Running head: PAYING FOR ADVICE

Do we listen to advice just because we paid for it?

The impact of cost of advice on its use

Francesca Gino
Harvard Business School
Soldiers Field Road
Boston, MA
fgino@hbs.edu
phone: 617 495 6788
fax: 617 496 4066

Draft: February 3, 2006

Financial support from the Harvard Business School Division of Faculty Research and Development is gratefully acknowledged. I am especially grateful to Max Bazerman, Amy Edmondson, Andrew McAfee, Gary Pisano, Alvin Roth, Massimo Warglien, and Enrico Zaninotto for their insightful feedback on earlier drafts.

Abstract

When facing a decision, people often ask others for advice. Whether people use advice in a way that is helpful to them is not well understood. How do people evaluate the usefulness of the advice they receive? Drawing on aspects of behavioral decision theory, this paper argues that the cost of advice, independent of its quality, will affect how it is used. Two experiments were conducted. In both studies, participants were asked to answer different sets of questions about American history. Before they answered some of the questions, I made available advice on the correct answers. In the first study, participants had the opportunity to choose whether to get this advice for free or to pay a certain amount of money for it. In the second study, participants received either free or costly advice by default. The results show that participants used costly advice significantly more than they used free advice.

Keywords: Advice taking, behavioral biases, sunk costs, information use

When facing a decision, people often ask for advice before making a final commitment. Advice taking and advice giving are common activities across a wide range of contexts in the business world, from consumer behavior to product forecasting. While advice often plays a critical role in decision making, whether people use advice in a way that is helpful to them is not well understood.

Previous behavioral decision-making research has demonstrated that people use information irrationally, as in the case of anchoring (Tversky and Kahneman, 1974) and herd behavior (Banerjee, 1992; Anderson and Holt, 1996; 1997). In related work, Bastardi and Shafir (1998) have shown that people often delay making a decision until they receive irrelevant information, and rarely are aware that the information is irrelevant. Having pursued the information (at no cost), they use it as if it were relevant.

Would decision makers behave differently if the information they pursue was not free? In this article, I will show that, holding constant the quality of information that people receive, they weigh information differently according to the amount of money they pay for it. In particular, I examine a specific type of information – advice – and consider whether the cost of advice affects the degree to which people use it, under controlled, laboratory conditions. I conducted two experiments. In both studies participants were asked to answer different sets of questions about American history. Before answering some of the questions, they received advice about the correct answer. In the first study, participants had the opportunity to get advice about the correct answer either for free or by paying a certain amount of money. In the second study, participants received either free or costly advice by default. [1]

The results show that people tend to weigh the advice they receive *significantly more* when advice is costly than when advice is free.

[1] By costly advice, I do not mean 'expensive' advice but advice that costs money.

Advice taking in decision making

The role of advice in decision making has been the subject of two main streams of research. These streams differ greatly both according to the tasks and type of advice studied, which may help explain why they reach quite different conclusions.

Studies in the first research stream (see, for instance, Bogachan, Shachar, and Schotter, 2002; Chandhuri, Schotter, and Sopher, 2002; Iyengar and Schotter, 2002; Merlo and Schotter, 1999; 2003; Schotter and Sopher, 2003) show that, in laboratory settings, individuals tend to follow the advice of naïve advisors. In these studies, participants receive word-of-mouth advice before they make their decisions and play "intergenerational games" in which non-overlapping "generations" of players sequentially engage in a stage game for a finite number of periods; players are allowed to communicate in two subsequent generations (Schotter, 2003). The information passed from generation to generation is of two types: the history of the game played by all (or some subset) of the previous generations, and advice about how successors should behave. Schotter and his colleagues found that word-of-mouth advice plays a relevant role in shaping the decisions that people make, even if naïve advisors have hardly any more expertise in the task at hand than do the other participants (Schotter, 2003).

Studies in the second group, in contrast, show that people tend not to listen to others' opinions; that is, they tend to egocentrically discount advice (e.g., Gardner and Berry, 1995; Harvey and Fischer, 1997; Yaniv and Kleinberger, 2000; Yaniv, 2003). Egocentric advice discounting is one of the most robust findings of the psychology literature on advice-taking, which is often referred to as Judge-Advisor System (JAS) research (see Sniezek and Buckley, 1995; Sniezek and Van Swol, 2001).[2] The JAS

[2] For a comprehensive review of the literature on advice giving and advice taking, see Bonaccio and Dalal (2006).

provides a representation of the decision-making processes that occur in natural environments such as organizations or markets (Sniezek and Buckley, 1995).

The two studies presented in this paper correspond to work within the JAS literature. Thus, I describe this stream of research in more detail below.

The Judge-Advisor System Research

In a Judge-Advisor System, a judge receives advice and must decide how to incorporate it into his or her final decision; thus, the judge is the decision maker. The advice comes in the form of information or recommendations from one or more advisors who share with the judge an interest in the decision problem. Studies in the JAS literature have employed either "choice" or "judgment" tasks (Bonaccio and Dalal, 2006). In the first case, judges must choose among several alternatives that are qualitative in nature. In the second case, judges instead must provide quantitative estimates. In most studies using choice tasks, the advice is a recommendation from the advisor in favor of a particular option (e.g., Sniezek & Buckley, 1995). For instance, if the judge must choose one of several options, the advice is expressed as "Choose Option X." By contrast, in experiments employing judgment tasks, advice usually takes the form of an estimate. For instance, if the judge must estimate the year in which a certain event happened (e.g., Yaniv, 2003; 2004a; 2004b) or the weight of a person based solely on her photograph (e.g, Gino & Moore, 2006), the advice becomes another participant's estimate when facing the same task, expressed as "The Advisor's estimate is Y."

In some studies, egocentric advice discounting has been attributed to differential information (Yaniv, 2003; Yaniv and Kleinberger, 2000); according to this view, judges know their own reasons for their opinion but have no access to the advisor's internal reasons (Yaniv, 2003; Yaniv and Kleinberger, 2000). Egocentric discounting also has been attributed to anchoring (Tversky & Kahneman, 1974). In this case, the judge's initial decision might serve as an anchor. After receiving advice, the judge adjusts the anchor but, in the case of advice taking, does so insufficiently (Harvey & Fischer, 1997; Lim & O'Connor, 1995). Finally, the effect has been explained by egocentric biases (Krueger, 2003): judges might prefer their own opinions if they believe them to be superior to those of others, including advisors' recommendations.

Several variables have been found to reduce egocentric advice discounting. For instance, when judges perceive or know advisors to be knowledgeable experts, judges tend to be more responsive to the advice (Goldsmith and Fitch, 1997; Harvey and Fischer, 1997; Sniezek, Schrah, and Dalal, 2004). Egocentric discounting also diminishes when judges perceive or know advisors to be older, better educated, wiser, or more experienced than themselves (Feng and MacGeorge, *in press*). Another variable considered in the JAS literature is the presence of financial incentives or rewards. Incentives based on performance reduce egocentric advice discounting (Dalal 2001). Finally, the amount of advice discounted is also affected by the quality of the advice; the higher the quality of advice, the less advice is discounted (Yaniv and Kleinberger, 2000). Nonetheless, judges often discount good advice (Gardner and Berry, 1995).

The impact of advice cost on advice use

Both research streams described above focus primarily on whether or not people follow advice and, if they do, on how and in what directions advice changes people's behavior. In all of the studies reviewed, advice is given for *free*. But what happens when people pay for advice? Are people more likely to listen to advice when such information cost money? In other words, does the investment of a certain sum to gather advice give people a reason to use the advice, a reason they do not possess when they receive advice for free? Addressing these questions, which have not received the attention of previous work, should lead to a better understanding of advice-taking in decision making. Research on individual decision making has shown that people's economic behavior is often influenced by sunk costs (Arkes and Blumer, 1985; Garland, 1990; Heath, 1995). According to a basic economic principle, only the incremental costs and benefits of current options should affect one's decision (Arkes and Ayton, 1999). Nevertheless, evidence shows that people attend to prior investments—sunk costs—as they consider what course of action to take or what decision to make.

The same faulty logic could be at work in the case of advice taking. If people pay for advice, the sunk-cost effect would predict that they might justify such payment by using the received information in order to avoid the regret of wasting money on unused advice. The cost of advice could enter buyers' mental accounting and become salient when they are asked to make their final decisions. Thus, people might give greater weight to costly advice than free advice. Doing otherwise might appear to waste the resources already expended to acquire information about others' opinions, while heeding the advice justifies the previous action of buying advice.

Related work has found that people often experience cognitive dissonance (Festinger, 1957) upon receiving information that is inconsistent with something they believe to be true and important about themselves. In the case of advice-taking, an inconsistency may exist between the possibility that already-paid-for advice actually could be unhelpful and the belief in oneself as a rational person who does not waste money on useless things (including advice). A person might avoid or resolve this inconsistency by developing a view of the advice as more worthwhile than she would otherwise and therefore following the advice.

Prior research also has demonstrated that people often rely on advice as an indicator of quality. Although this is often a poor heuristic to use when one is choosing among different product categories (Ordonez, 1998), marketing research has shown that consumers often infer quality from price based on the belief that the two factors are positively correlated (Monroe, 1973; Monroe and Petroshius, 1981). These findings follow lay wisdom and might lead to the conclusion that *information use depends on price*. People might assign a value to the price they paid to acquire advice and thus rely on the price as a signal of quality. As a result, they may tend to "overvalue" costly advice. In the two studies presented in this paper, I tried to eliminate the viability of this explanation by explicitly telling participants that the quality of the advice received did not change based on whether the advice was free or costly. Participants were also explicitly told what the advice consisted of and who provided it.

These arguments led to the following main hypothesis: *independent of its quality, people weigh advice differently depending on whether it is costly or free.* In particular, I hypothesize that people assign a significantly greater weight to costly advice than they do to free advice.

I designed two studies to test this hypothesis. In the experiments, advice is expressed in terms of quantitative estimates (for example, the advice a participant receives might be "The Advisor's estimate is 1978"). This type of advice is often relevant in different contexts at both an individual and an organizational level. For instance, inventory managers and forecasters, consultants, or engineers use numerical estimates to advise others on a specific strategy to implement or a decision to make in the face of uncertainty.

STUDY 1

Method

Experimental sessions were conducted on computers in the computer laboratory of a university in the northeastern United States. The procedure was identical across experimental sessions, each of which lasted approximately one hour. Participants received $10 for showing up and had the opportunity to win up to an additional $22 during the experiment. On average, participants earned a total of $21. All payoffs were expressed in "experimental dollars" (points) and converted into U.S. dollars at the end of the experiment.

Participants. Sixty-three graduate and undergraduate students, 30 males and 33 females ranging in age from 19 to 24, participated in this experiment. They were recruited via ads that offered money in exchange for participating in an experiment. To avoid having experts in the lab, history majors could not participate. Thirty-five students participated in Condition 1 and 28 students participated in Condition 2.

Procedure. The study was described to participants as an experiment on individual decision making. Upon arriving at the laboratory, participants were registered and randomly assigned to one of two experimental conditions. Participants received a copy of the instructions explaining the experiment and use of the computer; the researcher also read the instructions aloud and gave participants an opportunity to ask questions. Before being paid and leaving the lab, participants were asked to answer a questionnaire with demographic questions.

The experiment consisted of four phases. In each phase, participants were asked to estimate questions about the dates of specific events in American history (within the last 400 years) and to provide their estimates privately without communicating with other participants. In addition to guessing the correct year of a given event, participants were asked to provide lower and upper estimates of their 90% confidence interval.

In two of the four phases participants had the opportunity to choose whether to receive advice. They were told, accurately, that the advice would be another student's answer to the same question. Participants did not receive feedback after answering each question; hence, they had no opportunity to learn anything about advisors and their expertise. Thus, they could not weigh advice differently depending on the advisors' reputation. Moreover, as participants were explicitly told, advisors were randomly chosen by the researcher on a question-by-question basis. The lack of performance feedback and varied source of advice presumably made it very hard for participants – if not impossible – to assess the quality of the advice they received.

The experimental design of Study 1 is shown in Table 1. Next, I provide details of each phase in each of the two conditions.

Condition 1. In Phase 1, participants were asked to answer 15 questions about American history. In Phase 2, participants were asked to answer the same series of questions; this time, however, they could choose to obtain "advice" on the correct answer. For each question, the advisors estimate came from a pool of estimates collected in an earlier study in which respondents provided a best estimate for each question. The researcher randomly selected advisors and their estimates prior to the study; the person chosen as the advisor might have varied from one question to the next and across participants. In Phase 2, the advice was free; before the start of Phase 2, the researcher asked participants once whether they wanted it or not. If they accepted, they received advice for each of the 15 questions in the phase.

In Phase 3, the researcher asked participants to answer a different set of questions about American history; the rules were exactly the same in Phase 1. As in Phase 2, participants could choose to solicit advice in Phase 4. However, in Phase 4, participants had to pay for the advice. If they accepted the offer of costly advice, they had to pay 10,000 experimental dollars (equivalent to $4) in exchange for advice on all 15 questions. The fee was deducted from their final payoff, whether or not they followed the advice.

Thus, participants had to make the decision of whether or not to receive advice only twice: first, before the beginning of Phase 2; second, before the beginning of Phase 4; they were asked to receive advice not on a question-by-question basis but on a set of questions. This procedure was intended to create a feeling of commitment for those who chose to receive advice in the costly-advice condition.

Condition 2. Condition 2 was similar to condition 1, but in this case, to test for order effects, advice was costly in Phase 2 and free in Phase 4.

Participants' payoff. A participant's payoff was computed as the sum of the show-up fee plus bonuses based on the accuracy of point estimates in each phase. In phases 1 and 2, participants received 1,500 points in Condition 1 and 3,000 points in Condition 2 as a bonus every time their point estimate fell within a range of minus or plus 10 years surrounding the year the event took place. In Phases 3 and 4, participants received 3,000 points in Condition 1 and 1,500 points in Condition 2 as a bonus every time their point estimate fell within the same range of minus or plus 10 years surrounding the year the event took place.

Dependent measure. To measure the degree to which participants used the advice they received, I use the "weight of advice" (WOA) measure, which has been previously employed in several studies (see Harvey and Fisher, 1997, and Yaniv, 2003, in the context of advice taking; and Hell, Gigerenzer, Gauggel, Mall, and Muller, 1988, in the context of memory). WOA is defined as follows: $WOA = \frac{|final\ estimate - initial\ estimate|}{|advice - initial\ estimate|}$, in which the final estimate was the estimate provided in Phase 2 and 4 when advice was available and the initial estimate was the point estimate provided in Phase 1 and 3.

This measure is well defined if the final estimate falls between the initial estimate and the advice (Yaniv, 2003). In Study 1, this happened in 97% of the cases. With such a measure, the weights assigned to the initial estimate and to the advice are proportional to the shift of participants' point estimates toward or away from the advice in two subsequent phases (either Phase 1 and 2 or Phase 3 and 4).

The weight of advice reflects how much a subject uses the advice she receives; thus, it is inversely related to how much a participant discounts the advice (Yaniv, 2003). The WOA equals 0 when a participant entirely discounts the advice. In such a case, the final estimate equals the initial one, meaning that the participant did not change her initial decision after receiving the advice. The opposite is true when a participant's initial estimate shifts completely toward the advice, meaning that the advice was followed. In this case, WOA equals 1, since the final estimate equals the received advice. Finally, when WOA equals a value between 0 and 1, partial discounting results from the fact that the participant weighs both her initial estimate and the received advice positively. For instance, a WOA of 0.5 means that the participant averaged the advice with her initial estimate. On average, this is the optimal strategy people should use when making their final estimate (Soll & Larrick, 2004).

The WOA measure has a few drawbacks. First, when the advice equals the judge's initial estimate, the denominator is zero; thus, the WOA is an undefined value. In these cases, it is not possible to quantify how much the participant did or did not use the advice. Second, when the final and initial estimates are equal, the WOA equals 0 and is thus interpreted as 100% advice discounting. Yet, in these cases, the denominator might have been a very small number, meaning that the advice served as a confirmation that the initial estimate was a good proxy for the correct answer and therefore should be maintained. Third, the WOA measure does not allow one to distinguish cases in which the final estimate moves *toward* the advice from cases in which it moves *away from* the advice. Finally, the WOA measure has a lower bound of zero but not an upper bound; thus, it can be a number greater than 1, though this happens rarely (e.g., Harvey & Fischer, 1997).

Interval estimates. We can reasonably assume that the width of the interval estimates reflects participants' assessments of their own knowledge, their confidence in the answers they provided as their best guesses and, at least partially, the accuracy of their answers (Yaniv, 2003; Yaniv and Foster, 1995; 1997).

Results

The results of Study 1 are summarized in Table 2.

Distribution of the values for WOA. In the analysis of the WOA values, I left out cases in which the advice equals the initial estimate, since WOA in those cases equals a number divided by 0. As for cases in which the WOA is not well-defined (that is, those in which the final estimate does not fall between the initial estimate and the advice, and WOA is thus greater than 1), I chose to change values above 1 to 1. In particular, I changed values above 1 to 1 in 1% of the cases in the costly advice treatment and in 4% of the cases in the free-advice treatment.[3]

To analyze the distribution of WOAs, I rounded each value of WOA to the nearest decimal, then assigned them to one of the following three groups: low (0 - .3), medium (.4 - .6), and high (.7 - 1). There were 1,380 individual trials (63 participants x 30 questions, minus the trials of participants who opted not to receive advice). The percentage of low

[3] Another possibility would be to leave out cases in which WOA is not well defined. However, as long as the number of such cases is small, the method used should make little difference. The data of both Study 1 and Study 2 were analyzed using each of the two methods. The nature and significance of the results did not change.

WOAs was higher in the free-advice treatment (39.5%) than in the costly advice treatment (22.4%), χ^2 (1, N = 413) = 118.259, p<.0001. By contrast, the percentage of high WOAs was higher in the costly advice treatment (72.2%) than in the free-advice treatment (52.4%), χ^2 (1, N = 730) = 17.184, p<.0001. These results are consistent with the hypothesis that individuals tend to weigh advice more when they pay for it than when get it for free. The percentage of medium WOAs equaled 8.1% in the free-advice treatment and 5.4% in the costly advice treatment; this difference is statistically significant, χ^2 (1, N = 88) = 20.045, p<.0001.

Buying rate. In Condition 1, 34 participants out of 35 (97%) opted to receive advice in the free-advice treatment, and 15 (43%) bought advice in the costly-advice treatment. In Condition 2, 26 participants out of 28 (93%) opted to receive advice in the free-advice treatment, and 17 (61%) bought advice in the costly-advice treatment. It appeared that people who received the offer of free advice before costly advice bought advice at a lower rate than did people who were offered costly advice before free advice. Yet testing of whether sequence affected the buying rate found no effect, F(1,61)=.818, p=.369, η^2=.013.

I conducted further analyses to explore whether participants who bought advice differed from those who did not. From such analyses, I cannot conclude that people who paid for the advice were less confident in the answers they provided than were people who chose not to buy the advice. Moreover, there is no evidence that mean error is a good predictor of participants' willingness to buy advice in the next phase.

Impact of advice cost. As predicted, the cost of advice had a significant effect on the weight participants assigned to the advice itself. At the aggregate level, the mean WOA in the costly advice treatment was 0.74 (SD = 0.06). By contrast, in the free-advice treatment, WOA was only 0.56 (SD = 0.07). This difference is statistically significant, $t(58)=-5.811$, $p<.0001$, thus supporting my main hypothesis.

Table 3 reports the values for WOA, pooled by both condition and treatment. I analyzed the results for the impact of advice, presented in this table, by looking at the *mean WOA per question* both across conditions and across phases. In the analyses, I compared the weight that participants assigned to the advice when it was free and when it was costly. WOA was higher in the costly-advice treatment than in the free-advice treatment. The difference is statistically significant, both if one looks at the results "across phases" ($t[28]=-8.074$, $p<.0001$, for Phases 1 and 2, that is, Condition 1-Free vs. Condition 2-Costly; $t[28]=-2.096$, $p=.045$, for Phase 3 and 4, that is, Condition 1-Costly vs. Condition 2-Free) and if one analyzes them "across conditions" ($t[28]=-2.563$, $p=.016$, Free vs. Costly in Condition 1; $t[28]=-7.014$, $p<.0001$, Free vs. Costly in Condition 2). Note that the analysis conducted across phases holds constant the content of the questions, something that the within-subjects comparisons do not do.

Core test of main hypothesis. In the following analysis, I focus only on the participants who bought advice, in each condition. In other words, I analyze the data "controlling for people," considering the behavior of participants who chose to pay for costly advice and looking at their behavior when the advice was free. Computed values for WOA, pooled by condition, are reported in Table 4. As the table shows, in each condition, WOA is higher in the costly-advice treatment than in the free-advice treatment. A 2x2 ANOVA with advice (free vs. costly) as a within-subjects factor and condition (1 vs. 2) as a between-subjects factor reveals that the difference is statistically significant in both Condition 1 ($F(1, 28)=-8.909$, $p=.006$) and Condition 2 ($F(1, 28)= 24.438$, $p<.0001$). These results are the core test of my initial hypothesis.

Interval estimates. Participants might not feel equally confident about their answers to questions belonging to two different sets (15 questions asked in Phase 1 and 2 vs. 15 different questions asked in Phase 3 and 4). I checked whether this was the case by comparing the mean width of the interval estimates in Phase 1 and in Phase 3 and found no significant difference, $t(58)=-.676$, *ns*.

An alternative explanation: Do people buy advice based on their propensity to use it? The previous analyses have shown that advice use is influenced by advice cost. One might wonder whether the results described above are driven by a selection mechanism. Specifically, people who pay for advice may value it more than others because they are less certain of their initial estimates. Support for this alternative explanation comes from prior research showing that people's uncertainty about their initial decision is a good predictor of their advice-seeking behavior (Cooper, 1991; Gibbons, Sniezek & Dalal, 2003). A mixed model analysis was used to isolate the effect of paying for advice from the effect of how much participants might value that advice *ex ante*. In a mixed model analysis, the studied sample of participants is treated as a random selection from the general population; participants are thus included in the model as random effects. I estimated the following model:

$$WOA_{ij} = \alpha_0 + \beta_1 PAY_i + \beta_2 IE_{ij} + subjects'\ RANDOM\ \ EFFECTS + \varepsilon\ ,$$

where index *i* referred to participants and index *j* referred to questions. The dependent variable was the value for WOA for each subject and for each question. Explanatory variables were: (i) a dummy variable indicating whether the subject got advice for free or by paying for it (PAY_i), and (ii) the size of the initial range (IE_{ij}). The interval estimate was used as a proxy for how much a participant might have valued advice.

This analysis allows me to exploit an important feature of the experimental design: participants had to commit to paying for advice on all questions up front, but got to choose the degree to which they used that advice on a question-by-question basis.

Results are based on a total of 1,230 observations, each a question answered by a subject. The results reveal a significant and positive effect of the initial range size on WOA (β_2=0.001, *t*=7.090, p < .0001): as one might expect, the wider the initial range, the higher the WOA. It appears that participants do weigh advice more heavily when they are less confident of their own knowledge. However, even after I controlled for this effect, the effect of paying for advice is still in the expected direction and statistically significant (β_1=0.184, *t*=3.544, p < .0001).

Hence, the results of the mixed model analysis show that when participants pay for advice they weigh it more than one would expect given their confidence in their answers. As expected, WOA increases as the initial range increases; holding the initial range constant, however, the WOA is significantly higher when the advice is costly than when it is free.

Order effects. As described above, people who paid for the advice before receiving it for free used it more than did those who paid for it after previously receiving it for free. To test whether the order in which free vs. costly advice was received had an impact on WOA, I conducted an ANOVA in which participants' values for WOA served as dependent variable, pay (that is, free vs. costly advice) served as a within-subjects factor, and sequence (that is, Condition 1 vs. Condition 2) served as a between-subjects factor. The main effect for pay replicates the finding of the previous analyses: Participants tend to weigh costly advice significantly more than free advice, $F(1,28)=37.994$, $p<.0001$, $\eta^2=.576$. Results also reveal a main effect for sequence, $F(1,28)=11.134$, p=.002,

η^2 =.285, suggesting that participants who received the advice for free first were less prone to using the advice when they paid for it than were participants who were first offered the advice at a cost. The act of perceiving the advice as costly might have led participants in Condition 2 to implicitly assign more “expert” status to that information. If so, this data provides another example of the tendency of people to implicitly value things more when they cost more, in line with the marketing literature cited in the introduction. Yet, the analyses conducted in Study 2 do not confirm the order effects results discussed above.

Discussion

The results of Study 1 support the initial hypothesis: Participants use advice significantly more when it is costly than when it is free. Even after taking into account differences in confidence that might influence willingness to pay for advice, I find that advice is weighed more heavily when it is acquired at a cost. In addition, the results show that participants who received the advice for free first were less prone to overweight the advice when they paid for it than were participants who were first offered the advice at a cost.

The findings of the first study suggest two other facets of the paid-advice effect that should be further explored. The first concerns the buying rate in the costly-advice treatment: only about half of the participants bought the advice when they were given the option not to buy it in the costly-advice treatment. There might have been a difference between buyers and non-buyers that was not detected by Study 1. The results I presented help to address this concern, but do not eliminate it completely. The second facet concerns participants’ payoffs: the size of the bonus paid for the correct answers was indeed confounded with whether the advice was free or costly. Participants were asked to pay for advice when the bonus for the point estimate was 3,000 points per question. The advice was free when the bonus for the point estimate was 1,500 points per question. Thus, participants might have been more strongly influenced by the advice in the costly-advice treatment than in the free-advice treatment not because they paid for the advice but because they had more to gain if the advice helped them get the answers correct.

A second study was designed to explore these issues. In the traditional JAS, advice is imposed on decision makers. In other words, judges do not have the opportunity to decide whether or not to receive advice, but receive it by default. This feature was employed in Study 2. Thus, in this second experiment, I eliminate the endogeneity problem encountered in Study 1 at the expense of the choice of getting advice. As for the concern related to participants' payoff, in Study 2 the bonus based on the accuracy of point estimates did not differ in the free-advice and costly-advice treatments.

STUDY 2

Method

Experimental sessions were conducted in the computer laboratory of a northeastern U.S. university. The procedure was identical across experimental sessions, and each session was conducted on computer. Participants received $10 for showing up and also had the opportunity to win up to $25 during the experiment.

Participants. Participants were recruited via ads that offered money in exchange for participating in an experiment. On average, each experimental session lasted one hour and 15 minutes. Eighty-eight people agreed to participate (49% male and 51% female). The average age of participants was 26 (SD = 7). Most participants (93%) were students. There were 44 participants in each condition.

Procedure. Study 2 followed the same procedure used in Study 1 with only one difference. While in Study 1, participants were given the option of receiving advice (free or costly), in Study 2 they did not have this option; instead, advice was given *by default*. At the beginning of the experiment, participants were told that they would be randomly assigned to one of two conditions. If assigned to Condition 1, they would receive advice for free in Phase 2 and costly advice in Phase 4. If assigned to Condition 2, they would receive costly advice in Phase 2 and free advice in Phase 4.

Participants were explicitly told that the advice was given by various advisors. Specifically, the advice came from another participant's estimate when asked the same question in the corresponding previous phase. The advisor might have varied from one question to the next and was randomly selected. In the costly advice treatment, the advice cost the equivalent of $8. Participants were told that they would be charged this amount of money when receiving costly advice.

Participants' payoff. Participants' payoffs were computed as the sum of the show-up fee and some bonuses. As in Study 1, participants received bonuses based on the accuracy of their point estimates. They received $ .45 for each question in which their best guess fell within a range of 10 years minus or plus the true date. Participants were paid in cash at the end of the experiment according to what they had earned.

Results

The results of Study 2 are summarized in Table 5.

Distribution of the WOA values. As in my Study 1 analyses, I omitted cases in which advice equaled a participant's initial estimate. When the WOA was greater than 1, I changed values to 1. Specifically, I changed values above 1 to 1 in 1% of cases in the costly advice treatment and in 2% of cases in the free-advice treatment.

In the analysis of the distribution of WOA values, I compared the percentage of low, medium, and high WOAs in the two advice treatments (free and costly). The percentage of low WOAs was higher in the free-advice treatment (52%) than in the costly advice treatment (32%), χ^2 (1, *N* = 1010) = 57.984, p<.0001. As for the percentage of high WOAs, it was higher in the costly advice treatment (57%) than in the free-advice treatment (37%), χ^2 (1, *N* = 1124) = 52.968, p<.0001. These results are consistent with the hypothesis that people weight advice more heavily when it is costly than when it is free.

The impact of advice cost. My analyses for Study 1 were similar for those of Study 2. Overall, the results of Study 2 show further support for my initial hypothesis. At the aggregate level, the mean WOA equaled 0.62 (SD = 0.03) in the costly advice treatment and 0.42 (SD = 0.09) in the free-advice treatment. This difference is statistically significant, *t*(58)=-10.564, p<.0001.

Table 6 reports the WOA values, pooled by both condition and treatment. As in Study 1, I analyzed these results by looking at the *mean WOA per question* both across conditions and across phases. WOA in the costly advice treatment was higher than in the free-advice treatment. The difference is statistically significant, as shown by analyzing the data both "across phases" (*t*[28]=-5.476, p<.0001, Condition 1-Free vs. Condition 2-Costly; *t*[28]=-14.171, p<.0001, Condition 1-Costly vs. Condition 2-Free) and "across conditions" (*t*[28]=-6.258, p<.0001, Free vs. Costly in Condition 1; *t*[28]=-13.636, p<.0001, Free vs. Costly in Condition 2).

Order effects. In Study 2, participants who received the advice for free before paying for it (Condition 1) used it more than did those who received it for free after previously paying for it (Condition 2). To test for an order effect, I conducted an ANOVA in which participants' WOA values served as the dependent variable, pay (free vs. costly advice) served as a within-subjects factor, and sequence (Condition 1 vs. Condition 2) served as a between-subjects factor. The main effect for pay replicates the finding of the previous analyses: Participants tend to use costly advice significantly more than free advice, $F(1,28)=251.948$, $p<.0001$. Results also reveal a main effect for sequence, $F(1,28)=18.748$, $p<.0001$, suggesting that participants who received the advice for free before paying for it were more prone to use the free advice than were participants who were first given the advice at a cost. The impact of this sequence in Study 2 (differential use of *free* advice in the two conditions) does not confirm the order effects found in Study 1 (differential use of *costly* advice in the two conditions).

Interval estimates. I conducted further analyses to compare the confidence of participants in their answers in the free-advice treatment vs. in the costly advice treatment. The width of participants' interval estimates in the phase in which advice was not available provides a measure of participants' confidence. For each condition, I used a repeated-measures ANOVA in which participants' interval estimates served as the dependent variable and advice (free vs. costly) served as a within-subjects factor. Results reveal a significant main effect for advice only for Condition 1; participants felt more confident about their answers in the costly advice condition (M=26.38 years, SD=15.65) than in the free-advice condition (M=30.89, SD=14.19), $F(1,43)=3.511$, $p=.068$, $ή^2 = .075$.[4]

[4] As for condition 2, no effect was found, $F(1,43)=.696$, $p=.409$, $ή^2 = .016$. Moreover, in condition 2, the mean width of participants' interval estimates is higher in the free advice treatment (M=32.31, SD=11.28) than in the costly advice treatment (M=26.04, SD=9.23).

Discussion

The results of Study 2 are consistent with my initial hypothesis and thus provide further support for the findings from Study 1. Specifically, people weigh advice more heavily when the advice is costly than when it is free. In Study 2, as in several other studies from the JAS literature, advice was imposed on people, but they were able to choose the degree to which they followed the advice.

GENERAL DISCUSSION AND CONCLUSIONS

Most of the decisions people make on a daily basis result from weighing their own opinions with advice from other sources. The present work explored one factor that might affect the use of advice: advice cost. In particular, the initial hypothesis was that, independent of its quality, people would weigh advice significantly more when it is costly than when it is free. This hypothesis was tested in two experiments requiring participants to answer questions about American history with or without advice from others. The results of the two studies show that participants rely more heavily on advice when it is costly than when it is free.

The cost of advice affected the degree to which participants used advice but did not affect the value gained by following advice. In both studies, advice came from another participant who was randomly chosen on a question-by-question basis. On average, advisors were as informed or knowledgeable as judges. In fact, in at least a few rounds, the same participant likely played the role of advisor and judge at the same time. Moreover, participants had no opportunity to assess the accuracy of advisors' estimates. Nor had they the opportunity to assess the accuracy of their own estimates, since no performance feedback was provided. When advice was costly, participants behaved as if their own opinions were worth less than those of a randomly selected other. When advice was free, they instead behaved as if their estimates were more accurate than those of a randomly selected advisor.

The two experiments do not study the mechanism underlying this paid-advice effect. Yet one might expect this bias to occur for different reasons. I suggested earlier that when decision makers buy advice, they fall prey to the sunk-cost fallacy. Prior investments should not influence participants' evaluation of current options; instead, they should only consider incremental costs and benefits (Arkes and Ayton, 1999). In the experiments, the cost incurred to acquire the advice (either by choice, as in Study 1, or by default, as in Study 2) became salient to participants when I asked them to provide their final estimates and they assigned the advice a greater weight than when they received it for free. In the case of advice taking, the sunk-cost effect might be even stronger than in other contexts, since people might feel greater regret about not using information provided by others than about not using a product or service commonly available.

Research on cognitive dissonance (Festinger, 1957) might also predict the paid-advice effect. A person experiences cognitive dissonance when there is an inconsistency between the information one receives and what one believes to be true and important about oneself. In the experiments, an inconsistency may have existed between the possibility that already-paid-for advice actually could be unhelpful and the belief in oneself as a rational person who does not waste money on useless advice. People might avoid or resolve this inconsistency by developing a view of the advice as more worthwhile than they otherwise would. As a result, they would follow the advice.

Both of these explanations implicitly assert that the paid-advice effect occurs because participants overweigh the advisor's response relative to their own during the information-processing stage of decision making. However, the effect may also occur at the encoding stage. The price attribute that people attach to the advice might signal quality (or advisor expertise), which, in turn, affects how people automatically receive and encode the information into working memory. The two studies were designed to eliminate the viability of this explanation, but may have not ruled it out completely. Thus, the mechanism that is actually in place remains a question to be addressed by further studies.

My findings suggest several other areas of future research. First, a better understanding of decision makers' sensitivity to the cost incurred to acquire advice is needed. Either increasing or decreasing the cost of advice would enable an exploration of the extent to which the "size" of cost matters. An awareness that advice cost affects advice use might be of interest to the consultant or medical professions. In such fields, advisors must decide whether or not to charge their "clients" for the information and opinions they provide, as well as how much to charge. The type of advice that experts provide suggests another direction for future research. In the present study, I used numerical estimates as advice—that is, "quantitative" advice. Would the paid-advice effect occur in the case of qualitative advice? Researchers could also explore the paid-advice effect empirically in settings such as consulting, health care, or product forecasting. Finally, future research should investigate if the paid-effect effect holds when the conceptualization of advice is broader. Following most of the studies in the JAS literature, I defined advice in this paper at the operational level rather than at the construct level. Specifically, I modeled advice as a recommendation from a randomly selected advisor, expressed as the advisor's response to the same task. Research on advice-taking has not yet adequately defined the term "advice" (Bonaccio and Dalal, 2006). Once the definition of advice is broadened, a closer exploration of the finding that people use costly advice more heavily than free advice would be a fruitful endeavor.

REFERENCES

Anderson, L. R., & Holt, C. A. (1997). Information Cascades in the Laboratory. *The American Economic Review*, **87** (5), 847-862.

Anderson, L. R., & Holt, C. A. (1996). Classroom Games. Information Cascades. *Journal of Economic Perspectives*, **10** (4), 187-193.

Arkes, H. R., & Ayton, P. (1999). The sunk cost and concorde effects: are humans less rational than lower animals? *Pschological Bulletin*, **125** (5), 591-600.

Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational Behavior and Human Decision Processes*, **35**, 124-140.

Banerjee, A.V. (1992). A simple model of heard behavior. *Quarterly Journal of Economics*, **107** (3), 797-817.

Bastardi, A., & Shafir, E. (1998). On the pursuit and misuse of useless information. *Journal of Personality and Social Psychology*, **75** (1), 19-32.

Bogachan, C., Shachar, K., & Schotter, A. (2002). The advice puzzle: an experimental study of social learning where words speak louder than actions. Mimeo, New York University.

Bonaccio, S., & Dalal, R.S. (2006). Advice taking and advice giving in decision making: an integrative review of the literature. Working paper. Purdue University.

Chandhuri, A., Schotter, A., & Sopher, B. (2002). Talking ourselves to efficiency: coordination in inter-generational minimum games with private, almost common, and common knowledge of advice. Mimeo, New York University (submitted to *American Economic Review*).

Cooper, R.S. (1991). Information processing in the judge-adviser system of group decisionmaking. Unpublished master's thesis, University of Illinois, Urbana-Champaign.

Dalal, R.S., (2001). The effect of expert advice and financial incentives on cooperation. Unpublished master's thesis, University of Illinois, Urbana-Champaign.

Feng, B., & MacGeorge, E. L. (2005). (in press). Predicting responsiveness to advice: Characteristics of the problem, the advice-giver, and the recipient. *Southern Communication Journal*.

Festinger, L. (1957). *A Theory of Cognitive Dissonance*. Evanston, IL: Row, Perterson & Company.

Gardner, P.H., & Berry, D.C. (1995). The effect of different forms of advice on the control of a simulated complex system. *Applied Cognitive Psychology*, **9**, S55-S79.

Garland, H. (1990). Throwing good money after bad: the effect of sunk costs on the decision to escalate commitment to ongoing project. *Journal of Applied Psychology*, **75** (6), 728-731.

Gibbons, A.M., Sniezek, J. A., & Dalal, R. S. (2003, November). Antecedents and consequences of unsolicited versus explicitly solicited advice. In D. Budescu (Chair), Symposium in Honor of Janet Sniezek. Symposium presented at the annual meeting of the Society for Judgment and Decision Making, Vancouver, BC.

Gino, F. & Moore, D.A. (2006). When Do People Listen to Advice? Effects of Task Difficulty on Use of Advice. under review.

Goldsmith, D.J., & Fitch, K. (1997). The normative context of advice as social support. *Human Communication Research,* **23**, 454-476.

Harvey, N., & Fischer, I. (1997). Taking advice: Accepting help, improving judgment and sharing responsibility. *Organizational Behavior and Human Decision Processes*, **70**, 117-133.

Heath, C. (1995). Escalation and de-escalation of commitment to sunk costs: The role of budgeting in mental accounting. *Organizational Behavior and Human Decision Processes*, **62**, 38-54.

Krueger, J. L. (2003). Return of the ego—self-referent information as a filter for social prediction: Comment on Karniol (2003). *Psychological Review,* 110, 585-590.

Iyengar, R., & Schotter, A. (2002). Learning with a meddlesome boss: an experiment in sloppy principal-agent relationships. Mimeo, New York University (submitted to *Quarterly Journal of Economics*).

Lim, J.S., & O'Connor, M. (1995). Judgmental adjustment of initial forecasts: Its effectiveness and biases. *Journal of Behavioral Decision-Making*, 8, 149-168.

Merlo, A., & Schotter, A. (1999). A surprise quiz view of learning in economic experiments. *Games and Economic Behavior*, **28** (1), 25-54.

Merlo, A., & Schotter, A. (2003). Learning by not doing: an experimental investigation of observational learning. *Games and Economic Behavior*, **42** (1), 116-136.

Monroe, K. B. (1973). Buyer's subjective perceptions of price. *Journal of Marketing Research*, **10**, 70-80.

Monroe, K. B., & Petroshious, S. M. (1980). Buyer's perceptions of price: An update of the evidence. In T. Robertson & H. Kassarjian (Eds.), *Perspectives in Consumer Behavior* (pp. 43-55). Glenview, Scott: Foresman and Company.

Ordonez, L. D. (1998). The Effect of Correlation between Price and Quality on Consumer Choice. *Organizational Behavior and Human Decision Processes*, **75** (3), 258-273.

Savage, L. J. (1971). Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, **66** (336), 783-801.

Schotter, A. (2003). Decision Making with naïve advice. *American Economic Review*, **May**, 196-201.

Schotter, A., & Sopher, B. (2003). Social learning and convention creation in inter-generational games: an experimental study. *Journal of Political Economy*, **June**.

Sniezek, J.A., & Buckley, T. (1995). Cueing and cognitive conflict in judge-advisor decision making. *Organizational Behavior and Human Decision Processes,* **62,** 159-174.

Sniezek, J.A., Paese, P.W., & Swithers, F.S. (1990). The effect of choosing on confidence in choice. *Organizational Behavior and Human Decision Processes,* **46***,* 264-282.

Sniezek, J.A., & Van Swol, L.M. (2001). Trust, confidence, and expertise in a judge-advisor system. *Organizational Behavior and Human Decision Processes,* **84***,* 288-307.

Sniezek, J.A., Schrah, G.E., & Dalal, R.S. (2004). Improving judgment with prepaid expert advice. *Journal of Behavioral Decision Making,* **17***,* 173-190.

Soll Jack B. & Larrick Richard P. (2004) Strategies for Revising Judgment: How, and How Well, Do People Use Others' Opinions? Status: Current draft is under revision for resubmission.

Thaler, R. H. (1980). Towards a positive theory of consumer choice. *Journal of Economic Behavior and Organization*, **1**, 39-60.

Tversky, A., & Kahneman, D. (1974). Judgment under Uncertainty: Heuristics and Biases. *Science*, **185**, 1124-1131.

Winkler, R. L. (1969). Scoring rules and the evaluation of probability. *American Statistical Association Journal*, **September**, 1073-1078.

Yaniv, I., & Foster, D.P. (1995). Graininess of judgment under uncertainty: an accuracy-informativeness tradeoff. *Journal of Experimental Psychology: General*, **124**, 424-432.

Yaniv, I., & Foster, D.P. (1997). Precision and accuracy of judgmental estimation. *Journal of Behavioral Decision Making*, **10**, 21-32.

Yaniv, I., & Kleinberger, E. (2000). Advice taking in decision making: Egocentric discounting and reputation formation. *Organizational Behavior and Human Decision Processes*, **83** (2), 260-281.

Yaniv, I. (2004a). The benefit of additional opinions. *Current Directions in Psychological Science, 13,* 75-78.

Yaniv, I. (2004b). Receiving other people's advice: Influence and benefit. *Organizational Behavior and Human Decision Processes, 93*, 1-13.

Yaniv, I. (2003). Receiving other people's advice: Influence and benefit", *Organizational Behavior and Human Decision Processes*, **92**, 113-125.

Yaniv, I. (1997). Weighting and trimming: heuristics for aggregating judgments under uncertainty. *Organizational Behavior and Human Decision Processes*, **69** (3), 237-249.

Table 1. Experimental design.

	Condition 1	*Condition 2*
Phase 1	15 questions about American history	15 questions about American history (same questions as in Condition 1)
Phase 2	15 questions about American history (same questions as in Phase 1) together with **FREE advice** if wanted	15 questions about American history (same questions as in Phase 1) together with **COSTLY advice** if wanted
Phase 3	Other 15 questions about American history	Other 15 questions about American history (same questions as in Condition 1)
Phase 4	15 questions about American history (same questions as in Phase 3) together with **COSTLY advice** if wanted	15 questions about American history (same questions as in Phase 3) together with **FREE advice** if wanted

Table 2. Summary of results for Study 1.

	Total	Accepted free advice	Accepted costly advice
Condition 1			
N	35	34 (97%)	15 (43%)
Mean WOA		0.56	0.66
% Low		175 (39%)	58 (31%)
% Medium		43 (10%)	10 (5%)
% High		228 (51%)	119 (64%)
Condition 2			
N	28	26 (93%)	17 (61%)
Mean WOA		0.57	0.81
% Low		142 (40%)	38 (16%)
% Medium		22 (6%)	13 (5%)
% High		193 (54%)	190 (79%)
Combined			
N	63	60 (95%)	32 (51%)
Mean WOA		0.56	0.75
% Low	413	317 (39%)	96 (22%)
% Medium	88	65 (8%)	23 (5%)
% High	730	421 (52%)	309 (72%)

Table 3. WOA values pooled by condition and treatment (Study 1).

	Condition 1			*Condition 2*		
	Mean	**Std Dev**	**Median**	**Mean**	**Std Dev**	**Median**
Free Advice	0.56	0.09	0.58	0.57	0.11	0.56
Costly Advice	0.66	0.12	0.67	0.81	0.08	0.79

Table 4. WOA for participants who bought the advice, pooled by condition (Study 1).

	Condition 1			*Condition 2*		
	Mean	**Std Dev**	**Median**	**Mean**	**Std Dev**	**Median**
Free Advice	0.54	0.11	0.55	0.62	0.08	0.63
Costly Advice	0.66	0.12	0.66	0.81	0.08	0.79

Table 5. Summary of results for Study 2.

	Total	Free advice treatment	Costly advice treatment
Condition 1			
N	44		
Mean WOA		0.48	0.63
% Low		287 (46%)	194 (31%)
% Medium		80 (13%)	62 (10%)
% High		257 (41%)	362 (59%)
Condition 2			
N	44		
Mean WOA		0.36	0.61
% Low		339 (59%)	190 (33%)
% Medium		57 (10%)	70 (12%)
% High		183 (32%)	322 (55%)
Combined			
N	88		
Mean WOA		0.42	0.62
% Low	413	626 (52%)	384 (32%)
% Medium	88	137 (11%)	132 (11%)
% High	730	440 (37%)	684 (57%)

Table 6. WOA values pooled by condition and treatment (Study 2).

	Condition 1			*Condition 2*		
	Mean	**Std Dev**	**Median**	**Mean**	**Std Dev**	**Median**
Free Advice	0.48	0.09	0.52	0.36	0.07	0.36
Costly Advice	0.63	0.04	0.62	0.61	0.03	0.61