Faculty Publications, Department of Psychology

Psychology, Department of

May 1995

Learning Lessons From Sunk Costs

Brian H. Bornstein
University of Nebraska-Lincoln, bbornstein2@unl.edu

Gretchin B. Chapman
University of Illinois at Chicago

Learning Lessons From Sunk Costs

Brian H. Bornstein, *Department of Psychology, Louisiana State University*
Gretchen B. Chapman, *Departments of Medical Education and Psychology,*
University of Illinois at Chicago

Study participants rated the quality of several arguments for continuing an original plan in sunk cost situations in order to (a) avoid wasting resources, (b) learn to make better decisions, (c) punish poor decision making, and (d) appear consistent. The lesson-learning argument was perceived as most appropriate when adult teachers taught lessons to others, the original decision was carelessly made, or if it consumed comparatively more resources. Ratings of the lesson-learning argument were higher for teacher-learner than for adult-alone situations, regardless of whether the learner was a child or an adult. The implications for improving decision making and judging whether the sunk cost effect is a bias are discussed.

An unfortunate fact of life is that things frequently do not go as planned. Whatever the reason—poor planning, a change of mind, altered circumstances, or mere bad luck—people must then decide whether to try something different or continue with their original plans. If individual preferences or circumstances have changed, future utility would often be increased by changing behaviors. For example, imagine that you decide you want to learn how to play the cello. After buying a $1,000 cello and spending an additional $200 for 3 months of lessons, you lose interest in it and want to stop taking lessons. What do you do?

Decision makers display the *sunk cost* (or *escalation*) *effect* if they continue to pour resources into a plan or project even though future utility would be increased by terminating the plan or switching to an alternative behavior (Arkes & Blumer, 1985; Brockner, 1992; Thaler, 1980). In the above example, a frequent response to the dilemma is to continue taking

more expensive lessons simply because a great deal of money and time have already been spent and would otherwise be wasted. Such a response is nonnormative because the decision maker is influenced by the initial investment (i.e., the cost of the cello and the lessons), even though it is irretrievably lost and therefore irrelevant in weighing the future consequences of the current decision.

This kind of reasoning has been observed in a variety of situations. Sunk costs influence personal decisions (Arkes & Blumer, 1985; Brockner & Rubin, 1985; Garland & Newport, 1991), financial decisions in business (Garland, 1990; Staw, 1976), evaluations of employees' performance (Bazerman, Beekun, & Schoorman, 1982), and competitive behavior (Teger, 1980). These diverse situations are all cases in which the mere fact of an initial investment—usually money or time—can influence a decision maker's subsequent behavior, although future resources could be used more effectively in another manner.

Factors Affecting the Sunk Cost Effect

An obvious prediction that is consistent with the sunk cost effect is that the greater the initial investment, the stronger the effect should be; for example, one should be more likely to continue taking cello

lessons if the cello cost $1,000 than if it cost $100 or nothing. Experimental evidence has supported this prediction (Arkes & Blumer, 1985; Garland, 1990), although Garland and Newport (1991) have found that the absolute amount of the initial investment is less important than the relative cost compared with an individual's total available resources. The importance of relative costs is consistent with prospect theory (Kahneman & Tversky, 1979), according to which future losses (here the displeasure produced by continuing the lessons) result in a smaller decrease in subjective value as the amount already lost (i.e., the cost of the cello and the initial lessons) increases (see also Thaler, 1980).

Situational factors may also come into play. For example, Staw (1976) and Whyte (1993) reported that sunk cost behavior was more frequent when the decision maker felt personally responsible for any negative consequences engendered by the original plan of action. In contrast, Simonson and Nye (1992) found that accountability for one's decision decreased susceptibility to the sunk cost effect. Simonson and Nye's results were likely influenced by their use of business students, many of whom had been formally trained to avoid the sunk cost effect (cf. Larrick, Morgan, & Nisbett, 1990). Finally, Arkes and Blumer (1985) found no consistent relationship between sunk cost behavior and whether oneself or a third party made the decision. These studies indicate that the effect of personal involvement on behavior in such situations is still unclear.

Are There Rational Reasons for Sunk Cost?

The decision to continue with a plan because of sunk costs is usually viewed as a cognitive bias (Baron, 1990). Regardless of the arguments on which the decision is based, if all other factors are equal, then greater utility is achieved by ignoring past costs and attending solely to future consequences. However, there are circumstances in which sunk cost behavior may in fact be rational in the sense of increasing overall utility.

People typically justify sunk cost behavior because they do not wish to appear wasteful (Arkes & Blumer, 1985; Frisch, 1993; Teger, 1980). Because it is adaptive not to waste resources, a rule against waste is generally a good rule to have. However, it may also be maladaptive, as when the rule is over-

generalized to situations in which the resources have already been sunk and remaining resources might be more useful elsewhere (Baron, 1990). Such reasoning attends to irrelevant factors and is therefore counterproductive from a utilitarian perspective.

However, a number of potentially rational reasons might underlie the apparently irrational inclination not to waste irretrievably lost resources. Three such reasons are the opportunity to learn a lesson, punishment for making a bad decision, and the desire to appear to be a consistent decision maker. Below we address each of these reasons for continuing a failed plan.

One reason why individuals might continue with a failed plan (like playing an instrument despite not enjoying it) is to teach themselves that next time they should think carefully before making an expensive purchase. This argument is potentially rational because it may lead to improved future decisions. However, it implies that the decision maker has two "selves," one a teacher and the other a learner. A number of decision theorists have portrayed the decision maker as having multiple selves (e.g., Elster, 1986; Thaler & Shefrin, 1981). For example, Thaler and Shefrin's (1981) theory of self-control describes the decision maker as consisting of both a "myopic doer," who executes decisions but is influenced only by short-term consequences, and a "farsighted planner," who is concerned with lifetime utility. Self-control is achieved when the planner persuades the doer to act in accordance with long-term goals.

The learn-a-lesson argument for sunk cost behavior is similar, but it implies a teacher and a learner, rather than a planner and a doer. Such an argument is even more plausible if the teacher and learner actually are two people, especially in situations where lesson teaching is clearly appropriate, such as a parent teaching a child. Children's decision making is subject to a number of cognitive biases (Baron, Granato, Spranca, & Teubal, 1993; Klayman, 1985), which are unlikely to be corrected by mere experience (Baron, 1990). Hence, parents can play an important and potentially effective role in teaching their children how to be better decision makers.

A second reason for continuing a failed plan is to endure the failed outcome as penance for making a bad decision. Punishment decreases the likelihood that the decision maker will make the same re-

sponse in the future; therefore, it would also serve the function of teaching the decision maker a lesson. Like the learn-a-lesson argument, punishment implies a decision maker with learner and teacher–punisher components. Unlike the learn-a-lesson reasoning, however, punishment can deliver retribution with no explicit explanation for why the punishment is necessary (other than that the decision maker deserves it). Thus, punishment is a rational defense for attending to sunk costs if it provides deterrence against future bad decision making but not if it provides mere vengeance. Consequently, continuing in order to learn a lesson should be more effective at changing the learner's behavior than continuing as penance.

A final line of reasoning for continuing the plan is to present the appearance of a good, consistent decision maker. Changing a course of action is often interpreted as meaning that the original decision was poorly made—an admission of error—and consistency may bring admiration from others (Staw, 1981; Staw & Ross, 1987). Thus, this argument can be viewed as rational because of the high subjective utility associated with a favorable self-presentation (cf. Goffman, 1959). Like the learn-a-lesson and punishment arguments, the consistency argument also implies a multiple-self decision maker (i.e., an actor and an evaluator).

In addition to its influence on the opinion of others, continuing a plan already underway also allows one to view oneself as a good decision maker; that is, escalation is often motivated by self-justification (Bobocel & Meyer, 1994; Brockner, 1992). This would be the case if one adhered to a rule always to finish what one has started. However, such a personal rule (cf. Ainslie, 1986) would not be affected by situational variables shown to influence sunk cost behavior, such as the amount of the initial investment.

These potentially rational reasons for continuing an original plan will not apply to all sunk cost situations. For example, Arkes and Blumer (1985, Experiment 6) showed that most subjects would choose to eat the more expensive of two previously purchased, identical TV dinners. In this situation, more cannot be learned by eating the more expensive dinner, nor does one appear more consistent by doing so. Such rationales are feasible only in situations where one has invested an initial cost and is then faced with the options of continuing the original plan or switching to a more cost-effective alternative. However, one is infrequently forced to choose between two alternatives that are identical except for their cost; on the other hand, there are frequently times when changing circumstances require a choice between one's original plan and trying something different.

Experimental Overview

We conducted three experiments to explore potentially rational justifications for sunk cost behavior in such "continue-or-switch" situations. In each experiment, participants read a series of scenarios and rated five possible responses, including the normative response (switch to a more cost-effective alternative), the traditional sunk cost argument (the *waste* response), and the learn-a-lesson, punishment, and consistency arguments described above. These last four responses all involved a continuation of the original plan. We addressed the effect of a number of factors on the attractiveness of each of these responses. In Experiment 1, we manipulated the care with which the original decision had been made, the amount of the initial investment, and whether a parent had an opportunity to teach a lesson to a child. In Experiments 2 and 3, we investigated the role of personal responsibility and the effect on sunk cost behavior of the lesson learner's age and relationship to the teacher. If sunk cost (i.e., continuation) behavior occurs for some or all of these reasons, it should be more common in some situations than others.

Experiment 1

Study participants rated the attractiveness of each response for a number of scenarios. These scenarios varied as a function of three variables: Whether (a) the decision maker was portrayed as an adult acting alone or as a parent deciding with a child; (b) the decision was made carefully or carelessly; and (c) the initial investment was relatively large or small. We make specific predictions (see Table 1) for each of these variables.

The importance of learning a lesson is predicted to increase as the amount of the initial investment increases because it is most important to avoid repeating previous mistakes in high-stakes situations. The decision care variable is also clearly related to the issue of learning how to make good decisions, as decisions that are made thoroughly and carefully will on average produce better outcomes than those that are not. Thus, we predict that the learn-a-lesson response will receive higher ratings when the decision has been made carelessly because future decision making could thereby be improved more than when the decision was made well but turned out poorly anyway.

Whether the decision maker is an adult alone or a parent making a decision on behalf of a child is also directly relevant to the learn-a-lesson response. If lesson learning is used as a justification, then one should be more likely to continue with a suboptimal plan in order to teach a child a lesson than to teach oneself a lesson. Teaching oneself a lesson highlights the paradoxical nature of a multiple-self representation of decision making (Thaler & Shefrin, 1981): oneself presumably already knows the lesson (otherwise, how could one teach it?), whereas a child may not yet have had the opportunity to learn the lesson.

If punishment serves the same function as learning a lesson—that is, improving future behavior—then this response should also be affected by the level of initial resources, decision care, and whether the decision maker is a single person or a parent and child together. With regard to the amount of the initial investment, for example, an analogy can be drawn from tort law, in which greater deterrence (accomplished through the awarding of damages) is desired for acts resulting in relatively more harmful consequences (Landes & Posner, 1987). Because losing a large initial investment in a sunk cost situation is more harmful than losing a small investment, the attractiveness of punishment for deterrent purposes would be greater in the former case. However, we predict that the punishment response will be less attractive than the learn-a-lesson response overall because the latter contains a clearer explanation of the future benefits that will accrue.

The level of resources already invested is also predicted to affect ratings of the consistency response, on the grounds that it is more important to present an image as a competent decision maker when the stakes are relatively high than when little is at risk (Staw, 1981). With regard to decision care, carelessly made decisions might produce a negative image of the decision maker and thereby create greater self-presentation demands; on the other hand, a careless decision might lessen self-presentation demands, by providing a convenient excuse for the poor outcome that does not necessarily reflect a stable behavior pattern. Hence we make no prediction relating the care with which a decision was made to the consistency rationale. As with decision care, there are competing hypotheses for the expected effect of the decision maker on the consistency argument. On the one hand, children are less cognizant of self-presentation demands than adults (Goffman, 1959), making this argument less attractive in the parent–child condition. Yet parents may be more concerned with creating a favor-

Table 1. *Predicted Effect of Amount of Invested Resources, Identity of Decision Maker, and Decision Care on Five Different Responses*

Variable	Response				
	Normative	Waste	Lesson	Consistency	Punishment
Invested resources					
High		✓	✓	✓	✓
Low	✓				
Decision maker					
Parent and child		—	✓	—	✓
Adult alone	✓	—		—	
Decision care					
Carefully	✓	—		—	
Carelessly		—	✓	—	✓

Note. Check marks indicate the level of each variable for which a response is predicted to receive relatively higher ratings. Dash indicates there was no prediction for how that variable should influence ratings of that response.

able presentation for their children than for themselves. Therefore, we make no prediction about the effect of who the decision maker is on the consistency rationale.

In keeping with the basic sunk cost effect demonstrated in previous research (e.g., Arkes & Blumer, 1985; Garland, 1990), we predict that the normative response will be rated lower when a high level of resources has already been invested than when the initial investment was small and that the opposite will be true for the waste response. In addition, the normative response will receive lower ratings when a parent and child are making the decision together or when the decision was made carelessly. The attractiveness of various reasons for continuing the plan is predicted to increase in those circumstances, and the attractiveness of an option tends to decrease as the attractiveness of a competing alternative increases (Baron, 1994). Finally, we predict that the attractiveness of the waste response will not be affected either by the identity of the decision maker or by the care with which the decision was made. The desire not to waste resources should depend solely on the amount of resources that has already been invested and therefore is independent of these other two variables.

Method

Participants

The participants were 80 students from Louisiana State University who received course credit.

Materials and design

Eight scenarios were developed (summarized in the Appendix), each of which describes a situation in which a decision maker has invested resources (either time or money, or both) in an initial plan of action. Subsequently, circumstances or preferences change, forcing the decision maker to choose between continuing the original plan and switching to an alternative plan. For each scenario, eight conditions were constructed by crossing three binary variables: (a) amount of resources already invested (high or low); (b) decision care (whether the decision had been made "after careful consideration" or "on the spur of the moment"); and (c) the identity of the decision maker (whether the decision was made by an adult alone or by a parent and child together).

For example, see Scenario 1 with alternate versions in parentheses below:

Agatha (Agatha's young daughter Becky) decides after careful consideration (on the spur of the moment) that she wants to take cello lessons. Agatha spends $1,000 ($100) on a beginner cello and an additional $200 ($40) on the first 3 months (1 month) of cello lessons. After 3 months (1 month) of lessons, Agatha (Becky) realizes that she no longer enjoys the cello and wants to stop taking lessons. It is almost certain that if she signs up for more lessons, she will not enjoy them and will never enjoy playing the cello. What should Agatha do?

Each participant saw one version of each of the eight scenarios. Participants were randomly assigned to one of eight counterbalancing conditions, such that each participant read a set containing all eight scenarios and all combinations of the three factors described above. Across sets, the scenarios were always presented in the same order; the order of the conditions within each set was varied according to a Latin-square design, such that each condition appeared in each ordinal position one eighth of the time. Across all participants, all 64 combinations of scenarios and conditions were presented. The design is illustrated in Table 2.

Procedure

After reading each scenario, participants rated the desirability of all five possible responses, with a 10-point scale. Previous research on the sunk cost effect (e.g., Arkes & Blumer, 1985; Larrick et al., 1990) has required participants simply to choose one response from two alternatives (i.e., either continue the original plan or switch to the new alternative).

Table 2. *Design of Experiment 1*

	Scenario							
Set	1	2	3	4	5	6	7	8
A	psh	pch	psl	pel	ash	ach	asl	acl
B	pch	psl	pel	ash	ach	asl	acl	psh
C	psl	pel	ash	ach	asl	acl	psh	pch
D	pel	ash	ach	asl	acl	psh	pch	psl
E	ash	ach	asl	acl	psh	pch	psl	pel
F	ach	asl	acl	psh	pch	psl	pel	ash
G	asl	acl	psh	pch	psl	pel	ash	ach
H	acl	psh	pch	psl	pel	ash	ach	asl

Note. Letters refer to the level of each of three variables. Identity of the decision maker (parent-child = p and adult alone = a), decision care (carefully = c and carelessly = s), and amount of resources invested (high = h and low = l).

Table 3. *Percentage of Trials Receiving Highest Rating and Mean Ratings With Standard Deviations for Each Response*

| | Experiment | | | | | | | | |
| | 1 | | | 2 | | | 3 | | |
Response	%	*M*	*SD*	%	*M*	*SD*	%	*M*	*SD*
Normative	61.7	6.88***	3.14	44.9	5.66***	3.26	62.1	7.07***	3.06
Waste	15.6	4.76***	3.03	15.2	4.60	3.07	12.6	4.70***	2.97
Learn a lesson	8.9	3.79***	2.84	21.5	4.58***	2.99	12.9	4.02***	2.90
Consistency	8.8	3.37***	2.57	14.6	3.87***	2.82	8.0	3.36***	2.49
Punishment	5.0	2.28***	2.06	3.8	1.82***	1.75	4.4	2.50***	2.15

Note. Means were compared with the next lowest response or with the minimum rating of 1 if there was no lower response.

***p < .001 to correct for multiple comparisons.

We used a rating method to allow an evaluation of the different reasons that participants might rely on in choosing to continue the original plan. The rating scale has the additional advantage of permitting parametric analyses. The 10-point rating scale was anchored at 10, "a very good response, the one that [the decision maker] should definitely follow," and 1, "a very bad response, that [the decision maker] should definitely not follow." Use of the scale was unconstrained; that is, participants were not required to allocate a predetermined number of points across the five responses, and multiple responses could receive the same rating.

The five responses—which varied slightly depending on the identity of the decision maker—were *normative* (e.g., "Agatha should [let Becky] stop taking cello lessons because it would be a waste of time and money to take more lessons that she won't enjoy"); *waste* (e.g., "Agatha should [make Becky] continue with the lessons because otherwise she will have wasted the money and time already spent"); *learn-a-lesson* (e.g., "Agatha should [make Becky] continue lessons to teach her[self] that next time she should be more careful about what hobbies she selects for herself"); *punishment* (e.g., "if she was foolish enough to select a hobby that she doesn't enjoy, she deserves to suffer by continuing with her lessons"); and *consistency* (e.g., "if she stops taking lessons, that would mean she made a bad decision in deciding to take cello lessons. . . . If it was the right decision then, it is still the right decision"; modeled after an alternative used by Larrick et al., 1990). The order of the responses was determined randomly for each scenario, with the constraint that each response appear in each ordinal position at least once across the eight scenarios.

After completing the decision scenarios, participants provided information on their frequency of engaging in five sunk cost activities (e.g., continuing to watch a boring movie).[1] The entire questionnaire took 20 to 30 min to complete.

Results

Average ratings of the five responses

Four of the responses involved continuing the failed plan, whereas only one response (normative) involved switching to the new plan. This design may result in an overestimation of the tendency for participants to stick with the original plan, but it is unlikely to affect either the ordering of the four "continue" responses or the effect of the independent variables on the response ratings. Table 3 shows the average ratings of the five responses and the

[1] Participants provided this information in all three experiments. Each activity was rated on a 1-to-10 scale—with high numbers indicating a greater likelihood of abandoning a plan that is no longer optimal—and the five ratings were averaged to form one score. We examined the correlations between these activity scores and ratings of the five types of scenario responses. We predicted that activity scores would be negatively correlated with waste response ratings and positively correlated with normative response ratings. We did not obtain any consistent correlations between participants' reports of everyday behavior and their questionnaire responses.

Table 4. *Mean Response Ratings With Standard Deviations for Experiment 1*

	Response									
	Normative		Waste		Lesson		Consistency		Punishment	
Variable	M	SD	M	SD	M	SD	M	SD	M	SD
Invested resources										
High	6.50	3.24	5.12**	3.05	4.05**	2.92	3.55**	2.63	2.49**	2.28
Low	7.27**	2.98	4.39	2.96	3.52	2.73	3.20	2.51	2.07	1.79
Decision maker										
Parent and child	6.28	3.25	4.86	3.09	4.56**	3.02	3.57**	2.73	2.61**	2.33
Adult alone	7.49**	2.89	4.66	2.96	3.01	2.40	3.16	2.39	1.94	1.68
Decision care										
Carefully	6.92	3.13	4.66	2.86	3.62	2.78	3.33	2.54	2.25	2.06
Carelessly	6.84	3.14	4.86	3.18	3.96*	2.89	3.40	2.61	2.31	2.06

*$p < .07$. **$p < .05$.

percentage of trials on which each response was rated highest.

The normative response received the highest mean rating (6.88) and was also assigned the highest rating on the majority of trials (63%). The waste response was given a lower average rating (4.76), and the learn-a-lesson response a still lower rating (3.79). The consistency (3.37) and punishment (2.28) responses received the lowest ratings. The mean rating for each response was significantly different from all other responses ($p < .001$).[2] In addition, the order of the four "continue" responses (determined by their mean ratings) was the same both for trials where the normative response received the highest rating and for trials where one of the "continue" responses received the highest rating.

The normative response was negatively correlated with each of the four "continue" responses (r s ranged from $-.30$ to $-.51$, $n = 80$, p s $< .001$); whereas the four "continue" responses were positively correlated with one another (r s ranged from .41 to .50, p s $< .001$). Although the learn-a-lesson response is conceptually similar to the punishment response, the correlation between these two ($r = .47$) was not significantly higher than that between any other two "continue" responses. This pattern of correlations was replicated in Experiments 2 and 3.

Effects of decision care, decision maker, and amount of invested resources

Each of the five responses was used as the dependent variable in an analysis of variance (ANOVA) model containing the between-subjects variable of counterbalancing condition (i.e., stimulus set) and the within-subjects variables of scenario, decision maker (parent–child or adult alone), decision-care (carefully or carelessly), and level of invested resources (high or low). Although we did not predict any interactions, two-way interactions among the latter three variables were included in the models.[3] Mean ratings are presented in Table 4.

---

[2] For each participant, we also calculated the percentage of trials on which that participant rated each response highest and then averaged these percentages across participants (see Table 3). In all three experiments, pairwise comparisons of the percentages showed fewer significant differences than the pairwise comparisons of means, although the overall patterns were the same. Most importantly, the percentage measure also revealed systematic differences among the responses; for example, the learn-a-lesson response was rated highest on a significantly greater percentage of trials than the punishment response in all three experiments; Experiment 1: $t (79) = 1.87$; $p < .07$; Experiments 2 and 3: t s > 4.3; p s < .001. We limit our discussion to the analyses of response means.

[3] The counterbalancing variable of stimulus set did not have an effect in any of the three experiments. An additional counterbalancing measure, presentation order—used only in Experiment 3—also did not affect participants' responses. However, in all three experiments there was a main effect of scenario. Separate analyses of variance (ANOVAs) within each scenario indicated stronger manipulation effects for some scenarios than others; however, these effects were not systematic across experiments. Other researchers (e.g., Frisch, 1993) have also demonstrated differences between alternate sunk cost scenarios. As we were interested in generalities instead of responses to particular cases, we do not

As in other research (e.g., Arkes & Blumer, 1985; Garland, 1990), the waste response was rated as more appropriate when a large amount of resources had been sunk into the project than when a small amount had been invested, $F(1, 540) = 15.17$, $p < .0001$. Conversely, the normative response was viewed as more appropriate when a small amount of resources had been sunk into the project than when a large amount had been invested, $F(1, 540) = 13.95$, $p < .0001$. As predicted, the normative response was also viewed as more appropriate when the decision maker was an adult acting alone than when a parent and child acted together, $F(1, 540) = 32.91$, $p < .0001$, but contrary to our hypothesis, it was not affected by decision care ($F < 1$). Consistent with our predictions, ratings of the waste response were affected by neither decision care nor the identity of the decision maker.

The mean rating for the learn-a-lesson response for decisions made by a parent and child together was higher than for decisions made by an adult alone, $F(1, 540) = 73.39$, $p < .0001$. Figure 1 shows learn-a-lesson ratings for decisions made carefully or carelessly, plotted as a function of level of resources already invested. There was a main effect of invested resources, indicating higher ratings with a high level of investment, $F(1, 540) = 8.46$, $p < .01$. Also of interest, although the main effect of decision care was only marginally significant, $F(1, 540) = 3.53$, $p < .07$, there was an interaction between decision care and level of resources invested, $F(1, 540) = 6.78$, $p < .01$, indicating that for a high level of sunk cost (but not for a low level), the lesson response was rated higher for decisions made carelessly.

---

discuss scenario effects further. Higher order interactions were not included in the ANOVA models because they were not predicted by prior theory and would not be readily interpretable. For this and subsequent experiments, we also computed logistic regression analyses where the dependent variable for each of the five responses was whether that response received the highest rating. Because some of the responses virtually never received the highest rating of the five responses (e.g., punishment), we report only the ANOVA results on mean ratings, which showed the same general pattern of effects as the logistic regression analyses.
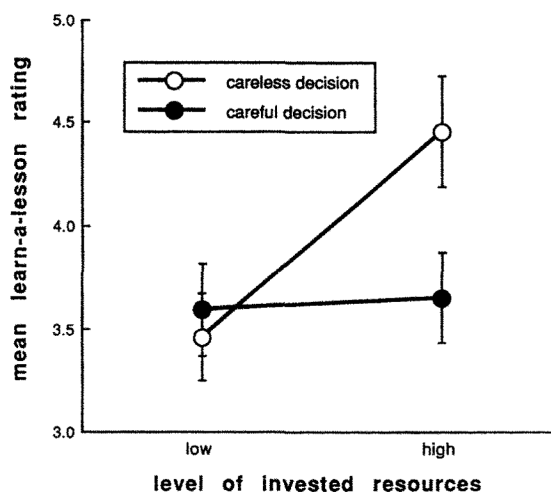

Figure 1. Mean ratings of the learn-a-lesson response with standard errors for Experiment 1 shown as a function of level of invested resources (high or low) and decision care (decision made carelessly or carefully)

Consistency and punishment ratings both showed main effects indicating higher ratings for a high versus a low level of invested resources and for a parent and child acting together versus an adult alone, F s$(1, 539) > 4.7$, p s $< .03$.

This general pattern of results was not limited to the parent–child scenarios. For example, within the adult-alone condition, the learn-a-lesson response was still perceived as more appropriate than the punishment response, $t(316) = 8.26$, $p < .001$, and roughly as good as the consistency response, $t(316) = -1.19$, $p > .2$; in addition, it was rated higher for careless than for careful decisions, $F(1, 227) = 10.57$, $p < .005$.

Discussion

These results clearly demonstrate the classic sunk cost effect. All four responses advocating the continuation of the original plan received higher ratings when a relatively large amount of resources had already been invested. In contrast, the normative response advocating termination of the original plan received higher ratings when a relatively small amount of resources had been invested. Reasons underlying the sunk cost effect can be determined by examining the ratings of the four continuation responses. Although the waste response had the highest rating of the four, indicating that a desire

to avoid wastefulness is the primary motivation behind sunk cost behavior (Frisch, 1993; Teger, 1980), the learn-a-lesson response received the next highest rating, indicating that it too can be an important justification.

In addition, the attractiveness of the learn-a-lesson response was influenced by two variables predicted to affect the appropriateness of this argument. Learning a lesson was viewed as more appropriate when a parent had an opportunity to teach a lesson to a child than for an adult alone and when the original decision had been made carelessly rather than carefully. A child is less likely to have already learned the lesson and has a longer life expectancy in which to benefit from the lesson compared with an adult. In addition, the parent–child situation provides social roles encouraging lesson teaching and also makes explicit the two selves involved in the decision. The normative response to switch to a new plan is correspondingly less attractive in parent–child situations.

Learning a lesson is more effective when a decision has been made carelessly because the decision maker has the opportunity to learn to be more careful. There may be no way to improve future decisions if the original decision was already made carefully. This suggests that participants are sensitive to the manner in which the original decision was made, as well as to the relationship between decision quality and decision outcome. Furthermore, a careless decision is of little consequence if it does not waste many resources; thus, learning a lesson is more important when careless decisions have incurred large costs.

This result suggests a continuum from trivial decisions (those involving few resources) to monumental decisions (involving extensive resources). Learning to make trivial decisions carefully is not worth the trouble, as it will save few future resources. In making a decision, factors other than the accuracy of the decision's outcome—such as the time and effort expended in making the decision—also need to be considered (Baron, 1994; Kunda, 1990; Shugan, 1980). In contrast, learning to make important decisions carefully is well worth the effort. Thus, although learning a lesson is influenced by the waste of resources, it is the waste of future resources, not past ones, that drives this argument. Consequently, what appears on the surface to be traditional sunk cost (and therefore irrational) reasoning can be justified because of its beneficial effects on future decision making.

However, just as ignoring sunk costs is not always the optimal response, continuing a failed plan in order to learn a lesson will not always be optimal either. In some situations, the cost of continuing will simply be too great. In deciding whether to continue or switch, one must weigh the costs and benefits of the two options (Staw & Ross, 1987). When the initial investment is large and the original decision was made carelessly, the long-term benefit of learning a lesson will often outweigh the cost of investing additional resources. Whether it does in a particular situation will depend on factors such as the amount of additional resources required and the likelihood of making a similar decision in the future (i.e., the probability that one will be able to use the lesson learned).

The learn-a-lesson response received higher overall ratings than both the punishment and consistency arguments. It was viewed favorably both for a parent deciding with a child and for an adult acting alone. In addition, the effect of decision maker was largest for this response, and it was the only response influenced by decision-making level of care, which otherwise had no effect on ratings of the various responses. These results point to the importance of learning a lesson in explaining sunk cost behavior. Consequently, in Experiments 2 and 3, we explored the learn-a-lesson response further by focusing on additional variables that might have an impact on its attractiveness as an argument for continuing a failed plan.

Experiment 2

In Experiment 1, the age of the child was not specified beyond characterizing the child as "young" to indicate a minor. If it is beneficial to teach a child a lesson, it seems reasonable to assume that this benefit will increase as the child's age decreases. Younger children's behavior patterns are less fully determined, so they are generally viewed as more malleable; they have had less opportunity to be taught the lesson already; and simply because they are younger, any learning will have a potentially greater payoff spread out over their expected lifespan. In Experiment 2, we manipulated the age of the child involved in a sunk cost situation to test the hypothesis that the learn-a-lesson response is more desirable for relatively young children. It can be argued that the learn-a-lesson response is actually less

desirable for very young children because they may be too young to learn to think more carefully before making future decisions. However, Elias, Branden-Muller, and Sayette (1991) have found that even elementary schoolchildren can learn how to improve their decision making. Consequently, the scenarios in Experiment 2 described children who were elementary school age or older.

Because of the conflicting literature on the effect of personal responsibility on the sunk cost effect (Arkes & Blumer, 1985; Staw, 1976; Whyte, 1993), we also varied whether the participant took the perspective of the decision maker or merely made a recommendation for a decision maker portrayed as a hypothetical other.

Method

Participants. The participants were 79 students at Louisiana State University who received extra course credit.

Materials and design. The parent–child versions of the four scenarios for which the learn-a-lesson effect was strongest in Experiment 1 were used. There were four different versions of each scenario, which resulted from crossing the two variables of amount of initial investment (same as Experiment 1) and age of the learner (7 or 8 years old vs. 15 or 16). Two younger ages and two older ages were used to make the within-subject manipulation of the learner's age less transparent. The same counterbalancing procedures were used as in Experiment 1, yielding four separate stimulus sets that each contained all four experimental conditions.

The perspective of the decision maker was added as a wholly between-subjects variable. Half of the participants were instructed to imagine themselves as the decision maker. These participants read scenarios referring to "you"; for example, "Your daughter Becky wants to take cello lessons." They answered the question, "What would you do?" in that situation. The remaining participants read scenarios describing another person; for example, "Agatha's daughter Becky wants to take cello lessons." They were asked, "What should Agatha do?"

Procedure. The procedure was the same as in Experiment 1 with the sole exception that participants read and rated four instead of eight sunk cost scenarios.

Results and Discussion

Average ratings of the five responses. The overall rankings of the responses were the same as in Experiment 1 (see Table 3). The normative response received the highest mean rating (5.66) and was also assigned the highest rating more often than any other response (45%). The waste response was given a lower average rating (4.60) and was rated highest on 15% of the trials. The learn-a-lesson response was given a similar average rating (4.58) and was assigned the highest rating on 22% of the trials, suggesting that this argument was the primary motivation behind sunk cost behavior. The comparatively high ratings assigned to the lesson response in this experiment were likely due to the fact that it had only parent–child scenarios. The consistency and punishment responses received average ratings of 3.87 and 1.82, respectively.

Effects of perspective, age of learner, and amount of invested resources. Each of the five responses was used as the dependent variable in an ANOVA containing the between-subjects variable of perspective of the decision maker (decision made by self or other) and the within-subjects variables of age of child and level of invested resources, as well as the two-way interactions among all three variables. Table 5 shows the mean response for each level of perspective, age, and amount of resources.

As in Experiment 1, the normative response was rated as more appropriate when a small amount of resources had been sunk into the project than when a large amount had been invested, $F(1, 229) = 6.26$, $p < .02$; whereas the waste response was rated as more appropriate when a large amount of resources had been invested compared with a relatively small amount, $F(1, 229) = 6.04$, $p < .02$. The normative and waste responses did not show any effect of perspective or age of child ($Fs < 1$). However, the waste response showed an interaction between perspective and age of child, $F(1, 229) = 3.95$, $p < .05$, indicating higher ratings for an older child in the "other" condition (5.09 vs. 4.49), but higher ratings for a younger child in the "self" condition (4.62 vs. 4.22).

As in Experiment 1, ratings for the learn-a-lesson response were higher for a high level of invested resources, $F(1, 229) = 9.59$, $p < .01$. There was no significant main effect of perspective, $F(1, 74) = 1.10$, $p > .29$, or of age of child ($F < 1$). Thus, participants

Table 5. *Mean Response Ratings With Standard Deviations for Experiment 2*

	Response									
	Normative		Waste		Lesson		Consistency		Punishment	
Variable	M	SD	M	SD	M	SD	M	SD	M	SD
Invested resources										
High	5.27	3.16	4.93**	3.22	5.01**	3.04	3.88	2.83	1.82	1.78
Low	6.03**	3.32	4.27	2.89	4.14	2.88	3.87	2.82	1.82	1.72
Age of child										
Younger	5.61	3.34	4.56	3.06	4.50	2.94	3.84	2.87	1.79	1.74
Older	5.70	3.18	4.65	3.10	4.65	3.04	3.91	2.78	1.82	1.77
Perspective										
Other	5.46	3.06	4.79	2.97	4.78	3.00	3.89	2.74	1.89	1.75
Self	5.84	3.45	4.42	3.17	4.38	2.98	3.86	2.90	1.75	1.76

**p < .05.*

did not view learning a lesson as more appropriate for younger children. Finally, consistency and punishment ratings did not show any effects of perspective, age of child, or level of invested resources, $Fs(1, 229) < 2.3$, $ps > .13$.

The sole effect of the self–other distinction was an interaction between perspective and age of child on ratings of the waste response. Although this interaction seemed to suggest that perspective is important when it comes to dealing with children (i.e., one's own are treated differently from someone else's), we caution against drawing any conclusions because of our failure to replicate it in ratings of any of the other responses. As Arkes and Blumer (1985) noted, it is difficult to construct a valid manipulation of personal involvement with questionnaire studies, as opposed to observations of genuinely consequential behavior. When responsibility is manipulated by varying whether the present decision maker or someone else made the usually hypothetical initial investment, the decision maker's decision to continue or switch is affected (e.g., Staw, 1976; Staw & Ross, 1977; Whyte, 1993); but when responsibility is manipulated by varying the participant's perspective (i.e., whether the present decision maker is the participant or a third party), it has less or no effect (Arkes & Blumer, 1985). Because learning a lesson would not be feasible in the former case—where the original decision maker may not be involved in the current decision—the present experiment manipulated the decision maker's perspective. Our failure to demonstrate an effect of perspective reflects the subtle nature of personal responsibility effects in this domain.

Although there is reason to predict that learning a lesson would be especially beneficial for a younger child, in Experiment 2 we found no effect of the child's age. This failure may indicate that there are other reasons beside a child's age as to why parent–child scenarios are particularly likely to evoke learn-a-lesson arguments. This possibility is explored in Experiment 3.

Experiment 3

An alternative explanation for why it is more important to teach a lesson to a child rather than oneself is that parents have a fiduciary responsibility to children that involves teaching lessons. In contrast, adults are not usually viewed as having a social responsibility to teach lessons to themselves. This explanation implies that teaching a lesson would be rated highly in other teacher–learner situations, even if the learner were not a young person. Although teaching a lesson might occur in adult–adult relationships as well, the learn-a-lesson rationale should be particularly attractive in a parent–child relationship, owing to parents' special responsibilities.

Experiment 3 was designed to test this hypothesis. Four scenarios from Experiment 1 were used. An additional four scenarios that involved adult learners were created. Therefore, one experimental factor was the type of scenario: half the scenarios described a situation that contained either a child learner or an adult alone, and half described situations involving either an adult learner or an adult alone. Thus, within each scenario, a second factor

was whether a teacher had an opportunity to teach a lesson to a learner or whether a single adult acted alone. A third within-scenario factor was the level of cost already sunk into the project. Participants responded to all eight scenarios and saw all combinations of the three factors.

Method

Participants. Participants were 86 undergraduate students at Louisiana State University who received extra course credit.

Materials and design. Eight scenarios were used. Four had been used in Experiment 1; the remaining four scenarios were new ones that described either a single adult or 2 adults making a decision together (see Appendix). For example, see Scenario 12 with alternate versions in parentheses below:

> Leroy is a private in the Army. His commanding officer, Jim, orders Leroy to figure out (and report to him) the best way to catalog all of the weapons in a large storehouse. . . . He must work alone, and without supervision (or must have his commanding officer's approval). . . . After working on the job for 2 weeks (or 2 months), Leroy figures out a new method that will allow him to do the whole job, from start to finish, in only 4 weeks. It will take another 2 months to finish doing the inventory the original way. (Leroy needs his commanding officer Jim's approval to change methods.) What should Leroy (or Jim) do?

It was not possible to construct plausible adult–child and adult–adult versions of the same scenario. Consequently, rather than comparing an adult–adult scenario directly with an adult–child scenario, each was compared with its own adult-alone control condition. The perspective in all scenarios was such that participants were asked, "What should [the decision maker] do?" as in Experiment 1. There were four versions of each scenario, which resulted from crossing the factors of the identity of the decision maker (teacher and learner versus adult alone) and the amount of the original investment (high versus low).

The two scenario types—parent–child/adult alone and adult–adult/adult alone—were presented in alternating order. The presentation order of the remaining conditions was counterbalanced by a partial Latin-square design, which yielded four different questionnaire sets. Although the order of the scenarios within each set was fixed, half of the participants completed each set in forward order, and half completed the set in backward order (see Footnote 3 on p. 257). Thus, each of the eight conditions appeared in each ordinal position an equal proportion of the time. Each participant responded to all eight scenarios and all combinations of the three variables.

Procedure. The procedure was identical to that used in Experiments 1 and 2.

Results and Discussion

Average ratings of the five responses. The overall ranking of the responses was the same as in Experiments 1 and 2 (see Table 3). The normative response received the highest mean rating (7.07) and was also assigned the highest rating on the majority of trials (62%). The waste response was given a lower average rating (4.70), and the learn-a-lesson response a still lower rating (4.02). The consistency (3.36) and punishment (2.50) responses received the lowest average ratings.

Effects of decision maker, type of scenario, and amount of invested resources. Each of the five responses was used as the dependent variable in an ANOVA containing the within-subjects variables of type of scenario (adult–adult/adult alone or parent–child/adult alone), decision maker (teacher–learner or adult alone), and level of invested resources, and two-way interactions among these variables. If the parent–child relationship is especially conducive to teaching a lesson, compared with adult–adult relationships, then learn-a-lesson ratings would show a significant interaction between scenario type and decision maker, indicating a greater effect of having a teacher and a learner (opposed to an adult alone) when the teacher and learner are parent and child. Conversely, ratings of the normative response should yield an interaction in the opposite direction. Mean responses are shown in Table 6.

As in Experiments 1 and 2, the waste response was rated as more appropriate, $F(1, 590) = 5.91$, $p < .02$, when a large (compared with a small) amount of resources had been sunk into the project; and the normative response was rated as more appropriate for a low compared with a high level of invested resources, although the latter effect was only marginally significant, $F(1, 590) = 2.92$, $p < .09$. As in Experiment 1, the normative response was rated as

Table 6. *Mean Response Ratings With Standard Deviations for Experiment 3*

	Response									
	Normative		Waste		Lesson		Consistent		Punishment	
Variable	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*	*M*	*SD*
Invested resources										
High	6.92	3.13	4.91**	3.00	4.31**	2.97	3.49	2.51	2.67**	2.31
Low	7.23	2.99	4.49	2.93	3.74	2.81	3.23	2.47	2.32	1.97
Decision maker										
Teacher and learner	6.62	3.19	4.87	3.04	4.51**	3.05	3.51**	2.57	2.80**	2.38
Adult alone	7.53**	2.87	4.54	2.89	3.54	2.66	3.20	2.40	2.19	1.85
Type of scenario										
Parent-child/control	6.85	3.15	4.82	2.98	4.16	2.89	3.44	2.55	2.55	2.18
2 adults/control	7.30**	2.96	4.58	2.96	3.89	2.91	3.27	2.43	2.45	2.13

**p < .05.

more appropriate when an adult acted alone than when a teacher and learner decided together, $F(1, 590) = 24.71$, $p < .0001$. It was rated similarly for adult–adult/adult-alone scenarios and for parent–child/adult-alone scenarios, $F(1, 6) < 1$. [4] As predicted, there was a significant interaction between type of scenario and decision maker, $F(1, 590) = 3.78$, $p = .05$ (see Figure 2), indicating that although the normative response was always more appropriate for single decision makers, this effect was especially pronounced when the teacher–learner version of the scenario involved a child. Thus, the normative response was least attractive when there was a child learner.

In corroborating the first two experiments, we found that the learn-a-lesson response was viewed as more appropriate for large amounts of invested resources, $F(1, 590) = 11.85$, $p < .001$. It was also rated higher when a teacher had an opportunity to teach a lesson to a learner than when the decision maker was acting alone, $F(1, 590) = 34.65$, $p < .0001$, replicating Experiment 1. As in Experiment 1, this overall pattern of results was not due simply to the teacher–learner condition. Within the adult-alone condition, the learn-a-lesson response was perceived as more appropriate than both the punishment, $t(344) = 10.43$, $p < .001$, and consistency responses, $t(344) = 2.36$, $p < .02$.

The type of scenario had no main effect on learn-a-lesson ratings, $F(2, 6) < 1$, whereas the interaction between type of scenario and decision maker was marginally significant, $F(1, 590) = 3.44$, $p < .07$. As predicted, the difference between the teacher–

learner and adult-alone conditions was larger for the scenarios in which the learner was a child ($M = 4.79$ vs. 3.52, respectively) than for scenarios in which the learner was an adult ($M = 4.22$ vs. 3.56, respectively; see Figure 2). Thus, teaching a lesson is seen as more appropriate for a child learner than for an adult learner. This interaction is consistent with the normative response ratings: the learn-a-lesson argument is most attractive and the normative response least attractive when there is a learner present and that learner is a child.

As predicted, even in the case of adult learners, teaching a lesson is a more appropriate rationale when two people are involved. Pairwise comparisons indicated that the learn-a-lesson response was rated higher in both the parent–child and adult teacher–adult learner conditions than in their respective adult-alone control conditions, F s$(1, 591) > 8.0$, p s $< .01$. In other words, the participants believed that one can teach an old dog a new trick.

The pattern of results for the punishment and consistency responses was very similar to that obtained in Experiment 1. The punishment response received higher ratings for the high relative to the low level of invested resources, $F(1, 590) = 7.47$, $p < .01$; for the consistency response, this effect was only marginally significant, $F(1, 590) = 3.10$, $p < .08$. Like the lesson response, the consistency and punishment responses were both rated higher when a teacher

[4] For the main effect of scenario type, we used the quasi F ratio recommended by Winer (1971) and Clark (1973), which incorporates the variance due to scenarios within each type.

had an opportunity to teach a lesson to a learner than when an adult acted alone, F s(1, 590) > 4.4, p s < .05, but the interactions with scenario type were not significant.

General Discussion

In all three experiments, the normative response was rated lower when a high level of resources had already been invested into the current project compared with a small initial investment; the opposite was true for the waste response. This result replicates the basic sunk cost effect demonstrated by Arkes and Blumer (1985) and others (e.g., Garland, 1990; Thaler, 1980). Although "throwing good money after bad" is usually perceived as an indication of biased or irrational thinking (Baron, 1990), we were able to uncover potentially rational reasons for engaging in sunk cost behavior. In addition to the traditional desire not to waste already expended resources, we explored three factors that might explain perseverance with a plan that is no longer optimal: the opportunity to learn a lesson, the provision

of punishment as a possible deterrent, and the desire to appear consistent.

Across the three experiments, participants' ratings showed a consistent preference for the normative response (i.e., switch to the alternative), followed in descending order by the following arguments for continuing the original plan: waste, learn-a-lesson, consistency, and punishment. Overall, choosing a non-normative response was quite common; across experiments, participants chose to continue the original plan an average of 43% of the time, a figure that is comparable with that found in previous studies (e.g., Arkes & Blumer, 1985; Larrick et al., 1990). More important than the overall frequency of sunk cost behavior, participants' willingness to endorse each possible argument varied as a function of a number of variables. In all three experiments, variables that increased the attractiveness of the various reasons for continuing the original plan made the normative response appear correspondingly less appropriate.

The waste response was affected only by the level of invested resources. Participants' greater disincli-
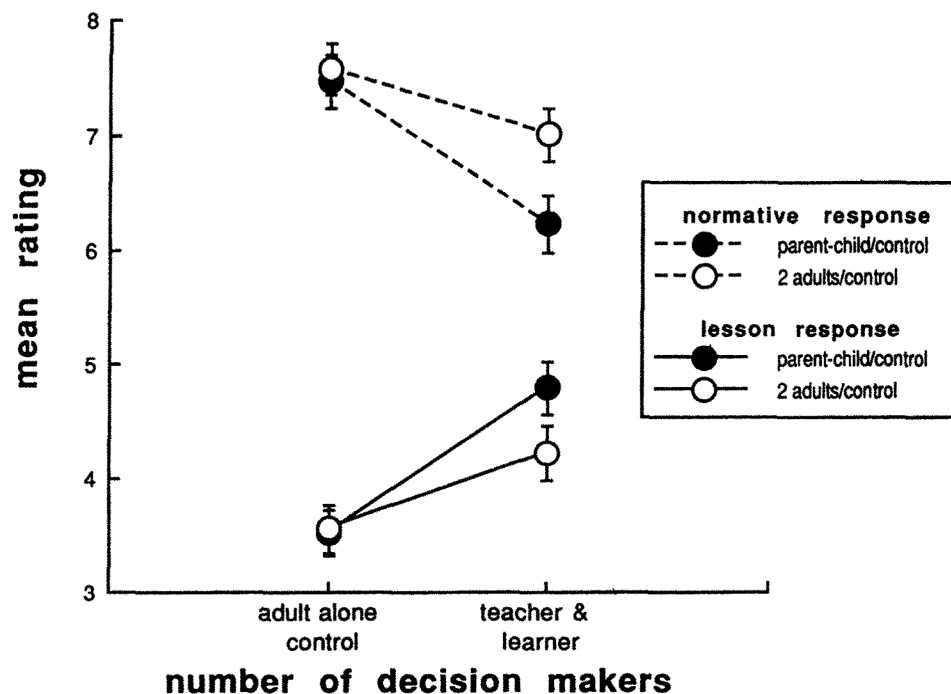

Figure 2. Mean ratings of the learn-a-lesson and normative responses with standard errors for Experiment 3 shown as a function of scenario type (parent–child/adult alone or adult–adult/adult alone) and number of decision makers (adult alone or teacher–learner).

nation to waste additional resources as the amount already invested increased is consistent with prospect theory (Kahneman & Tversky, 1979; Thaler, 1980). However, prospect theory does not predict an effect of variables such as whether the decision maker acted alone, the care with which the decision was made, or the nature of the relationship between teacher and learner. The other three responses were influenced by these variables.

The opportunity to learn a lesson was the strongest of the remaining alternatives. The learn-a-lesson response received higher ratings in situations where learning a lesson would be especially beneficial or appropriate. In Experiments 1 and 3, higher learn-a-lesson ratings occurred when a teacher had an opportunity to teach a lesson to a learner than when someone acted alone. To teach oneself a lesson, one must already know it, making the teaching paradoxical. Nevertheless, this "multiple self" approach (Elster, 1986; Thaler & Shefrin, 1981) is supported by the fairly high ratings received by the learn-a-lesson response compared with two other possible rationales (i.e., punishment and consistency), even when an adult acted alone.

In Experiment 2, the attractiveness of the learn-a-lesson response for teacher–learner scenarios was not influenced by the age of the learner or by whether the subject adopted the perspective of the teacher or of a third party. In Experiment 3, the teacher–learner effect extended to adult learners as well. The implication of these results is that the lesson response is viewed as more appropriate in teacher–learner situations not merely because of the learner's youth, but because of the social role that makes teaching a lesson appropriate. Nonetheless, the learn-a-lesson response was rated as more appropriate when the teacher and learner were parent and child rather than two adults, suggesting that the parent–child relationship is especially conducive to teaching a lesson. Parents' teaching role might be expected to lead them to instruct their children to ignore sunk costs and thereby make more cost-effective decisions (Baron, 1990). However, the results of the present experiments suggest that this tendency can be overridden by the long-term benefits that are derived from forcing a continuation of the original plan—namely, teaching the child a useful lesson that will improve the child's future decision making.

A final influence on ratings of the lesson response was revealed in Experiment 1, which demonstrated that learn-a-lesson ratings were higher when the original decision had involved a large cost and had been made carelessly. In this situation, learning to think more carefully would have a particularly high benefit by improving future decisions and saving large amounts of resources. In such cases, continuing a previous course of action may be rational in the long-term, though more costly in the short-term, because of its effect on future decision making.

The ratings of the consistency and punishment arguments indicate that the effect of having a separate learner is not specific to the learn-a-lesson argument, but also generalizes to other arguments for why decision makers should experience the consequences of their failed decisions. As predicted, the importance of appearing a consistent decision maker increased as the stakes involved increased. The other variable that affected ratings of the consistency response was whether the decision maker acted alone or in concert with another. This result suggests that parents (or other teachers) are more concerned with the learner's appearing consistent than they are with viewing themselves as consistent.

In Experiments 1 and 3, ratings of the punishment argument were also influenced by the amount of the initial investment and whether the decision maker acted alone. As predicted, this response seemed more attractive when the investment was relatively high and the decision maker was acting on behalf of another person. It is interesting to note that the punishment and learn-a-lesson arguments were affected in the same way by these two variables. Teaching a lesson and punishment can effectively serve the same purpose: improving the quality of one's future decisions. Despite this conceptual relationship, ratings of the learn-a-lesson and punishment responses were no more highly correlated than any other pair of responses. Punishment may also involve mere retribution without a clear explication of its deterrent effects on future decision making. The generally higher ratings for learning a lesson than for punishment suggest that the former is a more obvious means of improving future behavior. Additional research is called for that explores how the framing of punishment (i.e., as either a beneficial lesson or as just desserts) would affect the attractiveness of this

argument. A related topic for future research is what reasons subjects themselves would give for sunk cost behavior, if reasons were not provided by the experimenter.

The importance of learning a lesson does not apply to all sunk cost situations. For example, there is no clear benefit to teaching oneself a lesson in cases where someone else made the initial investment or where one is forced to choose between two identical purchases that differ only in the amount that each one cost (cf. Arkes & Blumer, 1985, Experiment 6). However, sunk cost situations frequently present an opportunity for the decision maker to learn a lesson. As Brockner and Rubin (1985) pointed out, "Entrapment—under certain special circumstances at least—can be construed as a phenomenon that can serve constructive ends" (p. 255). In judging the rationality of sunk cost behavior, one needs to weigh the costs and benefits of the options to switch or to continue, where those costs and benefits are not limited to measurable resources such as money and time. In many situations, the costs of continuing will outweigh the potential benefits—for example, if little was invested initially, continuing would exhaust all of one's remaining resources, the initial decision was carefully thought out, and the decision maker is unlikely to encounter a similar situation in the future. On the other hand, if the long-term benefits of learning and implementing the lesson outweigh the immediate costs of sticking to the original plan, then sunk cost behavior does not represent fallacious reasoning.

In commenting on "the rhetoric of irrationality," Lopes (1991) noted an overemphasis in the judgment and decision-making literature on the negative aspects of people's performance and a comparative neglect of situations in which they reason well (see also Christensen-Szalanski & Beach, 1984). What appears to be a bias in the laboratory may be functional behavior in a more realistic context (Funder, 1987; Hogarth, 1981), where a variety of justifications for the behavior can be considered. In general, ignoring sunk costs is an adaptive, cost-effective strategy. Yet what appears to be biased, irrational behavior—such as decreasing utility through attention to irretrievably wasted resources—can be described as "meta-rational" (Jungermann, 1986), assuming the benefits of learning and implementing the lesson outweigh the

costs of sticking to the original plan. However, it raises the interesting question of why continuing a failed plan is the best (or even a good) way to learn to make better decisions in the future. Perhaps one could both abandon the current unsuccessful plan and learn to think more carefully in future decisions. However, we argue that continuing with the unsuccessful plan is a particularly effective means of learning the lesson because the consequences of one's original decision are then "in-kind," thereby providing better feedback than merely noting that the original decision was made badly. In support of this conclusion, other research has shown that decision makers prefer punishments that fit the crime (Beattie & Baron, 1995).

Sunk cost behavior occurs in a variety of domains. Although the present studies, like most demonstrations of the phenomenon, have a laboratory setting, numerous examples can be drawn from real-world behavior. People's tendency to allow past decisions to influence current ones in a suboptimal fashion has been observed in interpersonal relationships (Brockner & Rubin, 1985); in political decision making (Janis, 1982; Teger, 1980); and in real organizational settings, where it can lead to escalation of commitment in financial decisions (Ross & Staw, 1993; Staw, 1981), performance appraisals (Schoorman, 1988), and group policy decisions (Brockner & Rubin, 1985; Janis, 1982). Thus, it is clearly important to understand why such behavior occurs, so that decision making in these and other situations might ultimately be improved.

Staw and Ross (1987) identified a number of procedures for reducing escalation behavior, such as bifurcating initial and subsequent decision making and lessening the consequences of failed decisions. Another technique that they suggested is to provide decision makers with negative feedback. An effective means of doing so would be for the decision maker to experience the negative consequences of the failed decision—in other words, to learn through failure. The importance of negative feedback in skill acquisition has been amply demonstrated in the learning literature (e.g., Bilodeau, 1969), and it has recently been applied to acquiring knowledge in complex domains such as medicine (Clancey, 1987). Thus, experiencing the consequences of a failed decision can be a very effective instructional device

for learning how to make better decisions in the future. [5] This approach is likely to be especially attractive in the early stages of learning, as in training a new employee; a more senior employee ought to have learned the lesson already. Those interested in improving decisions should emphasize all the consequences of continuing a plan, including learning a lesson. Our results suggest that in teaching decision makers to use better strategies, experiencing the consequences of one's actions can be used as one means of self-improvement.

In summary, decision makers may continue with unwise plans for reasons not previously considered. Because an initial investment of time or money in many situations is made under conditions of uncertainty, it is possible to improve decision making in such situations. A decision to continue with a failed plan is sometimes justified by the desire to teach the decision maker to be more careful in making future decisions. This response is attractive whether the potential learner is a young child, an older child, an adult, or even oneself. Finally, learning a lesson is perceived as especially important when the original decision has been made carelessly and relatively large resources have been invested, or when an adult teacher has the opportunity to teach a lesson to a child.

---

[5] One need not experience failure personally in order to learn the lesson; rather, one can also learn to make better decisions in sunk cost situations by observing a model continue a failed plan and then display regret at having done so (Brockner et al., 1984). However, an appropriate model is unlikely to be available much of the time.

References

Ainslie, G. (1986). Beyond microeconomics: Conflict among interests in a multiple self as a determinant of value. In J. Elster (Ed.), *The multiple self* (pp. 133–175). Cambridge, England: Cambridge University Press.

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

Baron, J. (1990). Harmful heuristics and the improvement of thinking. In D. Kuhn (Ed.), *Developmental perspectives on teaching and learning thinking skills* (pp. 28–47). Basel, Switzerland: Karger.

Baron, J. (1994). *Thinking and deciding (2nd. ed.)*. Cambridge, England: Cambridge University Press.

Baron, J., Granato, L., Spranca, M., & Teubal, E. (1993). Decision-making biases in children and early adolescents: Exploratory studies. *Merrill–Palmer Quarterly*, *39*, 22–46.

Bazerman, M. H., Beekun, R. I., & Schoorman, F. D. (1982). Performance evaluation in a dynamic context: A laboratory study of the impact of prior commitment to the ratee. *Journal of Applied Psychology*, *67*, 873–876.

Beattie, J., & Baron, J. (1995). Fitting the punishment to the crime: Nonconsequentialist decision making for punitive damages. *Journal of Experimental Psychology: Applied*, *1*, 136–151.

Bilodeau, I. M. (1969). Information feedback. In E. A. Bilodeau (Ed.), *Principles of skill acquisition* (pp. 255–296). New York: Academic Press.

Bobocel, D. R., & Meyer, J. P. (1994). Escalating commitment to a failing course of action: Separating the roles of choice and justification. *Journal of Applied Psychology*, *79*, 360–363.

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

Brockner, J., Nathanson, S., Friend, A., Harbeck, J., Samuelson, C., Houser, R., Bazerman, M. H., & Rubin, J. Z. (1984). The role of modeling processes in the "knee deep in the big muddy" phenomenon. *Organizational Behavior and Human Performance*, *33*, 77–99.

Brockner, J., & Rubin, J. Z. (1985). *Entrapment in escalating conflicts: A social psychological analysis. New*. York: Springer-Verlag.

Christensen-Szalanski, J. J. J., & Beach, L. R. (1984). The citation bias: Fad and fashion in the judgment and decision literature. *American Psychologist*, *39*, 75–78.

Clancey, W. J. (1987). *The knowledge engineer as student: Metacognitive bases for asking good questions (Technical*. Report STAN-CS-87-1183). Stanford, CA: Department of Computer Science, Stanford University.

Clark, H. H. (1973). The language-as-fixed-effect fallacy: A critique of language statistics in psychological research. *Journal of Verbal Learning and Verbal Behavior*, *12*, 335–359.

Elias, M. J., Branden-Muller, L. R., & Sayette, M. A. (1991). Teaching the foundations of social decision making and problem solving in the elementary school. In J. Baron & R. V. Brown (Eds.), *Teaching decision making to adolescents* (pp. 161–184). Hillsdale, NJ: Erlbaum.

Elster, J. (1986). *The multiple self. Cambridge,*. England: Cambridge University Press.

Frisch, D. (1993). Reasons for framing effects. *Organizational Behavior and Human Decision Processes*, *54*, 399–429.

Funder, D. C. (1987). Errors and mistakes: Evaluating the accuracy of social judgment. *Psychological Bulletin*, *101*, 75–90.

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

Goffman, E. (1959). *The presentation of self in everyday life. New*. York: Doubleday Anchor.

Hogarth, R. M. (1981). Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. *Psychological Bulletin*, *90*, 197–217.

Janis, I. L. (1982). *Groupthink. Boston:* Houghton Mifflin.

Jungermann, H. (1986). The two camps on rationality. In H. R. Arkes & K. R. Hammond (Eds.), *Judgment and decision making* (pp. 627–641). Cambridge, England: Cambridge University Press.

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, *47*, 263–291.

Klayman, J. (1985). Children's decision strategies and their adaptation to task characteristics. *Organizational Behavior and Human Decision Processes*, *25*, 179–201.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin*, *108*, 480–498.

Landes, W. M., & Posner, R. A. (1987). *The economic structure of tort law. Cambridge,*. MA: Harvard University Press.

Larrick, R. P., Morgan, J. N., & Nisbett, R. E. (1990). Teaching the use of cost–benefit reasoning in everyday life. *Psychological Science*, *6*, 362–370.

Lopes, L. L. (1991). The rhetoric of irrationality. *Theory & Psychology*, *1*, 65–82.

Schoorman, D. F. (1988). Escalation bias in performance appraisals: An unintended consequence of supervisor partici-

pation in hiring decisions. *Journal of Applied Psychology*, *73*, 58–62.

Shugan, S. M. (1980). The cost of thinking. *Journal of Consumer Research*, *7*, 99–111.

Simonson, I., & Nye, P. (1992). The effect of accountability on susceptibility to decision errors. *Organizational Behavior and Human Decision Processes*, *51*, 416–446.

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

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

Staw, B. M., & Ross, J. (1987). Behavior in escalation situations: Antecedents, prototypes, and solutions. In L. L. Cummings & B. M. Staw (Eds.), *Research in organizational behavior* (*Vol. 9*, pp. 39–78). Greenwich, CT: JAI Press.

Teger, A. I. (1980). *Too much invested to quit. New*. York: Pergamon Press.

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

Thaler, R., & Shefrin, H. M. (1981). An economic theory of self-control. *Journal of Political Economy*, *89*, 392–406.

Whyte, G. (1993). Escalating commitment in individual and group decision making: A prospect theory approach. *Organizational Behavior and Human Decision Processes*, *54*, 430–455.

Winer, B. J. (1971). *Statistical principles in experimental design. New*. York: McGraw-Hill.

APPENDIX

Scenarios in Experiments 1–3

Experiment 1 used Scenarios 1–8. Experiment 2 used Scenarios 1, 3, 6, and 7. Experiment 3 used Scenarios 5–12. The amount of the initial investment (high vs. low) is in parentheses. Alternate versions of each scenario, including adult-alone versions, were constructed based on the design of each experiment.

1. Agatha's young daughter Becky decides to take cello lessons. After Agatha buys a cello and pays for lessons ($1,200 for 3 months vs. $140 for 1 month), Becky finds she is no longer interested and wants to quit.

2. Dan's young daughter Edith selects a video to rent. After Dan pays for it and they watch the beginning ($4 for 45 min vs. $.99 for 10 min), Edith realizes she is not enjoying the movie and wants to turn it off.

3. Ira's young daughter Jill selects a school project. After Ira buys supplies and helps Jill work on it (1 month vs. 1 week), Jill discovers a better project that will take less time to complete.

4. At his urging, Harriet buys football tickets ($35 vs. $8) for her son George. Later, George's favorite player is hurt, so he doesn't want to go to the game.

5. After a large meal, Paul asks his parents to buy him a chocolate souffle ($7.95 vs. $1.50). After a few bites he finds he is too full to finish it.

6. At his son's request, Max buys ballet tickets ($80 vs. $15) for his son Luis. A week later, Luis is invited to a party at the same time as the ballet. Luis would prefer to go to the party.

7. Young Nathan urges his father Oscar to drive (4 hr vs. 30 min) to a state park for a hike. When they arrive it has turned cold and rainy. Nathan would not enjoy the planned hike and wants to go home.

8. Young Sonya joins the soccer team. After her mother Robin buys soccer equipment and Sonya attends several practices ($90 and 2 months vs. $25 and 1 week), Sonya decides she would rather play softball.

9. Colleen agrees to pay for bridge lessons ($250 vs. $50) for her elderly mother Sylvia. After several weeks of lessons, Sylvia loses interest and wants to quit.

10. An adult returning to school, Tanya selects a paper topic and gets her professor's approval. After working on the paper (3 weeks vs. 4 days), Tanya thinks of an alternative topic that is more interesting and would take less time to complete.

11. Yvonne asks Dr. Xavier to prescribe a new drug for her headaches. Yvonne buys a supply ($450 for 6 months vs. $90 for 6 weeks), but 1 week before finishing the medication decides that it isn't helping and wants to stop.

12. A private in the Army, Leroy plans a way to inventory weapons for his commanding officer, Jim. After working on it (2 months vs. 2 weeks), Leroy figures out a new method that will take less time to complete.