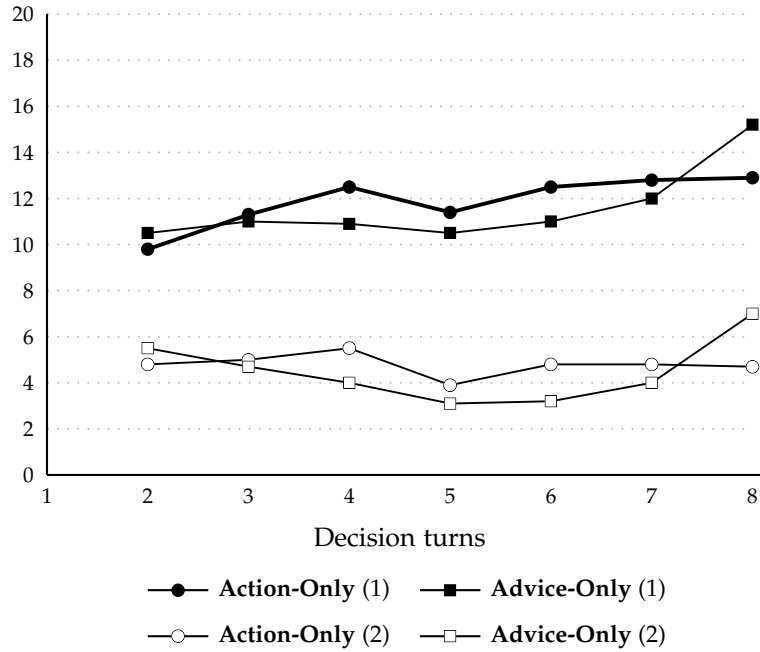


This graphs shows the conditional means where the conditioning is done on whether the subject's decision was a concurring decision.



This graphs shows the conditional means where the conditioning is done on whether the subject's decision was weakly concurring (neutral decisions are included).

FIGURE 5: The strength of disagreement



This graphs shows 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 6: Unconditional mean cutoffs





FIGURE 7: The distribution of contrary subjects



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

FIGURE 8: Mean cutoffs

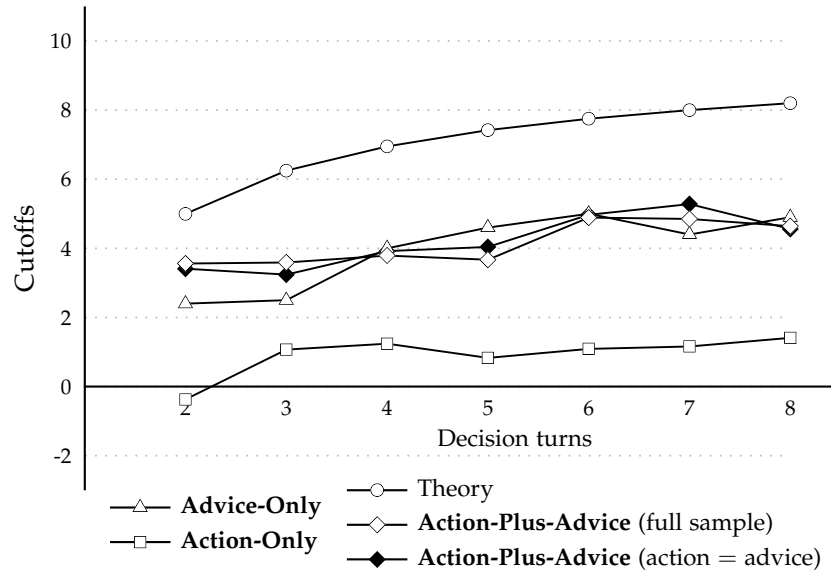



FIGURE 9: The distribution of overturning subjects



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