Organizational Behavior and Human Decision Processes 107 (2008) 234–245

Contents lists available at ScienceDirect

Organizational Behavior and Human Decision Processes

journal homepage: www.elsevier.com/locate/obhdp

Do we listen to advice just because we paid for it? The impact of advice cost on its use

Francesca Gino

Carnegie Mellon University, CMU/Tepper, 5000 Forbes Avenue, Pittsburgh, PA 15217, USA

A R T I C L E I N F O

Article history:
Received 2 April 2007
Accepted 14 March 2008
Available online 25 April 2008

Accepted by William Bottom

Keywords:
Advice taking
Judge–Advisor System
Sunk costs
Paid-advice effect
Information use

A B S T R A C T

When facing a decision, people often rely on advice received from others. Previous studies have shown that people tend to discount others' opinions. Yet, such discounting varies according to several factors. This paper isolates one of these factors: the cost of advice. Specifically, three experiments investigate whether the cost of advice, independent of its quality, affects how people use advice. The studies use the Judge–Advisor System (JAS) to investigate whether people value advice from others more when it costs money than when it is free, and examine the psychological processes that could account for this effect. The results show that people use paid advice significantly more than free advice and suggest that this effect is due to the same forces that have been documented in the literature to explain the sunk costs fallacy. Implications for circumstances under which people value others' opinions are discussed.

Introduction

> *If it's free, it's advice; if you pay for it, it's counseling; if you can use either one, it's a miracle.* Jack Adams (1838–1918)

When facing a decision, people often consult others for their opinion before making a final commitment. How much do people value others' advice? Previous laboratory studies have generally found that people give more weight to their own opinions than those of others, although the appropriate use of advice leads to better judgments (Gardner & Berry, 1995; Harvey & Fischer, 1997; Yaniv & Kleinberger, 2000). This result holds even when people have the same information and knowledge as their advisors (Harvey & Fischer, 1997).

While these laboratory studies provide evidence suggesting that people tend to discount advice from others, data collected outside the lab shows that, in fact, there are circumstances in which people do listen to others' opinions. For instance, individuals pay high fees to professional therapists (Prince, 2005) even if research has shown that their recommendations are not that helpful (Dawes, 1994). Furthermore, people pay a great deal of money to professional investors (Bogle, 1999) despite the fact that movements in the stock market rarely can be predicted (Malkiel, 2003). Similarly, shareholders invest in mutual funds based on the professional investment advice they buy from fund managers (Freeman & Brown, 2001), when their investments would be more profitable if they simply purchased broad stock indexes that track the overall market (Bazerman, 2001). Firms, for their part, pay high fees to management consultants in exchange for advice in solving their business problems, despite a shortage of evidence that consultants' advice provides real value (Micklethwait & Wooldridge, 1996; Pfeffer & Sutton, 2006).

In laboratory studies that have shown advice discounting, participants generally received advice for free (e.g., Gino & Moore, 2007; Yaniv, 2004). By contrast, in real-world settings, individuals and organizations pay substantial amounts of money for expert recommendations. This paper explores the possibility that, even if one holds constant the quality of the advice people receive under controlled, laboratory conditions, the cost of advice moderates how much people value it. In three studies, participants were asked to answer different sets of questions about US history. Before answering some of the questions, they received "advice" on the correct answer. As participants were told, the advice came from another (randomly selected) participant's best estimate when asked the same question. 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 paid advice by default. The third study investigated potential mechanisms behind the paid-advice effect demonstrated in the first two studies.

The results of the three studies consistently show that people are significantly more receptive to advice that costs money than to free advice. Furthermore, the results suggest that this effect is due to the same forces that have been documented in the literature to explain the sunk costs fallacy.

E-mail address: fgino@andrew.cmu.edu

0749-5978/$ - see front matter © 2008 Elsevier Inc. All rights reserved.
doi:10.1016/j.obhdp.2008.03.001

Advice taking in decision making

Many advice-taking studies rely on the so-called Judge–Advisor System (JAS) paradigm (hereafter, JAS; Sniezek & Buckley, 1995; Sniezek & Van Swol, 2001). In the JAS, a "judge" is responsible for making a final decision. Before committing to the decision, however, she is exposed to advice from one or more "advisors" who share an interest in the decision problem. Within the JAS literature, studies have employed either "choice" or "judgment" tasks (Bonaccio & Dalal, 2006). In the first case, judges must choose from among several qualitative alternatives; in the second case, judges must provide quantitative estimates. In most JAS studies using choice tasks, an advisor chooses which piece of advice to give from among several options (e.g., Sniezek & Buckley, 1995); that advice might be expressed as "Choose Option X," for instance. By contrast, in experiments using judgment tasks, the advice usually is an estimate from another participant (e.g., Yaniv, 2003, 2004). For example, if a judge must estimate the year in which a certain event occurred (e.g., Yaniv, 2004) or a person's weight based solely on a photograph (e.g., Gino & Moore, 2007), the advice offered is the estimate of another participant facing the same task, expressed as "The Advisor's estimate is Y."

In the traditional JAS, judges are not allowed to decide whether or not to receive advice; they receive it by default. In some studies, however, advice is provided at the decision maker's request (e.g., Gardner & Berry, 1995). In the three studies presented here, participants received advice in the form of a quantitative estimate, as in many previous JAS experiments employing judgment tasks. In all of the studies, the advisors were no more or less informed on average than the participants. In Studies 1 and 3 participants had the option of whether or not to receive advice, while in Study 2 they received advice automatically.

Discounting others' opinions

One of the primary findings of the JAS literature is that people tend to discount the advice they receive from others (e.g., Bonaccio & Dalal, 2006; Yaniv, 2004; Yaniv & Kleinberger, 2000). Researchers have attributed the discounting of advice to three causes: differential information (Yaniv, 2004; Yaniv & Kleinberger, 2000), anchoring (Tversky & Kahneman, 1974), and egocentric bias (Krueger, 2003). According to the differential information explanation, decision makers discount advice because they lack access to the advisor's internal thought processes, while having privileged access to the rationale behind their own opinions (Yaniv, 2004; Yaniv & Kleinberger, 2000). By contrast, the anchoring explanation asserts that a decision maker's initial estimate or choice serves as an anchor that he or she subsequently adjusts in response to received advice; such adjustment typically is insufficient and thus results in egocentric discounting of advice (Harvey & Fischer, 1997; Lim & O'Connor, 1995). Finally, according to the egocentric bias explanation, decision makers tend to believe that their own opinions and choices are superior to those of others, including the opinions and recommendations of advisors (Krueger, 2003).

Several variables have been found to reduce egocentric advice discounting. When judges perceive or know advisors to be knowledgeable experts, they tend to be more responsive to the advice (Goldsmith & Fitch, 1997; Harvey & Fischer, 1997; Sniezek, Schrah, & Dalal, 2004). Egocentric discounting also diminishes when judges perceive or know advisors to be older, better educated, wiser, or more experienced than themselves (Feng & MacGeorge, 2006). In the JAS literature, the presence of financial incentives or rewards based on performance also has been found to reduce egocentric advice discounting (Dalal, 2001; Sniezek et al., 2004). Egocentric advice discounting is also affected by the quality of the advice received; the higher the quality of the advice, the less the advice is discounted (Yaniv & Kleinberger, 2000). Nonetheless, judges often discount good advice (Gardner & Berry, 1995). Finally, judges weigh advice differently depending on the difficulty of the judgment they are facing: while they overweight advice from equally informed others on difficult tasks, they underweight it on easy tasks (Gino & Moore, 2007).

The impact of advice cost on advice use

In the studies reviewed above, advice was given for free. To the best of my knowledge, only two studies in the JAS literature (Patt, Bowles, & Cash, 2006; Sniezek et al., 2004) have included conditions in which participants paid for the advice they received.

Sniezek et al. (2004) employed a reward condition in which judges were paid based on the accuracy of their estimates of the prices of backpacks. Judges had the power to allocate part of the potential reward between themselves and their advisors, either before task inception or after task completion. The authors found that committing money for expert advice increases both advice use and estimation accuracy. Building on this work, Patt et al. (2006) showed that prepayment for expert advice enhances its credibility and, as a result, increases advice use.

While both studies show a significant effect of prepayment for *expert advice* on advice use, the use of expert advice might have led participants to believe that the advice was valuable, thus confounding the impact of advice quality with the effect of advice cost. In addition, the researchers did not thoroughly investigate the psychological mechanisms behind this effect. This paper addresses both issues by examining whether the cost of advice affects the degree to which people use it, *holding constant advice quality*, and by exploring why the paid-advice effect occurs. In the studies, participants received advice from other participants, who, on average, were equally informed and knowledgeable. Such investigation might lead to a better understanding of advice taking in decision making.

Hypotheses development

Research on individual decision making has shown that people's economic behavior is influenced by sunk costs (Arkes & 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 decisions (Arkes & 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 (e.g., Arkes & Blumer, 1985; Staw, 1976). Thus, 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.

Related work in social psychology 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 resolve this inconsistency by developing a view of the advice as more worthwhile than she would otherwise and therefore using the advice.

Prior research has also demonstrated that people often rely on price as an indicator of quality. Although this can be 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 & Petroshious, 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" paid advice.

These arguments led to the following main hypothesis:

Hypothesis 1: People assign a significantly greater weight to paid advice than they do to free advice.

Two studies were conducted to test this hypothesis. A third experiment was then designed to investigate the psychological mechanism behind the paid-advice effect. In the three studies, advice was expressed in terms of quantitative estimates (e.g., the advice a participant received might have been "The Advisor's estimate is 1978"). This type of advice is 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 what specific strategy to implement or what decision to make in the face of uncertainty.

The first two studies manipulated the cost of advice within subjects: advice was either free or costly. This manipulation was crossed with a between-subjects factor varying the order in which advice of a certain type was presented to participants: free advice first or costly advice first. I included sequence as factor in the experimental design of the first two studies because of its potential relevance. One might expect that people who pay for the advice first use it more or less than those who pay for it after previously receiving it for free. Independent of their direction, sequence effects could be potentially interesting in the research presented here, particularly given the non-expert nature of the advice. Note, however, that I did not develop specific hypotheses with respect to sequence effects on advice taking.

Study 1

Method

Experimental sessions were conducted on computers in the 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 $24 during the experiment.

Participants

Seventy-three graduate and undergraduate students, 35 males and 38 females ranging in age from 19 to 26, participated in this study. They were recruited via ads that offered money in exchange for participating in an experiment on decision making. To avoid having experts in the lab, history majors could not participate. Thirty-seven students participated in Sequence 1, and 36 students participated in Sequence 2.

Procedure

Upon arriving at the laboratory, participants were registered and randomly assigned to one of two experimental conditions. Participants received a copy of 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 given a questionnaire and asked to answer questions about their gender, occupational status, and age.

The experiment consisted of four phases. In each phase, participants were asked to estimate dates of specific events in US history (within the last 400 years) and to provide their estimates privately, without communicating with other participants. In addition to estimating the correct year of a given event, participants were asked to provide lower and upper bounds of their 90% credible interval. In particular, participants were asked to provide a low and a high estimate such that they were 90% sure that the true date of the event fell within the specified range.

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: "The advice comes from the estimate another participant provided for same question. In particular, for each question the advisor's estimate comes from a pool of estimates collected in an earlier study in which respondents were asked to provide a best estimate for each question."

The values used as advice were held constant across participants and across the two conditions (free vs. costly advice). So, for instance, participants received the same advice for the question "In what year was NATO formed?" asked in Phase 2 of both Sequence 1 (when advice was free) and Phase 2 of Sequence 2 (when advice was costly).

Participants did not receive feedback after answering each question; hence, they had no opportunity to learn anything about the advisors and their expertise. Thus, they could not weigh advice differently depending on the advisors' reputations. Moreover, as participants were explicitly told, advisors had been randomly chosen by the researcher prior to the study on a question-by-question basis. The lack of performance feedback and varied source of advice presumably made it very hard—if not impossible—for participants to assess the quality of the advice they received.

The experiment employed a 2 (advice cost: free vs. costly advice) $\times$ 2 (sequence: free advice first vs. costly advice first) mixed design in which advice cost was a within-subject factor and sequence was a between-subject factor. The design of Study 1 is shown in Table 1. In each condition, participants were given information about all four phases in the instructions they read at the beginning of the study. Next, I provide details about each phase in each of the two sequence conditions.

Sequence 1

In Phase 1, participants were asked to answer 15 questions about US history (e.g., "In what year was the Cuban missile crisis?").[1] The order in which questions were presented to participants was counterbalanced. The questions are listed in Appendix A. In Phase 2, participants were asked to answer the same series of questions; this time, however, they could choose to obtain "advice" regarding the correct answers.

As participants were explicitly told, for each question, the advisor's estimate came from a pool of estimates collected in an earlier study in which respondents provided a best estimate for each question. Yaniv and Kleinberger (2000) noted that adequate ecological validity is maintained when estimates are sampled from pools of data.

For each question, the experimenter randomly selected one value to be used as advice from among answers of 50 people who had previously participated in a pilot study. In the pilot study, participants were asked to answer the 30 questions to be used in Study 1 in exchange for $5. The pilot study was conducted with the sole aim of collecting a pool of estimates to use in Study 1. Participants in the pilot study were recruited from the same population used to

[1] I initially chose 60 questions heuristically. To avoid the use of any kind of normalization in the analysis, I chose questions that used the same response scale, i.e., questions about specific dates of events in US history within the last 400 years. Then I conducted a pilot study in which I asked 30 participants to answer the set of 60 questions. This allowed me to screen out some of the questions. In selecting the 30 questions used in the experiment out of the 60 used in the pilot, I followed two rules. First, I kept questions that people could intelligently estimate—that is, they had some rough idea about when the historical event occurred, yet did not know (with high probability) exactly when. Second, I aimed for roughly comparable variances in the responses.

Table 1
Experimental design used in both Study 1 and Study 2

	Sequence 1	Sequence 2
Phase 1	Fifteen questions about US history	Fifteen questions about US history (same questions as in Condition 1)
Phase 2	Fifteen questions about US history (same questions as in Phase 1) together with *FREE advice* if wanted [Study 1] or received by default [Study 2]	Fifteen questions about US history (same questions as in Phase 1) together with *PAID advice* if wanted [Study 1] or received by default [Study 2]
Phase 3	Other 15 questions about US history	Other 15 questions about US history (same questions as in Condition 1)
Phase 4	Fifteen questions about US history (same questions as in Phase 3) together with *PAID advice* if wanted [Study 1] or received by default [Study 2]	Fifteen questions about US history (same questions as in Phase 3) together with *FREE advice* if wanted [Study 1] or received by default [Study 2]

recruit participants in the actual study. Students with a major in history could not participate in the pilot study. This is an important feature of my three studies: while in Sniezek et al. (2004) and Patt et al. (2006) participants received advice from experts, in my experiments, participants received advice from people who, on average, were as equally knowledgeable as they were.

In Phase 2 of Sequence 1, the advice was free; before the start of Phase 2, the computer asked participants once whether they wanted the advice or not. If they accepted, they received advice for each of the 15 questions in the phase.

In Phase 3, participants were asked to answer a different set of questions about US history; the rules were exactly the same as in Phase 1. In Phase 4, as in Phase 2, participants could choose whether to receive advice. However, in Phase 4, participants had to pay for the advice. If they accepted the offer of costly advice, they had to pay $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. In other words, participants were offered advice not on a question-by-question basis but for sets of questions. This procedure was intended to create a feeling of commitment for those who chose to receive advice in the paid-advice condition.

Sequence 2

Sequence 2 was similar to Sequence 1, but in this case, to test for sequence effects, advice cost money in Phase 2 and it was free in Phase 4.

Payment

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 each phase, participants received 40 cents 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 used the "weight-of-advice" (WOA) measure, which has been previously employed in several studies (see Gino & Moore, 2007; Harvey & Fischer, 1997; Yaniv, 2004, in the context of advice taking; and Hell, Gigerenzer, Gauggel, Mall, & Muller, 1988, in the context of memory). WOA reflects how much a subject uses the advice she receives (Yaniv, 2004). WOA is defined as follows: $\mathrm{WOA} = \frac{|\text{final estimate} - \text{initial estimate}|}{|\text{advice} - \text{initial estimate}|}$. In the studies presented here, the final estimate was the point estimate provided in Phase 2 and Phase 4 when advice was available, and the initial estimate was the point estimate provided in Phase 1 and Phase 3. 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 Phases 1 and 2 or Phases 3 and 4).

The WOA is equivalent to 0 when a participant entirely discounts the advice. In such a case, the participant's final estimate equals her initial one, meaning that she did not change her initial decision after receiving the advice. The opposite is true when a participant follows the advice: her initial estimate shifts completely toward the advice. In this case, WOA equals 1, because the final estimate equals the received advice. Finally, when WOA equals a value between 0 and 1, the participant weighs both her initial estimate and the received advice positively, and partial discounting results. For instance, a WOA of 0.5 means that the participant took an average of the advice and her initial estimate. On average, this is the optimal strategy that people should use to determine their estimates, if opinions are expressed as quantitative estimates and assuming that judge and advisor are equally well informed (Larrick & Soll, 2006). Because it cancels out errors, averaging tends to produce more accurate estimates (Larrick & Soll, 2006). Evidence of the benefits of averaging comes from a wide range of fields, from psychiatry and meteorology to economics and forecasting (for references of such studies, see Larrick & Soll, 2006; Surowiecki, 2003).

Interval estimates

As mentioned earlier, in addition to estimating the correct year of a given event, participants were asked to provide lower and upper bounds of their 90% credible interval. It can be reasonably assumed that the size of a participant's interval estimate reflects her assessments of her own knowledge, her confidence in the answers she provided as her best estimates, and, at least partially, the accuracy of her answers (Yaniv, 2004; Yaniv & Foster, 1995, 1997). I will refer to the interval estimates as "credible intervals."

Results

In all of the studies, I first conducted analyses with gender, age, and occupational status as explanatory variables. No main effects or interaction effects involving these variables were found in any of the studies. Therefore, I only report analyses in which these factors are collapsed and do not discuss their effect any further.

The results of Study 1 are summarized in Table 2. Following the procedure used in prior studies (e.g., Gino & Moore, 2007; Yaniv, 2004), in the analyses of the WOA values discussed below, I left out cases in which the advice equaled the initial estimate, since WOA in those cases equaled a number divided by 0. In general, in these cases, it is not possible to quantify how much a participant did or did not use the advice. As for cases in which the WOA was not well defined (that is, those in which the final estimate did not fall between the initial estimate and the advice, and WOA was thus greater than 1), I chose to change values above 1 to 1. In general, this might happen when a participant provides a final estimate (e.g., 1980) that is further away from the advice (e.g., 1965) than his initial estimate (e.g., 1970). In experiments on advice taking, cases in which WOA is not well defined are usually rare.

In Study 1, I changed values above 1 to 1 in 3% of the cases. Another possibility would have been to leave out cases in which WOA was not well defined. However, as long as the number of such cases is small, the method used should make little difference. I analyzed the data of the three studies using each of the two methods. The

Table 2
Summary of results for Study 1 and Study 2, pooled by sequence and combined across sequence conditions

	Study 1			Study 2		
	Total	Accepted free advice	Accepted paid advice	Total	Free advice	Paid advice
Sequence 1						
N	37	36 (97%)	17 (46%)	44		
Mean WOA		0.49	0.71		0.48	0.63
(Std. Dev.)		(0.08)	(0.10)		(0.09)	(0.04)
Sequence 2						
N	36	33 (92%)	24 (67%)	44		
Mean WOA		0.44	0.74		0.36	0.61
(Std. Dev.)		(0.08)	(0.05)		(0.07)	(0.03)
Combined						
N	73	69 (95%)	41 (56%)	88		
Mean WOA		0.46	0.72		0.42	0.62
(Std. Dev.)		(0.09)	(0.08)		(0.09)	(0.03)

Note. The WOA measures the extent to which participants used advice and varies from 0 (100% discounting of advice) to 1 (0% discounting of advice). To compute the mean WOA, I first calculated the average WOA across participants per question and then calculated the average of such values across questions. Obviously, in Study 1, I only had a measure for WOA if the participant opted to receive the advice. The nature and significance of the results do not change if I compute the mean WOA by calculating WOA across questions per participant and then calculate the average of such values across participants.

nature and significance of the results did not change. For Study 1, the results for the analyses conducted by leaving out ill-defined WOA values are mentioned separately in each of the main analyses presented below.

Buying rate

In Sequence 1, 36 participants out of 37 (97%) opted to receive advice in the free-advice treatment, and 17 (46%) bought advice in the paid-advice treatment. In Sequence 2, 33 participants out of 36 (92%) opted to receive advice in the free-advice treatment, and 24 (67%) bought advice in the paid-advice treatment. It appeared that people who were offered free advice before paid advice bought advice at a lower rate than did people who were offered paid advice before free advice. Yet testing of whether sequence affected the buying rate found only a marginal effect, $\chi^2(1, N = 73) = 3.18$, $p = .07$, Cramer's $V = .21$.

Impact of advice cost

Hypothesis 1 predicted that the cost of advice had a significant effect on the weight participants assigned to the advice itself. To test this hypothesis, the values for WOA in the paid-advice treatment and in the free-advice treatment were compared. At the aggregate level, the mean WOA per question in the paid-advice treatment was 0.72 ($SD = 0.08$). By contrast, in the free-advice treatment, WOA was only 0.46 ($SD = 0.09$). This difference was statistically significant, $t(58) = 11.82$, $p < .001$, thus supporting Hypothesis 1.

Note that when ill-defined WOA values were eliminated, the nature and significance of the results did not change. At the aggregate level, the mean WOA per question was significantly higher in the paid-advice treatment ($M = 0.72$, $SD = 0.08$) than in the free-advice treatment ($M = 0.43$, $SD = 0.08$), $t(58) = 13.19$, $p < .001$.

Next, I analyzed the results for the impact of advice (presented in Table 2) by looking at the mean WOA per question both across sequences and across phases. Note that the analysis conducted across phases, unlike the within-subjects comparisons, held constant the content of the questions. In the analyses, I compared the weight that participants assigned to the advice when it was free to the weight they assigned to it when it cost money. As predicted, WOA was higher in the paid-advice treatment than in the free-advice treatment. The difference was statistically significant, both when analyzing the results "across phases"($t[28] = 9.68$, $p < .001$, for Phases 1 and 2, that is, Sequence 1-Free vs. Sequence 2-Paid; $t[28] = 7.69$, $p < .001$, for Phase 3 and 4, that is, Sequence 1-Paid vs. Sequence 2-Free) and when analyzing them "across sequence conditions" ($t[28] = 6.34$, $p < .001$, Free vs. Paid in Sequence 1; $t[28] = 11.49$, $p < .001$, Free vs. Paid in Sequence 2).

Note that when ill-defined WOA values were eliminated, the nature and significance of the results did not change. The difference in the average WOA between the paid-advice condition and the free-advice condition was statistically significant, both "across phases" ($t[28] = 10.79$, $p < .001$, for Phases 1 and 2, that is, Sequence 1-Free vs. Sequence 2-Paid; $t[28] = 8.24$, $p < .001$, for Phase 3 and 4, that is, Sequence 1-Paid vs. Sequence 2-Free) and "across sequence conditions" ($t[28] = 7.33$, $p < .001$, Free vs. Paid in Sequence 1; $t[28] = 12.32$, $p < .001$, Free vs. Paid in Sequence 2).

Testing for selection bias

One potential problem with these analyses is that most of the participants in the free-advice condition were included in the analyses (because most of them chose to receive advice), while only a smaller percentage of participants in the paid-advice condition were included in the analyses (because fewer of them chose to purchase the advice). This is problematic because those who chose to purchase advice and those who did not may differ in ways that could influence the results. For example, if those who are inclined to purchase advice are also less confident than the average participant, then my results may not be driven by the effect of paid advice. Instead, they could be attributed to a selection bias, if I were comparing a random sample of the population (in the free-advice condition) with a sub-sample of the population that tends to be less confident (in the paid-advice condition).

To address this concern, I conducted further analyses to explore whether participants who bought advice differed from those who did not. Specifically, I investigated whether differences existed between the two groups in the width of their initial credible intervals, which provide a measure of participants' confidence in their answers prior to having the opportunity to buy advice. The width of participants' initial credible intervals for each question served as the dependent variable in a repeated-measure ANOVA with buying costly advice (yes vs. no) as the between-subjects factor (repeated measure on question). This analysis showed that people who paid for the advice were as confident in the answers they provided as people who chose not to buy the advice, $F(1,71) = 1.26$, $p = .27$.

Furthermore, the data allowed for a conservative test of the selection bias explanation. In the following analysis, I focused only on the participants who bought advice in each condition. In other words, I analyzed the data "controlling for people," considering the behavior of participants who chose to pay for costly advice to

Table 3
Descriptive statistics of WOA values for participants who bought the advice, pooled by condition (Study 1)

	Sequence 1 ($N = 17$)		Sequence 2 ($N = 24$)	
	Mean	Std. Dev.	Mean	Std. Dev.
Free advice	0.40	0.06	0.42	0.09
Paid advice	0.71	0.10	0.74	0.05

Note. The WOA measures the extent to which participants used advice and varies from 0 (100% discounting of advice) to 1 (0% discounting of advice).

their behavior when the advice was free. Computed values for WOA, pooled by condition, are reported in Table 3. As the table shows, in each condition, the mean WOA was higher in the paid-advice treatment than in the free-advice treatment. A repeated-measure ANOVA with WOA values as the dependent variable, advice (free vs. paid) as a within-subjects factor and sequence (1 vs. 2) as a between-subjects factor (repeated measure on question) revealed that the difference was statistically significant, $F(1,8) = 9.93$, $p = .014$, $\eta^2 = .55$. These results are the core test of Hypothesis 1.

Note that when ill-defined WOA values were eliminated, the nature and significance of the results did not change. A repeated-measure ANOVA with WOA values as dependent variable, advice (free vs. paid) as a within-subjects factor, and sequence (1 vs. 2) as a between-subjects factor (repeated measure on question) revealed that the difference was statistically significant, $F(1,6) = 7.90$, $p = .031$, $\eta^2 = .57$.

An alternative explanation

The previous analyses have shown that, as predicted, 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 initial decisions is a good predictor of their advice-seeking behavior (Cooper, 1991; Gibbons, Sniezek, & Dalal, 2003). I used a mixed-model analysis 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 was treated as a random selection from the general population; thus, I included participants in the model as random effects. I estimated the following model:

$$\mathrm{WOA}_{ij} = \alpha_0 + \beta_1 PAY_i + \beta_2 IE_{ij} + participants' \ RANDOM \ EFFECTS + \varepsilon,$$

where index i referred to participants and index j referred to questions. The dependent variable was the value for WOA for each participant and for each question. Explanatory variables were: (i) a dummy variable indicating whether the participant 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. Because I was interested in the effect of free vs. paid advice, the cost of advice was a fixed effect. Because the study participants were a sample from the larger population of interest, they were a random effect. While there was likely to be participant-to-participant variation in people's use of advice, I was not directly interested in that variation in this analysis. Furthermore, this analysis allowed me to exploit an important feature of the experimental design: participants had to commit to paying upfront for advice on all questions in a set, 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 1520 observations, each observation being a question answered by a participant. The results revealed a significant and positive effect of the initial range size on WOA ($\beta_2 = 0.001$, $t = 2.89$, $p = .004$): as one might expect, the wider the initial range, the higher the WOA. It appears that participants did weigh advice more heavily when they were less confident of their own knowledge. However, even after I controlled for this effect, the effect of paying for advice was still in the expected direction and statistically significant ($\beta_1 = 0.249$, $t = 10.49$, $p < .001$).

Hence, the results of the mixed-model analysis show that when participants paid for advice, they weighed it more than one would expect given their confidence in their answers. As expected, WOA increased as the initial range increased; holding the initial range constant, however, the WOA was significantly higher when the advice costs money than when it was free.

Sequence effects

As described above, people who paid for 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 sequence in which free vs. paid advice was received had an impact on WOA, I conducted a repeated-measure ANOVA in which participants' values for WOA served as dependent variable, pay (that is, free vs. paid advice) served as a within-subjects factor, and sequence (that is, Sequence 1 vs. Sequence 2) served as a between-subjects factor (repeated measure on question). The main effect for pay replicated the finding of the previous analyses: Participants weighed paid advice significantly more than free advice, $F(1,8) = 9.93$, $p = .014$, $\eta^2 = .554$. The main effect for sequence was not significant ($F(1,8) < 1$, $p = .58$), nor was the Pay $\times$ Sequence interaction ($F(1,8) < 1$, $p = .47$). These results suggest that participants who received the advice for free first were equally prone to using the advice when they paid for it as were participants who were first offered the advice at a cost.

Discussion

The results of Study 1 support Hypothesis 1: Participants use advice significantly more when it costs money 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.

The findings of the first study suggest a facet of the paid-advice effect that should be further explored. This concerns the buying rate in the paid-advice treatment: only about half of the participants bought the advice when they were given the option not to buy it in the paid-advice treatment. There might have been a difference between buyers and non-buyers that was not detected by Study 1. The analyses I presented help to address this concern, but do not eliminate it completely.

I designed a second study to explore this issue. 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 eliminated the endogeneity problem encountered in Study 1 at the expense of the choice of getting advice.

Study 2

Method

Experimental sessions were conducted in the computer laboratory of a northeastern US 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 $24 during the experiment.

Participants

Participants were recruited via ads that offered money in exchange for participating in an experiment on decision making. On average, each experimental session lasted 45 min. Eighty-eight people agreed to participate (49% male). The average age of participants was 26 ($SD = 7$). Most participants were students (93%). 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 paid), 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 sequence conditions. They were also told that, if assigned to Sequence 1, they would receive advice for free in Phase 2 and paid advice in Phase 4. If assigned to Sequence 2, they would receive paid advice in Phase 2 and free advice in Phase 4. As in Study 1, in the paid-advice treatment, the advice cost was $4. Participants were told that they would be charged this amount of money when receiving paid advice. Note that, as in Study 1, participants were given information about how advice was generated. In particular, they were told that the experimenter had randomly selected the advice on each question prior to the study, choosing from a pool of estimates collected in a previous session. The values used for advice in Study 2 were the same as in Study 1, as were the questions that participants were asked to answer. The order in which questions were presented to participants was counterbalanced.

Payment

Participants' payoffs were computed as the sum of the show-up fee and available bonuses that they achieved. As in Study 1, participants received bonuses based on the accuracy of their point estimates. They received 40 cents for each question in which their best estimate fell within a range of 10 years minus or plus the true historical 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 2. As in the analyses conducted for Study 1, 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 paid-advice treatment and in 2% of cases in the free-advice treatment. Analyses were also conducted by leaving out cases in which WOA was not well defined. The nature and significance of the results did not change, as mentioned in the analyses reported below.

The impact of advice cost

Overall, the results of Study 2 showed further support for Hypothesis 1. At the aggregate level, the mean WOA per question was higher in the paid-advice treatment ($M = 0.62$, $SD = 0.03$) than in the free-advice treatment ($M = 0.42$, $SD = 0.09$), $t(58) = 10.56$, $p < .001$. Table 2 reports the WOA values, pooled by both sequence and treatment. As in Study 1, I analyzed these results by looking at the mean WOA per question both across sequence conditions and across phases. WOA in the paid-advice treatment was higher than in the free-advice treatment. The difference was statistically significant, as shown by analyzing the data both ''across phases''($t[28] = 5.48$, $p < .001$, Sequence 1-Free vs. Sequence 2-Paid; $t[28] = 14.17$, $p < .001$, Sequence 1-Paid vs. Sequence 2-Free) and ''across sequence conditions'' ($t[28] = 6.26$, $p < .001$, Free vs. Paid in Sequence 1; $t[28] = 13.64$, $p < .001$, Free vs. Paid in Sequence 2).

Note that when eliminating ill-defined WOA values, the results are as follows. In Sequence 1, the average WOA per question was higher in the paid-advice treatment ($M = 0.63$, $SD = 0.04$) than in the free-advice treatment ($M = 0.47$, $SD = 0.08$), $t(28) = 6.760$, $p < .001$. Similarly, in Sequence 2, the average WOA per question was higher in the paid-advice condition ($M = 0.60$, $SD = 0.03$) than in the free-advice treatment ($M = 0.34$, $SD = 0.07$), $t(28) = 13.588$, $p < .001$.

Sequence effects

In Study 2, participants who received the advice for free before paying for it (Sequence 1) used it more than did those who received it for free after previously paying for it (Sequence 2). To test for a sequence effect, I conducted a repeated-measure ANOVA in which participants' WOA values served as the dependent variable, pay (free vs. paid advice) served as a within-subjects factor, and sequence (Sequence 1 vs. Sequence 2) served as a between-subjects factor (repeated measure on question). Consistent with the previous analyses, I found a main effect for pay: Participants used paid advice significantly more than free advice, $F(1,18) = 7.33$, $p = .014$, $\eta^2 = .289$. However, the main effect for sequence was not significant, $F(1,18) < 1$, $p = .83$, nor was the Pay × Sequence interaction, $F(1,18) = 1.63$, $p = .22$.

Discussion

The results of Study 2 are consistent with Hypothesis 1 and thus provide further support for the findings from Study 1. Specifically, people weigh advice more heavily when it costs money than when it is free. In Study 2, as in several other studies from the JAS literature, advice was imposed on participants, but they were able to choose the degree to which they followed the advice. Given that advice was provided by default, the distinction between paid and free advice was essentially illusory in this second study. Participants across conditions faced the same judgment task, with the same potential payoff. Yet those participants told that a set of advice was free treated it differently from those told that it was paid. While this uniformity follows directly from removing choice from the paradigm used in Study 1, it also serves to illustrate how psychological the paid-advice effect is.[2]

Study 1 and 2 demonstrated how paying for advice affects the degree to which people value it. The third study investigated the mechanism underlying this paid-advice effect.

Study 3

While Study 1 and 2 showed that the effect of paid advice is strong, they did not explain why people weighed paid advice more heavily than free advice. Study 3 examined this question.

As I suggested earlier, when buying advice, people might fall prey to the sunk-cost fallacy. In the first two studies, 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. As a result, participants were more receptive to paid advice than to free advice. 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.

As discussed in the Introduction, the paid-advice effect can be predicted also by research on cognitive dissonance (Festinger, 1957). Both the sunk costs and the cognitive-dissonance explanations implicitly assert that the paid-advice effect occurs because

[2] I thank one of the anonymous reviewers for making this insightful observation.

a participant overweighs an advisor's response relative to her own during the information-processing stage of decision making.

However, the effect also may occur at the encoding stage. The price attribute that people attach to the advice might signal quality (or advisor's expertise), which, in turn, affects how people automatically receive the information and encode it into working memory. Study 1 and 2 were designed to eliminate the viability of this explanation, but may not have ruled it out completely. Indeed, in the first two studies, participants were explicitly told that the quality of the advice received did not change based on whether the advice was free or cost money. Participants also were told explicitly what the advice consisted of and who provided it. By holding constant advice quality and information source, the studies tried to eliminate the viability of the "price as a signal of quality" explanation. However, despite the clear instruction participants received, they still might have perceived paid advice as more valuable than free advice because of its "price."

To eliminate this possibility and to examine the psychological processes behind the paid-advice effect, Study 3 included measures of participants' perceptions of advice quality together with measures of cognitive dissonance and sunk costs.

Measuring the sunk-cost fallacy

Prior research has found that people routinely succumb to the sunk-cost fallacy, pursuing a goal more doggedly if they have invested significant money, time, or effort than if they have not (Arkes & Blumer, 1985, p. 124). For instance, Thaler (1980) speculated that a family would be more likely to drive through a snowstorm to attend a basketball game if they had purchased $40 tickets to the game than if they have been given the same tickets for free.

Prior research has suggested various explanations for the sunk-cost effect. Among them there are the desire to not appear wasteful (Arkes & Blumer, 1985), the need to justify a prior course of action (Brockner, 1992; Staw, 1976, 1981), and the tendency to be risk seeking in light of previous losses (Garland & Newport, 1991; Kahneman & Tversky, 1979; Thaler, 1980). Although these explanations are driven by a different psychological process, they share one important feature: each requires an individual to track the costs and benefits associated with a particular transaction or at least to be aware of them (Gourville & Soman, 1998). In short, people's mental accounting decisions help to explain the sunk-cost effect (Gourville & Soman, 1998). Gourville and Soman (1998) use Thaler's example (1980) to describe such decisions:

> [I]n the case of the paid for tickets, the family opened a mental account on making payment for those tickets, with the expectation of closing that account on seeing the game. If they do not attend the game, and the tickets cannot be resold or exchanged, the family is forced to close their account with no offsetting benefit, resulting in a perceived loss of $40. In contrast, had the tickets been free, the family can forgo the game and close their mental account without the prospect of a loss. Contrasting these two outcomes, all else being equal, there is greater pressure on the family to attend the game in the case of the paid for tickets (to avoid the perceived loss of $40) than in the case of the free tickets (where there is no loss to avoid) (Gourville & Soman, 1998, p. 162).

Extending this reasoning, one can conclude that the psychological impact of any past payment on future consumption should increase monotonically with the size of that payment (Arkes & Blumer, 1985). Thus, Thaler's hypothetical family will be more likely to go to the game not only if they spent $40 rather than $0 on their tickets, but also if they spent $40 instead of $20 (Gourville & Soman, 1998).

In line with this reasoning suggested by Gourville and Soman (1998), Study 3 included a measure for sunk costs by employing three conditions for advice cost: in the free-advice condition, participants received advice for free; in the low-cost-advice condition, participants paid $1 to receive advice; and in the high-cost-advice condition, participants paid $2 in exchange for advice. Details on each condition are given in the method section below.

Based on the findings regarding the sunk-cost effect, I predict that:

> *Hypothesis 2:* People in the high-cost-advice condition will weigh advice from others significantly more heavily than people in the low-cost-advice condition and in the free-advice condition. Furthermore, people in the low-cost-advice condition will weigh advice from others significantly more heavily than people in the free-advice condition.

Results consistent with this hypothesis would suggest that the most probable cause for the paid-advice effect demonstrated in the first two studies is the consideration of sunk costs. In other words, empirical support for Hypothesis 2 would suggest that the paid-advice effect is due to the same forces that have been documented in the literature to explain prior instances of the sunk costs fallacy.

Measures for alternative explanations

Study 3 also included a measure of the perceived quality of the advice and a measure of participants' cognitive dissonance. Perceived quality of the advice was measured by asking participants to indicate how accurate the advice they received was, using a 7-point likert scale (from 1 = not accurate at all to 7 = very accurate). Participants' cognitive dissonance was measured using an adapted version of a valid and reliable multidimensional scale developed by Sweeney, Hausknecht, and Soutar (2000) to measure cognitive dissonance after purchase. In the cognitive-dissonance inventory, participants were asked to indicate how much they agreed with each of the scale items (e.g., "After I bought the advice I resented it."), using a 7-point scale (from 1 = strongly disagree to 7 = strongly agree).

Based on prior research on cognitive dissonance, the authors developed a scale that measured three distinct dimensions of cognitive dissonance: emotional, wisdom of purchase, and concern over deal. The first dimension refers to the psychological discomfort a person experiences after the purchase decision. The second dimension is defined as "a person's recognition after the purchase has been made that they may not have needed the product or may have not selected the appropriate one" (Sweeney et al., 2000). Finally, concern over deal refers to a person's recognition after the purchase has been made that they may have been influenced against their own beliefs by the people selling the product. The items included in the cognitive-dissonance inventory are listed in Appendix B.

Based on the price as a signal of quality explanation, one might expect ratings participants provided for perceived advice quality to have an impact on the weight-of-advice measure, which indicates the extent to which participants used the advice. Finally, based on the cognitive-dissonance explanation, one might expect the cognitive-dissonance measure to have an impact on advice use (measured by the WOA).

Method

Participants

As in Study 1 and 2, participants were recruited through ads offering money to participate in an experiment on decision making. One-hundred sixty-eight individuals agreed to participate (58% male). The average age of participants was 23 years

(*SD* = 7). Most participants were students from local universities (88%). Participants were randomly assigned to one of two conditions: the high-cost-advice condition and the low- cost-advice condition. Details on each of these conditions are provided below.

Procedure

Given that no sequence effects were found in the first two experiments, Study 3 did not include a sequence condition. The experiment was conducted on paper and consisted of two phases. The task people faced in each phase was the same one used in Study 1 and 2: Participants were asked to estimate dates of specific events in US history (within the last 400 years) and to provide their estimates privately without communicating with other participants. In addition to estimating the correct year of a given event, participants were asked to provide lower and upper bounds of their 90% credible intervals.

In Phase 1, participants were asked to answer 10 questions about US history. The questions are listed in Appendix C. Participants were asked to answer the same set of questions in Phase 2. The order in which questions were presented to participants in each phase was counterbalanced.

Once participants provided their estimates for each question in Phase 1, they read the following instructions:

> The second task of the study is to take the same history trivia quiz. As before, your goal is to answer as many questions correctly as you can. Differently from before, you have now the option of receiving advice on the correct answer. The advice comes from another participant's best estimate when asked the same question in a previous session. The advice has been randomly selected from all who have provided estimates in a previous session. If you choose to receive advice, then you will receive advice on the correct answer for each of the questions of the trivia quiz.
>
> In case you choose to receive advice, the experimenter will toss a coin: if the result of the coin toss is TAILS then you will receive advice for FREE; if the result of the coin toss is HEAD then you will HAVE TO PAY to receive the advice. The cost of the ADVICE is $2 [$1]. This $2 [$1] cost will be subtracted from your final payoff. As in the first part of the study you will receive $.50 every time your best estimate for the correct answer to a certain question is within 1% of the true date.
>
> Please indicate whether you want to receive advice on the answers to the trivia quiz you will take in the next part of the study by checking one of the following:
>
> O I want to receive advice. O I do not want to receive advice.

For participants in the high-cost-advice condition, these instructions reported a cost of $2 for receiving advice in Phase 2. For participants in the low-cost-advice condition, the instructions reported a cost of $1. As the instructions indicate, before Phase 2, participants were given the option of receiving advice. Participants who chose to receive advice were given another person's best estimate for each question in Phase 2. The experimenter randomly selected values used as advice from among answers of 50 people who had previously completed Phase 1. These values were held constant across participants. Participants who chose not to receive advice simply answered the same series of questions a second time in Phase 2. As in Studies 1 and 2, participants did not receive feedback until the end of the experiment.

Before answering Phase 2 questions, participants who received advice were shown the advice for each of the 10 questions. Then they were asked to indicate for each question how accurate the advice they received was, on a 1–7 scale (from 1 = not accurate at all to 7 = very accurate). These ratings were used as a measure of the perceived quality of the advice. After providing their ratings for perceived advice quality, participants were asked to answer the questions in the cognitive-dissonance inventory, which were designed to measure the degree of cognitive dissonance they experienced after choosing to receive advice.

After providing their rating measuring cognitive dissonance, participants who chose to receive advice were asked to answer the same 10 questions they had answered in Phase 1. After completion of Phase 2, participants in all conditions answered a final questionnaire with demographic questions. Finally, they were debriefed, thanked and dismissed.

Payment

As in Studies 1 and 2, participants' payoff was based on the accuracy of their estimates. Participants received $.50 each time their best estimate of the correct answer to a certain question fell within 1% of the true date. In addition, participants received $4 as a show-up fee for their participation in the study.

Results

The results of Study 3 are summarized in Table 4. There were 84 participants in the high-cost-advice condition and 84 in the low-cost-advice condition. Thirty-six participants (43%) in the high-cost-advice condition chose to get advice before the beginning of Phase 2; 23 of them ended up paying $2 for the advice, while the remaining 13 received it for free. In the low-cost-advice condition, 44 participants (52%) chose to get advice before the beginning of Phase 2; 23 of them ended up paying $1 for the advice, while the remaining 21 received it for free.

In the analyses presented below, I only included participants who decided to receive advice. Indeed, a measure of WOA can be computed only if a participant opted to receive advice. Note that all WOA values were well defined in Study 3.

The impact of advice cost

Hypothesis 1 predicted that that the cost of advice would have a significant effect on the weight participants assigned to the advice

Table 4
Summary of results for Study 3, pooled by condition

	LOW-cost-advice condition		HIGH-cost-advice condition	
N	84		84	
Opted to receive advice before Phase 2	44 (52%)		36 (43%)	
	Free advice	Paid advice	Free advice	Paid advice
N	21	23	13	23
Mean WOA	0.38	0.51	0.47	0.65
(Std. Dev.)	(0.11)	(0.11)	(0.11)	(0.12)

Note. The WOA measures the extent to which participants used advice and varies from 0 (100% discounting of advice) to 1 (0% discounting of advice). In order to compute the mean WOA, I first calculated the average WOA across participants per question and then calculated the average of such values across questions. The nature and significance of the results do not change if I compute the mean WOA by calculating WOA across questions per participant and then calculate the average of such values across participants. The table reports the number of participants in each condition and the descriptive statistics for the WOA measure.

itself. To test this hypothesis, the values for WOA in the paid-advice treatment and in the free-advice treatment were compared within each of the two paid-advice treatments (i.e., high-advice cost vs. low-advice cost). In the high-cost-advice condition, the mean WOA per question in the paid-advice treatment was significantly higher ($M = 0.65$, $SD = 0.12$) than in the free-advice treatment ($M = 0.47$, $SD = 0.11$), $t(18) = -3.50$, $p = .003$. Similarly, in the low-cost-advice condition, the mean WOA per question in the paid-advice treatment was significantly higher ($M = 0.51$, $SD = 0.08$) than in the free-advice treatment ($M = 0.38$, $SD = 0.09$), $t(18) = -2.72$, $p = .014$. These results provide further support for Hypothesis 1.

Based on Hypothesis 2, I expected that people in the high-cost-advice condition would weigh advice from others significantly more heavily than would people in the low-cost-advice condition and that people in the low-cost-advice condition would weigh advice from others significantly more heavily than would people in the free-advice condition. I thus compared the mean WOA values per question across conditions and found that the mean WOA per question in the paid-advice treatment was significantly higher in the high-cost-advice condition than in the low-cost-advice condition, $t(18) = -2.64$, $p = .017$. Note that the mean WOA per question in the free-advice treatment was higher in the high-advice cost condition than in the low-advice cost condition, but this difference was only marginally significant, $t(18) = -1.85$, $p = .081$.

Taken together, these results support Hypothesis 2 and suggest that the paid-advice effect is due to the same forces that have been documented in the literature to explain prior instances of the sunk costs fallacy.

Additional analyses were conducted to investigate the psychological mechanisms behind the paid-advice effect. Participants' WOA values were used as the dependent variable in a repeated-measure ANOVA with advice cost (paid vs. free) and advice condition (high-cost-advice condition vs. low-cost-advice condition) as between-subjects factors (repeated measure on question). This analysis revealed a significant main effect of advice cost, $F(1,64) = 9.57$, $p = .003$, $\eta^2 = .130$. As predicted in Hypothesis 2, and consistent with the consideration of sunk costs, the main effect of advice condition was also significant, $F(1,64) = 6.02$, $p = .017$, $\eta^2 = .086$. Instead, the Advice Cost × Advice Condition interaction was not significant, $F(1,64) < 1$, $p = .45$. These results are presented graphically in Fig. 1.

The impact of perceived quality and cognitive dissonance

Study 3 also included measures of perceived quality of the received advice and participants' cognitive dissonance after receiving the advice. The items used to measure cognitive dissonance were highly correlated and highly reliable (Cronbach's alpha = 0.85); thus, I computed the average of the items in the scale and used this value as a measure of cognitive dissonance. This measure, together with the ratings participants provided for advice quality, allowed me to examine whether cognitive dissonance and price as a signal of quality might explain, at least in part, the paid-advice effect.

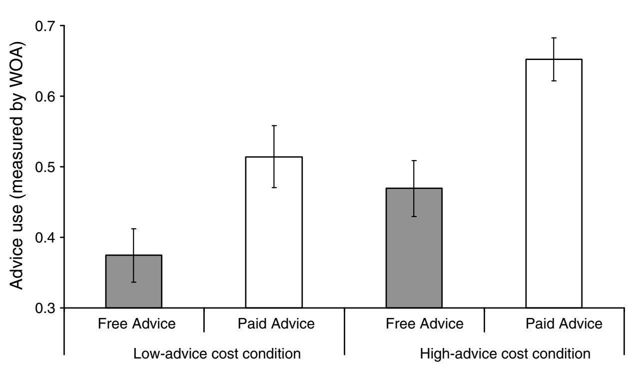

Fig. 1. Mean WOA values, pooled by condition (Study 3). Error bars represent standard errors of the means.

To investigate such a possibility, I first conducted partial correlation analyses. These analyses revealed a positive but insignificant correlation between perceived advice quality and advice use as measured by WOA ($r = .02$, $p = .59$). The analyses also revealed a positive but insignificant correlation between cognitive dissonance and WOA ($r = .03$, $p = .35$). Next, I conducted a mixed-model analysis and estimated the following model:

$$\text{WOA}_{ij} = \alpha_0 + \beta_1 \mathit{ADVICE\ COST}_i + \beta_2 \mathit{COND}_i + \beta_3 \mathit{QUALITY}_{ij} + \beta_4 \mathit{DISSONANCE}_i + \beta_5 \mathit{IE}_{ij} + \mathit{participants'\ RANDOM\ EFFECTS} + \varepsilon,$$

where index i referred to participants and index j referred to questions. The dependent variable was the value for WOA for each participant and for each question. Explanatory variables were: (i) a dummy variable indicating whether the participant got advice for free or by paying for it ($\mathit{ADVICE\ COST}_i$), (ii) a dummy variable indicating whether the participant was in the high-cost-advice condition or in the low-cost-advice condition (COND_i), (iii) the rating for the perceived quality of the received advice ($\mathit{QUALITY}_{ij}$), (iv) the rating for cognitive dissonance that participants experienced after receiving advice ($\mathit{DISSONANCE}_i$), and (v) 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 *ex ante*.

Results were based on a total of 800 observations, each a question answered by a subject. The results revealed a significant and positive effect of the initial range size on WOA ($\beta_5 = 0.002$, $t = 5.03$, $p < .001$): as one might expect, the wider the initial range, the higher the WOA. It appears that participants did weigh advice more heavily when they were less confident of their own knowledge. However, even after I controlled for this effect, both the effect of advice cost ($\beta_1 = 0.149$, $t = 4.64$, $p < .001$) and the effect of advice condition ($\beta_2 = 0.111$, $t = 3.82$, $p < .001$) were still in the expected direction and statistically significant. Instead, neither the effect of perceived quality ($\beta_3 = 0.001$, $t < 1$, $p = .92$) nor the effect of cognitive dissonance ($\beta_4 = 0.016$, $t = 1.23$, $p = .22$) were significant, despite the fact that they were positive, as one might expect.

Hence, the results of the mixed-model analysis show that when participants paid for advice, they weighed it more than one would expect given their confidence in their answers. As expected, WOA increased as the initial range increased; holding the initial range constant, however, the WOA was significantly higher when the advice cost money than when it was free. More important, this analysis provides support for Hypothesis 2, which is consistent with the prediction that the consideration of sunk costs might be the cause for the paid-advice effect. In addition, this analysis does not provide evidence in support of alternative explanations (i.e., price as a signal of quality and cognitive dissonance) for the demonstrated effect.

General discussion and conclusions

Many 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 costs money than when it is free. This hypothesis was tested in three experiments requiring participants to answer questions about US history with or without advice from others. The results of the studies show that participants relied more heavily on advice when it cost money than when it was free. The results also suggest that this paid-advice effect is due to the same forces that have been documented in the literature to explain prior instances of the sunk costs fallacy.

The cost of advice affected the degree to which participants used advice but did not affect the value gained by following advice. In the studies, advice came from another participant who was randomly chosen on a question-by-question basis. On average, advisors were as equally informed or knowledgeable as judges. In fact, individuals who were history experts could not participate in the studies. 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, as no performance feedback was provided. When advice cost money, participants weighed their personal opinions less than others'. When advice was free, they instead weighed their personal opinions more than others'.

My findings suggest several areas for 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, in line with the procedure used in Study 3, would enable an exploration of the extent to which the "size" of cost matters. In particular, 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?

Future research could also further examine the psychological mechanisms behind the paid-advice effect demonstrated in the present work. As suggested earlier, the results of my third study suggest that the paid-advice effect is due to the same psychological processes behind the sunk-cost effect. While prior research has distinguished several of these forces (e.g., people's desire to not appear wasteful), they all require individuals to make some mental accounting decisions. Further investigations into the influence of these mental accounting processes on the paid-advice effect are warranted. Such investigations would indeed provide a better understanding of the boundary conditions of the paid-advice effect.

The type of decision the person receiving advice faces might be another fruitful direction for future work. In many contexts, for instance, people are faced with "yes" or "no" decisions (e.g., "Should I have surgery?"). In such cases, a person must either follow advice given to them (for instance, by a doctor) or ignore it. Future work could investigate the paid-advice effect when the flexibility of using advice to influence one's own choices is not an option (i.e., in the case of yes/no decisions). Researchers could also explore the paid-advice effect empirically in settings such as consulting, healthcare, or product forecasting.

Another direction for future research is the study of advice that comes in the form of emotional or social guidance (e.g., in the case of a close family member offering advice). If such qualitative advice were being offered in addition to advice from a paid source (e.g., psychotherapy), at what point might cost effects be over-ridden by factors such as family trust? Similarly, in many real-world settings, people offer their opinions together with reasons for suggesting a certain course of action. Future studies could explore the impact such reasons have on judges' final decisions.

Finally, future research could investigate whether the paid-advice 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 (Bonaccio & Dalal, 2006). 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 & Dalal, 2006). Once the definition of advice is broadened, a closer exploration of the finding that people use paid advice more heavily than free advice would be a fruitful endeavor.

Acknowledgments

The author gratefully acknowledges the support of the staff and facilities of the Computer Lab for Experimental Research at the Harvard Business School and the Center for Behavioral Decision Research at Carnegie Mellon University. I am especially grateful to Max Bazerman, Dolly Chugh, Amy Edmondson, Andrew McAfee, and Don Moore for their insightful feedback on earlier drafts. I also want to thank the reviewers and the associate editor for their excellent feedback and for their wise guidance during the revision process. Special thanks to Zachary Burns for help with data collection.

Appendix A

Questions about US history used in both Study 1 and Study2.

Phase 1 and 2

1. When did the Congress declare war on Mexico?
2. In what year did Vietnam fall to Communists?
3. In what year was the Korean armistice signed?
4. In what year was NATO formed?
5. In what year was the Cuban missile crisis?
6. In what year did the Pilgrims reach Cape Cod?
7. When was the Truman Doctrine announced?
8. When was the Berlin wall built?
9. In what year was OPA (Office of Price Administration) established?
10. In what year was the first transcontinental railroad completed?
11. In what year the Women's rights convention at Seneca Falls take place?
12. In what year was the Presidential Succession Act?
13. When did Texas declare its independence (Battle of the Alamo)?
14. In what year was the National Labor Union formed?
15. In what year was the first US satellite in orbit?

Phase 3 and 4

1. When was the Standard Oil Trust organized?
2. In what year was the Civil Rights Act?
3. In what year was the Voting Rights Act?
4. When did the first American astronaut orbit earth?
5. In what year was SEC (Securities Exchange Commission) created?
6. When was the Emancipation Proclamation?
7. When was the Panama Canal opened to shipping?
8. In what year was Louisiana purchased?
9. When was the Haymarket Riot?
10. When did Mussolini seize power in Italy?
11. When was the Massachusetts Bay Colony founded?
12. When did the Korean War start?
13. In what year were US troops sent to South Vietnam?
14. When was the American Constitution first drafted?
15. When was the Bill of Rights ratified?

Appendix B

Cognitive-dissonance inventory used in Study 3. First dimension: items 1–3; second dimension: items 4–6; third dimension: items 7 and 8.

1. After I bought the advice I resented it.
2. After I bought the advice I felt disappointed with myself.
3. After I bought the advice I felt I'd let myself down.
4. I wonder if I really needed the advice.
5. I wonder if I made the right choice.
6. I wonder if I have done the right thing in buying the advice.
7. After I bought the advice I wondered if I'd been fooled.
8. After I bought the advice I wondered whether there was something wrong with the deal I got.

Appendix C

Questions about US history used in Study 3.

Phase 1 and 2

1. In what year was Louisiana purchased?
2. In what year was the first US satellite in orbit?
3. When did Texas declare its independence (Battle of the Alamo)?
4. In what year was the Presidential Succession Act?
5. In what year was the first transcontinental railroad completed?
6. In what year was OPA (Office of Price Administration) established?
7. When was the Truman Doctrine announced?
8. In what year was the Cuban missile crisis?
9. In what year was NATO formed?
10. When did the Congress declare war on Mexico?

References

Arkes, H. R., & Ayton, P. (1999). The sunk cost and concorde effects: Are humans less rational than lower animals? *Psychological 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.

Bazerman, M. H. (2001). *Smart money decisions: Why you do what you do with money (and how to change it for the better)*. New York: Wiley.

Bogle, J. C. (1999). *Common sense on mutual funds: New imperatives for the intelligent investor*. New York: Wiley.

Bonaccio, S., & Dalal, R. S. (2006). Advice taking and advice giving in decision making: An integrative literature review, and implications for the organizational sciences. *Organizational Behavior and Human Decision Processes, 101*, 127–151.

Brockner, J. (1992). The escalation of commitment to a failing course of action: Toward theoretical progress. *Academy of Management Review, 17*(1), 39–61.

Cooper, R. S. (1991). Information processing in the judge–adviser system of group decision-making. 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.

Dawes, R. M. (1994). *House of cards: Psychology and psychotherapy built on myth*. New York: The Free Press.

Feng, B., & MacGeorge, E. L. (2006). Predicting responsiveness to advice: Characteristics of the problem, the advice-giver, and the recipient. *Southern Communication Journal, 71*, 67–85.

Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston: Row & Peterson.

Freeman, J. P., & Brown, S. (2001). Mutual fund advisory fees: The cost of conflicts of interest. *Journal of Corporation Law, 26*(3), 609–673.

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.

Garland, H., & Newport, S. (1991). Effects of absolute and relative sunk costs on the decision to persist with a course of action. *Organizational Behavior and Human Decision Processes, 48*, 55–69.

Gibbons, A. M., Sniezek, J. A., & Dalal, R. S. (2003). 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. (2007). Effects of task difficulty on use of advice. *Journal of Behavioral Decision Making, 20*, 21–35.

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

Gourville, J. T., & Soman, D. (1998). Payment depreciation: The behavioral effects of temporally separating payment from consumption. *Journal of Consumer Research, 25*, 160–174.

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.

Hell, W., Gigerenzer, G., Gauggel, S., Mall, M., & Muller, M. (1988). Hindsight bias: An interaction of automatic and motivational factors? *Memory and Cognition, 16*, 533–538.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica, 47*, 313–327.

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

Larrick, R. P., & Soll, J. B. (2006). Intuitions about combining opinions: Misappreciation of the averaging principle. *Management Science, 52*, 111–127.

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

Malkiel, B. G. (2003). *A random walk down Wall Street*. New York: Norton.

Micklethwait, J., & Wooldridge, A. (1996). *The witch doctors: Making sense of the management gurus*. New York: Times Books.

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

Monroe, K. B., & Petroshious, S. M. (1981). 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.

Patt, A., Bowles, H., & Cash, D. (2006). Mechanisms for enhancing the credibility of an advisor: Prepayment and aligned incentives. *Journal of Behavioral Decision Making, 19*(4), 347–359.

Pfeffer, J., & Sutton, R. I. (2006). *Hard facts, dangerous half-truths, and total nonsense*. Boston, MA: Harvard Business School Press.

Prince, R. (2005). After the tide goes out. *The American Journal of Psychoanalysis, 65*(4), 319–326.

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

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.

Staw, B. (1976). Knee-deep in the big muddy: A study of escalating commitment to a chosen course of action. *Organizational Behavior and Human Performance, 16*(1), 27–44.

Staw, B. (1981). The escalation of commitment to a course of action. *Academy of Management Review, 6*(4), 577–587.

Surowiecki, J. (2003). *The wisdom of crowds*. New York: Doubleday.

Sweeney, J. C., Hausknecht, D., & Soutar, G. N. (2000). Cognitive dissonance after purchase: A multidimensional scale. *Psychology & Marketing, 17*(5), 369–385.

Thaler, R. H. (1980). Toward 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.

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

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

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.