Decision Making with Naive Advice

By ANDREW SCHOTTER*

In many of the decisions we make we rely on the advice of others who have preceded us. For example, before we buy a car, choose a dentist, choose a spouse, find a school for our children, or sign on to a retirement plan, we usually ask the advice of others who have experience with such decisions. The same is true when we make major financial decisions. Here people easily take advice from their fellow workers or relatives as to how to choose stock, balance a portfolio, or save for a child's education. Although some advice we get is from experts, most of the time we make our decisions relying only on the rather uninformed word-of-mouth advice we get from our friends or neighbors. I call this "naive advice."

Despite the prevalence of reliance on advice, economic theory has relatively little to say about it. In this paper I will survey a set of experimental results (see Bogachan Celen et al., 2002; Ananish Chaudhuri et al., 2002; Raghuram Iyengar and Schotter, 2002; Schotter and Barry Sopher, 2002, 2003) that indicate that word-of-mouth advice is a very powerful force in shaping the decisions that people make and tends to push those decisions in the direction of the predictions of rational theory. More precisely I will demonstrate the following:

(i) Laboratory subjects tend to follow the advice of naive advisors (i.e. advisors who are hardly more expert in the task at hand than they are).
(ii) This advice changes their behavior in the sense that subjects who play games or make decisions with naive advice play differently than those who play identical games without such advice.
(iii) The decisions made in games played with naive advice are closer to the predictions of economic theory than those made without it.
(iv) If given a choice between getting advice or the information upon which that advice was based, subjects tend to opt for the advice, indicating a kind of underconfidence in their decision-making abilities that is counter to the usual egocentric bias or overconfidence observed by psychologists.
(v) The reason why advice increases efficiency or rationality is that the process of giving or receiving advice forces decision-makers to think about the problem they are facing in a way different from the way they would do if no advice were offered.

In all but two of the experiments reported on below, subjects engage in what are called "intergenerational games." In these games a sequence of nonoverlapping "generations" of players play a stage game for a finite number of periods and are then replaced by other agents who continue the game in their role for an identical length of time. Players in generation t are allowed to see the history of the game played by all (or some subset) of the generations who played it before them and can communicate with their successors in generation $t + 1$ and advise them on how they should behave. This advice is in two parts. First, in most of the experiments discussed below, subjects offer their successors a strategy to follow. After this they are free to write a free-form message giving the reasons why they are suggesting the strategy. These messages provide a treasure trove of information about how these subjects are thinking the problem through, and because they have incentives to pass on truthful advice (they are paid half of what their successors earn) I feel confident that this advice is in earnest. Hence, when a generation-t player makes a move, she has both history and advice at her disposal.

Actually my colleagues and I investigate three experimental treatments: In one we call the "baseline," when generation t replaces generation $t - 1$, players are allowed to see the history of play of all previous generations and

* Department of Economics, New York University, 269 Mercer Street, New York, NY 10003.

receive advice from their predecessors. This advice is almost always private between a generation $t - 1$ player and his progeny. In a second treatment, called the "history-only" treatment, subjects can see the entire history but receive no advice from their predecessors. Finally, in our third treatment, called the "advice-only" treatment, subjects can receive advice but can only view the play of their immediate predecessor's generation. In addition, players care about the succeeding generation in the sense that each generation's payoff is a function not only of the payoffs achieved during their generation but also of the payoffs achieved by their successors in the game after they retire. By comparing the play of subjects in these three treatments we can measure the impact of advice on behavior.

I. The Impact of Advice in Ultimatum Games (Schotter and Sopher, 2002)

Consider an Ultimatum Game with a \$10 endowment played as an intergenerational game where each generation plays once and only once before it is retired. In our experiments we had 81, 79, and 66 generations play this game under the three treatments described above, respectively.

Since this game is played intergenerationally with each generation playing once and only once, when a Proposer arrives in the lab he or she will see on the computer screen an amount advised to be sent. A Receiver will receive advice advising the minimum offer he or she should accept. Economic theory predicts that only a small amount, say, \$0.01 will be sent.

A. *Was Proposer Advice Followed?*

In Figure 1 we have plotted the amounts advised to be sent as well as the amounts actually sent by each generation in two treatments of our intergenerational Ultimatum game experiment: the baseline treatment where subjects can both receive advice and see the entire history of all generations before them and the advice-only treatment where subjects can receive advice but only see the history of their immediate predecessors. As can be easily seen, by and large subjects simply sent the amount they were advised to send. Advice was followed in a very direct way.

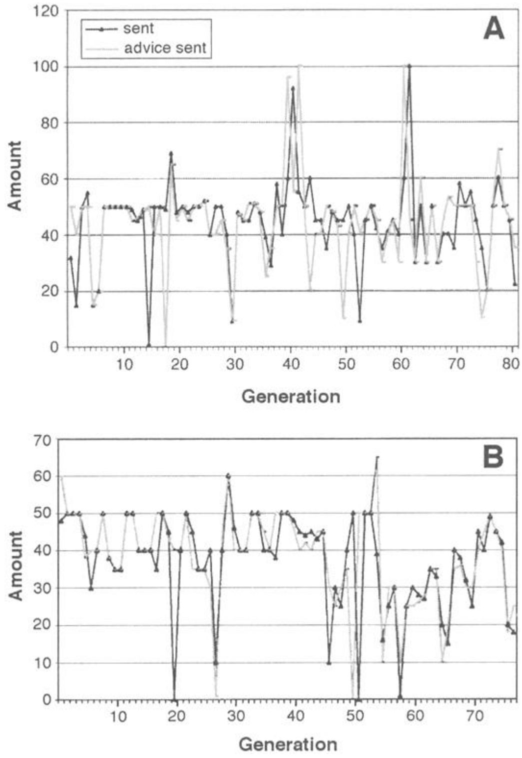

FIGURE 1. AMOUNT SENT AND ADVICE: (A) BASELINE TREATMENT; (B) ADVICE-ONLY TREATMENT

B. *Was Behavior Changed by Advice?*

Advice had a significant impact on behavior. For example, while the mean amount offered in the advice-only experiment over the last 40 generations was 33.68 it was 43.90 in the history-only treatment. Regressing offers on time indicates that only in the advice-only treatment does the time variable have a significant (and negative) coefficient. If one were to look at the histograms of offers generated by our advice-only and history-only experiments, one would see that the impact of advice is to truncate the right tail of the offer distribution. In fact while only 10 percent of the offers in the advice-only treatment were above 50, in the history-only treatment 10 percent of the observations were above 80. The variance of offers is significantly higher in the history-only treatment than in the advice-only treatment, indicating that history does not seem to supply a sufficient lesson for subjects to

TABLE 1—BATTLE OF THE SEXES GAME

	Column player	
Row player	1	2
1	150, 50	0, 0
2	0, 0	50, 150

guide their behavior in a smooth and consistent manner.

C. *Rejection Behavior*

Rejection behavior is also affected by advice. In Schotter and Sopher (2002) we estimate a logit model estimating the probability of acceptance as a function of the amount sent:

$$\Pr(x \text{ accepted}) = e^{a+bx}/(1 + e^{a+bx})$$

where x is the amount offered and the left-hand variable is a $\{0, 1\}$ variable taking a value of 1 if x is accepted and 0 otherwise. The results of this model estimation clearly indicate that the probability of accepting low offers (offers below 25) is significantly higher in the history-only treatment than in either the baseline of advice-only treatment. For example, while the probability that an offer of 10 is accepted is about 0.10 in the advice-only treatment, that probability increases to about 0.19 and 0.53 in the baseline and history-only treatments, respectively.

II. The Impact of Advice in Coordination Games: (Schotter and Sopher 2003)

Consider the Battle of the Sexes game shown in Table 1 when played in the lab as an intergenerational game.

A. *Was Advice Followed?*

In the baseline treatment of our Battle of the Sexes game, advice appears to be followed quite often, but the degree to which it is followed varies depending on the state last period. On average, for the row players it is followed 68.75 percent of the time, while for the column player it is followed 70 percent of the time.

In these experiments we measured the beliefs of each generation concerning their expectations of what strategies they expect their opponent to choose. We did this using a proper scoring rule, and this enabled us to define what a subject's best response was to those beliefs. Since in some instances the advice offered a subject was counter to their best-response action, we can measure the relative strength of advice by comparing how often the subjects chose one over the other. When advice and best responses differ, subjects are about as likely to follow the dictates of their best responses as to follow the advice they are given.

B. *Was Behavior Changed by Advice?*

One puzzle that arises from our Battle of the Sexes experiments is the following. While in the baseline we observe equilibrium outcomes 58 percent of the time (47 out of 81 generations), when we eliminate advice, as we do in history-only treatment, we only observe coordination in 29 percent of the time (19 out of 66 generations). When we allow advice but remove history, the advice-only treatment, coordination is restored and occurs 49 percent of the time (39 out of 81 generations).

These results raise what we call the "advice puzzle" which is composed of two parts. Part 1 is the question of why subjects would follow the advice of someone whose information set contains virtually the same information as theirs. In fact, as stated above, the only difference between the information sets of successive generations in our baseline experiment is the advice that predecessors received from their predecessors.

Part 2 of the advice puzzle is that, despite the fact that advice is private and not common-knowledge cheap talk, as in Russell Cooper et al. (1989), it appears to aid coordination in the sense that the frequency of equilibrium occurrences in our baseline (58 percent) and advice-only treatment (49 percent), where advice was present, is far greater than observed in the history-only treatment (29 percent), where no advice was present. While it is known that one-way communication in the form of cheap talk can increase coordination in Battle of the Sexes games (see Cooper et al., 1989), and that two-way cheap talk can help in other games, (see Cooper et al., 1992), how private communication of the type seen in our experiment works is an unsolved puzzle for us.

Finally, note that the desire of subjects to follow advice has some of the characteristics of an information cascade, since in many cases a subject is not relying on his or her own beliefs, which are based on the information contained in the history of the game, but is instead following the advice of a predecessor, who is just about as much a neophyte as is the present subject.

III. Would People Rather Have Advice or Data? (Celen et al., 2002)

If you had to make a decision would you rather be given some information concerning that decision or would you rather get advice from someone who has seen it. One would think that what you decide will depend on your estimate of your abilities as a decision-maker compared to those of the advice-giver as well as the informativeness of the data you might expect to process yourself.

To get at this question, Celen et al. (2002) investigated a social-learning experiment with a design that differed slightly from the intergenerational game experiments described above. In this experiment, eight subjects were brought into a lab and took turns making decisions sequentially in a random order. A round started by having the computer randomly select eight numbers from the set of real numbers $[-10, 10]$. The numbers selected in each round were independent of each other and of the numbers selected 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 her private signal. In practice, subjects observed their signals up to two decimal points.

The task of subjects in the experiment was to choose one of two decisions labeled A and B. Decision A was the correct decision to make if the sum of the eight private signals was positive, while B was correct if the sum of the private signals was negative. A correct decision earned \$2, while an incorrect one earned \$0. This problem was repeated 15 times, with each of the eight decision-makers receiving a new and random place in the line of decision-makers in each round.

Celen et al. (2002) used three treatments that differed in the information they allowed subjects to have. In one treatment (the action-only treatment), subjects could see the decision made by their predecessor in the line of decision-makers (so the fifth decision maker could see the decision of the fourth etc.) but no others and could not receive any advice from their predecessors. In another treatment (the advice-only treatment), subjects (except for the first one) could receive advice from their immediate predecessors telling them to either choose A or B, while in the final treatment (the advice-plus-action treatment), subjects could see both the decision their predecessor made and receive advice from him or her. Subject payoffs were equal to the sum of their payoffs over the 15 rounds in the experiment plus the sum of what their successors earned so each subject had an incentive to leave good advice.

TABLE 2—AGREEMENT AND CONTRARINESS IN ACTION-ONLY AND ADVICE-ONLY EXPERIMENTS

	Percentage of actions		
Subject type	Concurring	Neutral	Contrary
Sees action	44.2	16.6	39.2
Receives advice	74.1	9.1	16.8

The final feature of the experimental design, and the one that distinguishes it from other social-learning experiments, was that subjects did not directly choose a decision A or B, but rather set a cutoff level between -10 and 10. Once this cutoff was typed into the computer, it took action A for the decision-maker if her signal was above the cutoff specified and action B if it was not.

This design can help us answer the question stated above; would people prefer to have advice or information. For example Table 2 compares the actions of subjects who can only see the actions chosen by their immediate predecessor to those who cannot see what they have done but can receive an advised action. I have broken down the actions of subjects into those actions which agree with the action or advice of the predecessor (concurring decisions), those which disagree (contrary decisions), and those which neither agree or disagree with the actions or advice of one's predecessor (such actions are possible in this experiment since the subject can always set a zero cutoff which allows him to choose A or B with equal probability). By "agree" I mean that the subjects sets a negative cutoff when he is told or observes the A action and sets a positive cutoff when he is told or observes the B action.

As can be seen, subjects take actions that agree with the advice they receive 74.1 percent of the time yet copy the actions of their predecessors only 44.2 percent of the time. Actions disagree with advice only 16.8 percent of the time as compared with 39.2 percent for the experiment where only actions could be seen. In addition, behavior in the advice-only treatment was more efficient in the sense that subjects made more correct decisions and hence earned more money when advice was available. While earnings in the action-only experiment averaged \$18.8, they averaged \$21.8 in the advice-only experiments. These differences are significant at the 5-percent level using a Wilcoxon test. Finally, one of the main reasons why advice increases the payoffs and hence the welfare of our subjects is that it has a dramatic impact on the subjects' inclination to herd (make the same decision as their predecessor). While in our action-only experiments we observed herding of at least five subjects in a row in only 8 of the 75 rounds (10.7 percent), in the advice-only sessions herding occurred in 25 (33.3 percent) of the rounds. Finally, it is remarkable to note that in all experiments with advice all herds turned out to be on the correct decision.

In summary, it appears that, in this informational setting, words speak louder than actions, in the sense that subjects are more likely to follow the advice of their predecessors to take specific actions than they are to copy their behavior.

IV. Why Follow Advice? (Iyengar and Schotter, 2002)

The results above lead to the question of why advice should be so beneficial. Why should people give better advice to their successors than they gave to themselves? An experiment run by Iyengar and Schotter (2002) attempts to answer this question.

In this experiment subjects had to choose a number, e, between 0 and 100 called their decision number. They were told that they were playing against a computerized partner who would always choose the number 37. After this number is chosen, a random number is independently generated from a uniform distribution over the interval $[-a, +a]$ for both the subject and his computerized opponent. These numbers (the decision number and the random number) are then added together and a "total number" is defined for each the real and computerized players. Payoffs are determined by comparing the total numbers of the real and computerized subjects and awarding the real player a fixed payoff of M if his or her total is larger than that of the computerized opponent. If his or her total number is smaller, then he or she receives a payoff of m, $m < M$. The cost of the decision number chosen is given by a convex function $c(e) = e^2/r$, where r is a constant. This amount is then subtracted from these fixed payments to determine a subject's final payoff. By letting $r = 500$, $a = 40$, $M = 29$, and $m = 17.2$, and holding the computerized player's choice fixed at 37, our subjects face a rather simple decision problem with a quadratic payoff function whose peak is at 37.

This task was used by Antonio Merlo and Schotter (1999, 2003) to test the impact of information on learning. We used a "surprise quiz" method to test how well subjects learned the task put in front of them. In these experiments subjects performed the exact task as described above 75 times and received payoffs each period. When the 75 rounds were over they were surprised and told that they would play the game once more, but this time the stakes were multiplied by 75 so that they could earn for this one trial an amount equal to the sum of what they learned in all of the previous 75 rounds. Their choice in this high-stakes round should be a sufficient statistic for all that they have earned in the previous 75 rounds since the only way they can maximize their earnings in this round is by choosing that decision number that they feel is best.

In Iyengar and Schotter (2002) this exact experiment is repeated but in a slightly different manner. Instead of having one subject do the experiment alone, we sat another "advisor" subject next to him or her at the computer. The advisor makes written suggestions to the subject doing the experiment as to what he or she thinks is the best choice for that round. The chooser is free to follow this advice or not, but in one treatment the chooser is penalized for not doing so with a quadratic penalty function based on the difference between the action chosen and the action advised. In another treatment, no penalty is assessed for not following advice; it is simply cheap talk. The advisor's payoff is equal to three-quarters of that of his advisee. After the initial 75-round experiment run in this manner, both the adviser and advisee are separated and given surprise quizzes.

Miraculously, it appears as if the process of giving advice and receiving it greatly enhances the decision-making abilities of the Iyenga and Schotter subjects. For example, while the surprise-quiz choices of subjects in the no-advice treatment of Merlo and Schotter (1999) was 51.33, it was 36.1 for advisees and 33.5 for advisers in the Iyengar and Schotter (2002) experiment. The optimal decision number was 37. In other words, the process of giving advice seemed to focus the attention of advisers on the problem at hand in a manner that leads to greater learning on their part. Subjects seem to learn better when they give advice and when they receive it. We think this is true because giving and accepting advice causes a decision-maker not only to think through the problem another time, but to do so in a manner different from when making decisions alone.

This result offers a possible explanation of why it is advantageous to follow advice when it is offered and why advice is better than actions. The reason is that subjects learn better when they give advice, and that advice is therefore worth listening to. Further, a person receiving advice must contemplate whether or not to follow it, and this process may also foster learning. The process of advice-giving makes us think about the problems facing us differently than we tend to do when we are actually engaged in them.

V. Conclusion

This paper surveyed a set of papers all of which concentrate on the impact of naive advice on decisions-making. We have found that advice tends to be followed, changes behavior, and is welfare-improving. We have also explained that the process of giving and receiving advice fosters learning. Thus, in some sense, it is not surprising that subjects behave in a more rational manner when they make decisions under the influence of advice.

REFERENCES

Celen, Bogachan; Kariv, Shachar and Schotter, Andrew. "The Advice Puzzle: An Experimental Study of Social Learning Where Words Speak Louder than Actions." Mimeo, New York University, 2002.

Chaudhuri, Ananish; Schotter, Andrew and Sopher, Barry. "Talking Ourselves to Efficiency: Coordination in Inter-generational Minimum Games with Private, Almost Common, and Common Knowledge of Advice." Mimeo, New York University, 2002.

Cooper, Russell; DeJong, Douglas V.; Forsythe, Robert and Ross, Thomas W. "Communication in the Battle of the Sexes Game: Some Experimental Results." *Rand Journal of Economics*, Winter 1989, *20*(4), pp. 568–87.

______. "Communication in Coordination Games." *Quarterly Journal of Economics*, May 1992, *107*(2), pp. 738–71.

Iyengar, Raghuram and Schotter, Andrew. "Learning with a Meddlesome Boss: An Experiment in Sloppy Principal–Agent Relationships." Mimeo, New York University, 2002.

Merlo, Antonio and Schotter, Andrew. "A Surprise-Quiz View of Learning in Economic Experiments." *Games and Economic Behavior*, July 1999, *28*(1), pp. 25–54.

______. "Learning By Not Doing: An Experimental Investigation of Observational Learning." *Games and Economic Behavior*, January 2003, *42*(1), pp. 116–36.

Schotter, Andrew and Sopher, Barry. "Advice and Behavior in Intergenerational Ultimatum Games: An Experimental Approach." Mimeo, New York University, 2002.

______. "Social Learning and Coordination Conventions in Inter-generational Games: An Experimental Study." *Journal of Political Economy*, 2003 (forthcoming).