PREMATURE DE-ESCALATION IN RESPONSE TO FAILED INVESTMENT: A TEST OF ESCALATING COMMITMENT, MARGINAL DECISION-MAKING, AND MENTAL BUDGETING

CHIP HEATH
Graduate School of Business
University of Chicago
1101 East 58th Street
Chicago, IL 60637
(312) 702-9677

MITCHELL A. PETERSEN
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208-2006
(708) 467-1281

PREMATURE DE-ESCALATION IN RESPONSE TO FAILED INVESTMENT: A TEST OF ESCALATING COMMITMENT, MARGINAL DECISION-MAKING, AND MENTAL BUDGETING

Abstract

Using sophisticated subjects in an environment that should make the norms of economic behavior highly salient (MBAs in a corporate finance class), we test three theories about how people respond to previous investments: escalating commitment, mental budgeting, and marginal decision making. The results support mental budgeting.

PREMATURE DE-ESCALATION IN RESPONSE TO FAILED INVESTMENT: A TEST OF ESCALATING COMMITMENT, MARGINAL DECISION-MAKING, AND MENTAL BUDGETING

One of the key tasks of managers in organizations is to allocate resources effectively. Thus it is not surprising that management scholars have long been interested in the idea that people are prone to <u>mis</u>allocate resources when they encounter failure or negative feedback. A variety of research has started from the premise that when people's investments fail to produce a payoff, they respond to failure by becoming increasingly willing to invest <u>additional</u> resources (Arkes & Blumer, 1985; Brockner & Rubin, 1985; Garland, 1990; Staw, 1976; Staw & Ross, 1989; Thaler, 1980; Whyte, 1986).

Researchers have also proposed a variety of psychological mechanisms to explain why people "escalate commitment." For example, researchers have proposed that people escalate commitment: because they want to justify a previous investment that appears to have been in error (Brockner, 1992; Brockner & Rubin, 1985; Staw, 1976; Staw & Ross, 1989); because they do not want to waste their previous investments (Arkes & Blumer, 1985), because they become risk-seeking in the domain of losses (Garland, 1990; Thaler, 1980; Whyte, 1986), and because they see a given incremental investment as psychophysically "smaller" when it occurs in the context of larger absolute investments (Garland, 1990; Garland & Newport, 1991). In addition, researchers have attempted to determine which of a large number of procedures limit escalating commitment most effectively (Brockner & Rubin, 1985; Simonson & Staw, 1992).

Recent studies show that escalation is not a universal reaction to failed investments. In a series of seven studies Armstrong, Coviello, and Safranek (1993) fail to replicate the traditional escalation effect. McCain (1986) and Garland, Sandefur, and Rogers (1990) showed that people will de-escalate commitment over time in response to negative feedback.

Heath (1995) provides evidence that not only will people de-escalate commitment, but that they will do so improperly, in situations where they should continue or increase investment. The examples of de-escalation provide an interesting counterpoint to the examples of escalation that have shaped theory on how people allocate resources and respond to failure, and they challenge theories, like those above, which have been developed to explain escalation.

In this paper, we extend previous work on escalation in two ways. First, previous work on escalation has typically defined escalation in terms of how people respond to repeated loss, negative feedback or failure. For example, Brockner's (1992) review of escalation research defines escalation situations in terms of "repeated (rather than one-shot) decision making in the face of negative feedback about prior resource allocations, uncertainty surrounding the likelihood of goal attainment, and choice about whether to continue." However, researchers have never defined precisely how people should respond to failure. Is it always bad to work harder in response to negative feedback? When should we learn from our failures and quit and when should we ignore failure and persist? In this paper, we use principles from economics and finance to describe the correct procedure for making repeated decisions and to highlight situations where people should respond to negative feedback and where they should not.

Second, in a laboratory setting we explore the behavior of sophisticated subjects who have been specifically trained to make correct economic decisions and we give subjects a great deal of experience with the investment environment. Like most laboratory research, previous research on escalation has typically used undergraduate subjects. Cynics might argue that such lab research does not indicate how people would behave in real contexts because such subjects have insufficient training to understand the decision context. In these experiments, we ask MBAs in corporate finance class to participate in a class exercise on investment decisions. Thus, our subjects have received more training than the average corporate manager, and they make their decisions in a context which should highlight economic rules for correct decisions.

Using the principle of competitive theory tests (Platt, 1964), our experimental design allows us to test the predictions of three different approaches to describing people's investment behavior: marginal decision making from economics, escalating commitment, and mental budgeting. The decisions of our subjects are not consistent with the prescriptions of economic and financial models or with the prediction of escalating commitment. However, decisions can be described by mental budgeting: people frequently stop investing too early.

LEARNING FROM NEGATIVE FEEDBACK: THE ECONOMIC APPROACH TO REPEATED DECISIONS

The literature on escalation has treated persistent investment in the face of negative feedback as a prima facie case for errors in judgment. Indeed, it does seem unreasonable to invest repeatedly when investments are not successful. However, there are situations where repeated investment is the only way to achieve success. For example, consider pharmeceutical research and development where billions of dollars may be spent to develop a new drug. When we observe such extravagant investments, how can we gauge whether people are investing wisely or unwisely? In this section, we use models from finance and economics to describe how such decisions should be made. By better understanding how decisions should be made, we will sharpen our ability to describe actual behavior.

Northcraft and Wolf's (1984) critique of escalation research argued that researchers have typically assumed that escalation was always unreasonable when in fact a marginal decision often favors escalation of invesment. Northcraft and Wolf described how people should use information about the timing and magnitude of cash flows to calculate whether to continue to invest. For example, when people receive a payoff only when a project is complete, it is reasonable to escalate investment as the project approaches completion because completing the project will produce a high payoff and not completing the project will produce nothing. The cash flows considered by Northcraft and Wolf are one key component of an investment decision. The other key component is the probability of receiving a payoff. Below, we describe how people should use negative feedback or failure to calculate the probability that an investment will succeed.

Economics tells us that to make a good decision we should weigh marginal benefits (i.e., current and future benefits) against marginal costs and choose the action where benefits most outweigh costs. Past costs (and benefits) are irrelevant to the current decision--they are "sunk." However, in some situations our past investments provide us with information about what to expect in the future. Some failures are informative and should be considered, others should not. In either case, however, our decisions should respond to the information contained in the feedback from past investments, not to the past costs and benefits themselves.

When failure is not informative. We start with a simplified investment project where negative feedback is not informative because there is nothing to learn from the feedback. Investments in this project cost C dollars and "succeed" with a known probability, p. If investments succeed they pay S dollars, and if they fail, they pay nothing. To make this project more concrete, the reader can imagine a new drug development process, where companies invest money in research and succeed if they discover a new drug, or an oil exploration process, where companies invest money to drill wells and succeed if they discover oil.

The economic recommendation in this situation is simple: The company should invest if and only if the the expected value of success exceeds the investment cost (i.e., $pS>C$). In the language of finance, investments like this have a positive net present value (NPV). Since this calculation remains the same across sequential investments, negative feedback in this situation should <u>not</u> affect future decisions. If payoffs are sufficiently large and the probability of success is sufficiently small, the company should continue to invest even if they accumulate failure after failure. Imagine, for example, the opportunity to invest invest one dollar for a 1/1000 chance of earning 100,000 dollars. Although this is a very attractive investment, an investor might accumulate a large number of "failures" before succeeding. In fact, the probability of accumulating 100 failures without a success is over 90 percent.

<u>When failure is informative</u>. In the situation above, negative feedback is irrelevant because the investment domain is known completely. Therefore, there is nothing to learn about the future based on our past experience. This is not a realistic model of most decision environments. In most situations, there is some uncertainty about the parameters of the decision, so feedback from our past investments may teach us whether future investments are advisable--whether the costs are appropriate, whether the benefits are adequate, and whether the chances of success are sufficiently high.

In a particularly important class of situations, the probability of success is not precisely known and success or failure informs us about the true probability of future success. Imagine, for example, a region where oil fields either produce good wells 100% of the time or 0% of the time. Wells are expensive, so an oil company would prefer to drill only on good fields. They sink a well at a new field and it comes up empty. Should the company drill another? The answer is obviously no. The single failure is tremendously diagnostic and the company learns a great deal from it because oil in this region is either <u>always</u> present or <u>always</u> absent. The single failure proves that the true probability of oil on that field is 0%.

Because the alternatives in the situation above are so distinct, the correct response is obvious. In more realistic circumstances, failure may be partially but not completely diagnostic. Suppose, for example, that there are are richer and poorer fields; rich fields produce 45% of the time and poor fields produce 5% of the time. Again, the company sinks a well at a new field and it comes up dry. The dry well provides some evidence that the company might be drilling on a poor field, but the evidence is not conclusive.

A formal procedure known as Bayesian updating allows us to calculate exactly how much we should learn from a given failure (see appendix for details). In general, when we do not know the probability of success, our estimate of p should decrease after each failure. When p drops below C/S (i.e., when the NPV becomes negative), we should no longer invest. Suppose, for example, that our oil company finds it worthwhile to drill only if there is at least a 10% chance of success. For the field that is either an 5% or 45% producer, the estimated probability of success for the first well is 25%. This probability drops to 19.7% for the second well and with each subsequent failure it drops further. After four dry wells (but not before), the company should quit drilling. Figure 1 displays the estimated probability of future success after each of the first eight failures. The appendix contains an extended explanation of this example. For more detail, see Anderson, Sweeney, and Williams (1993, p. 829-838) or most introductory finance books. Even without the details of the mathematics, we can still make qualitative predictions. For example, the optimal number of investments increases with payoff, decreases with investment cost, and increases with our estimate of the success probability.

Insert Figure 1 about here.

Again note that this argument does not imply that people should respond to sunk costs. Decisions should not respond to the sunk costs associated with previous failures but they should respond to the information contained in the failures themselves. For example, a large investment that fails will have the same effect on the updated probability as a small investment that fails. More generally, there are many costs that do not inform us about the true probability of success. For example, imagine that the oil company had to pay a legal fee to clear the title on the land before they could begin drilling. The legal fee tells the company nothing about the quality of the field and should not affect the decision whether or not to drill an additional well.

Thus in economic recommendations for decision making, "failure" or negative feedback is not always informative. When parameters are unknown and failure does inform us about future costs and benefits, there are procedures to calculate exactly how much information we should learn from a given piece of negative feedback.

THEORETICAL PREDICTIONS ABOUT REPEATED DECISIONS

In the experiments below, we present people with a number of investments that share the characteristics described above. After we describe the potential investments, we develop predictions from three different theories of resource use: marginal decision making, escalating commitment, and mental budgeting.

For each experiment, we constructed a number of investment cases where success is probabilistic and where the probability of success is unknown. In some cases, for example, subjects are told that the true probability of success is either 5%, 15% or 25% with equal probability. Because subjects are initially uncertain about the true probability, as they invest, they learn more about the true probability of success. If their investments fail repeatedly, subjects should become more and more confident that the true probability of success is not 25%, but 15% or perhaps 5%. We design our experiments so that subjects should not continue to invest after the probability drops too close to the lower end of the distribution.

Researchers have pointed out that escalation situations are characterized by "uncertainty about the future chances of success" (e.g., Brockner, 1992; Staw, 1996), and that laboratory situations that precisely describe probabilities and payoffs may not allow people enough psychological leeway to demonstrate escalation (Staw, 1996). These researchers would presumably question the generalizeability of previous results which have provided people with precise probabilities and payoffs and have failed to find escalation (e.g., Heath, 1995). Because we wanted to compare subjects' choices to normative responses, we had to give people enough information to potentially calculate the correct decision. However, by making the true probability of success uncertain, we allow people room to rationalize their decisions and escalate. In fact, our design involves much more ambiguity about success than other procedures that have been used to study escalation (e.g., the counter game that Brockner and Rubin, 1985, use to study entrapment).

In order to gain more power to distinguish the three theories of investment behavior, we added a sunk cost which provides no information about the true probability of success. On some projects, before they could invest, subjects first had to pay a fee to purchase the "right to invest" in the project. Although this purchase fee provides subjects with no information about the true probability of success, it increases the losses that they must endure if their initial investments fail.

Our analysis will allow us to look at reactions to the two kinds of investments. The three theories predict different signs on coefficients for these terms--the coefficient on the uninformative sunk cost and the coefficient on the informative investments.

Escalation of commitment

Theories of escalation or entrapment assume that when people encounter negative feedback; they seek to justify their initial decisions by investing additional resources (Brockner, 1992; Brockner & Rubin, 1985; Staw, 1996; Staw & Ross, 1989). Brockner (1992), for example, describes this facet of escalation as resulting from "decision makers' unwillingness to admit that they were mistaken in having become committed to the intially chosen course of action." Thus rather than learning from their failures and quitting, theories of escalation predict that people will invest additional resources in the hope that they can turn the situation around. Furthermore, theories of escalation do not distinguish between informative and non-informative losses. Under the economic model, a company's decision about how long to drill for oil should not be affected by the legal fees it pays to clear the title to the land. However, according to theories of escalating commitment, the fees add to the "losses that have resulted from an original course of action" (Staw & Ross, 1989), and should prompt further escalation.

Of course in the long run, even strong escalators might eventually realize the futility of their investments. However, in the short-run, at least, people should be willing to invest more resources when they have accumulated greater losses in the past. This should be particularly true in our experimental design because people are highly uncertain about the true probability of success and thus have leeway to rationalize their escalation.

> Hypothesis 1a: Theories of escalating commitment predict that people will escalate commitment in response to informative losses.
>
> Hypothesis 2a: Theories of escalating commitment predict that people will escalate commitment in response to non-informative losses.

Economic marginal decision making

Economic theories assume that people behave rationally in their resource allocations. Thus they predict that real decision makers will follow the normative rules outlined in the previous section. For example, when negative feedback is informative, economic theory predicts that people will respond to failure by decreasing investment. Furthermore, economic theory predicts that people will distinguish between non-informative failures (like purchase fees) and informative ones.

In fact, some studies have shown evidence of de-escalation that might be interpreted as evidence in favor of the predictions of the economic model. McCain (1986) argued that escalation occurs only in the very early stages of a project when people are still trying to interpret the feedback provided by early failures. He points out that in Staw's paradigm even though people tend to escalate immediately after bad news, this tendency reverses over time (e.g., Staw & Fox, 1977). McCain replicates this scenario over multiple rounds and shows that people de-escalate investment later in an investment sequence. Another study by Garland, Sandefur, and Rogers (1990), showed that professional oil field geologists were less likely to invest money to drill on a particular field as sunk costs (dry wells) increased. As the number of dry wells increased from one to four, their estimated probability that the next well would succeed dropped from 51 percent to 19 percent. The decrease in estimated probability provides reasonably strong evidence that people are doing something like Bayesian updating--using information about past failures to estimate the probability of future success.

Although behavior in the McCain (1984) and Garland et al (1990) studies is consistent with the predictions of economic models, these demonstrations, like previous research on escalation, did not give subjects enough information about the costs and benefits of additional investments to allow them to calculate whether they should continue or stop investing. The current study provides an opportunity to extend their results by giving subjects enough information to calculate exactly when they should continue and stop. Because previous studies omitted this information, we cannot determine conclusively whether people de-escalate because they learn rationally from their experience that investment is a bad thing, or for some other reason.

> Hypothesis 1a: Theories of marginal decision making predict that people will respond to informative losses by decreasing investment.
>
> Hypothesis 2a: Theories of marginal decision making predict that people will not alter their behavior in response to non-informative losses.

Mental budgeting

Recently, Heath (1995) proposed a model of mental budgeting that assumes that people set budgets for their investments and stop investing when they exceed their budget. Mental budgeting predicts that people will sometimes de-escalate investment even when they are not learning anything from their failures. For example, in one experiment using real money, undergraduates prematurely decreased investment even in a situation where negative feedback was completely non-diagnostic. Heath gave subjects the option to invest in a project with known payoff, costs, and probability of success (e.g., the chance to invest $0.50 for a 25% chance to win $4.00). The probability of success was defined by a chance device controlled by the subject, and subjects played their choices for real money. In this situation, because everything about the investment is constant--costs, payoffs, and probabilities do not change--there is nothing to be learned from experience. The expected value of investments were always positive, and subjects should have continued to invest until they achieved the payoffs. (This assumes that subjects do not run out of money to invest, a possibility eliminated by the experimental design). However, 80 percent of the subjects prematurely decreased investment on more than one project.

Because the organizational behavior literature has emphasized escalation, the results of this experiment may seem surprising. At another level, however, they should be very intuitive: de-escalation results from a budgeting process which is pervasive in our investment and personal decisions (Heath, 1995; Heath & Soll, in press). In our personal decisions, for example, we allocate money toward this week's entertainment or this month's household expenses, and we resist exceeding our pre-determined allocations. This kind of budgeting behavior has been documented in descriptions of consumers for over five decades (Zelizer, 1993).

Mental budgeting argues that people allocate resources by (1) setting budgets for the investments they undertake and (2) tracking ongoing investments against their budget. People decrease investment when total investments exceed their budget. Counter to the recommendations of economic models (and similar to previous behavioral models) the mental budgeting model argues that sunk costs affect people's future investments. However, it predicts that people will frequently respond to sunk costs by decreasing future investment because the sunk costs push people toward their budget limit.

The model assumes that to set a budget people assess the expected benefits of the opportunity, and mentally earmark the money, time, or other resources that they will invest to acquire those benefits. Earlier research on escalation indicated that budget setting is frequent, and that it effectively limits escalation. For example, in one of their studies of entrapment, Brockner, Shaw, & Rubin (1979) found that 80% of control subjects spontaneously set a "commitment limit" for their investment.

How then, do people decide what commitment limits or budgets to set? Although different people may desire a different rate of return on their investments (some people may be willing to invest \$90 to receive \$100, while others may be willing to invest only \$60), the total benefits expected from a project provide a very clear upper bound on mental budgets--people are unlikely to budget \$110 to gain \$100. Thus, in this paper, we assume that people budget less than the total benefit they may receive from the project.

Because the MBAs we use in the current experiments are much more sophisticated than the undergraduates Heath used in his previous study, they may use at least some economic information (e.g., payoffs and probabilities) to set their budgets. However, to predict the overall pattern of investments, we do not need to specify exactly how subjects will combine information to set their budget. We need only point out that all investments will push people toward their budget limit and make them less likely to continue investing. Thus, in situations where people accumulate failed investments, mental budgeting, like economic theory, predicts that people will respond to failure by decreasing investment. However, unlike economic theories, mental budgeting predicts that people will not distinguish between non-informative failures (like purchase fees) and informative ones. Both kinds of investments reduce the amount of money available in the budget, and therefore both kinds of investments should decrease future investment.

> Hypothesis 1a: Mental budgeting predicts that people will respond to informative losses by decreasing investment.
>
> Hypothesis 2a: Mental budgeting predicts that people will respond to non-informative losses by decreasing investment.

STUDY 1

In this study, we ask subjects, MBAs in a corporate finance class at the University of Chicago, to make decisions about dynamic investment projects. In contrast to the undergraduates typically used in previous studies of escalation (Arkes & Blumer, 1985; Brockner & Rubin, 1985; Garland, 1990; Heath, 1995; Staw, 1977), the subjects in our studies are quite sophisticated. By the time subjects perform this exercise they had taken at least one micro-economics class (which covered correct responses to sunk costs), one course in probability and statistics, a class in capital markets (investments), and several weeks of corporate finance (which covered techniques of capital budgeting including net present value and the Bayesian techniques necessary to calculate the answer to the problems above). Some researchers have shown that even short courses on cost-benefit analysis significantly improve people's personal decisions, and that these improvements generalize to unrelated contexts (Larrick, Morgan, Nisbett, 1990). In selecting our subjects, we follow this strategy to its logical conclusion by exploring the decisions of people who have received about as much training on the normative model as any manager ever encounters. Furthermore, because they perform the exercise as part of their corporate finance class, the norms of economic behavior should be highly salient at the moment they are making their decisions.

In addition, this population is highly motivated to perform well. The class is an elective that is typically taken by students who plan to pursue a career in corporate finance. Furthermore, the exercise was presented as a course homework assignment. Subjects were told that their grade would depend on how well their decisions upheld the general principles of the economic analysis. (Subjects lost points for purchasing negative NPV projects or declining positive NPV projects. They also lost points for systematically quitting too early or too late.)[1]

Because subjects are sophisticated and economic norms are prominent, we may find it difficult to document any deviations from the investment patterns recommended by the economic model. Thus our design provides an especially stringent test of the psychological theories of escalating commitment and mental budgeting. However, studying these subjects in this environment may help us better understand what errors to expect in organizational environments where decision makers are explicitly trained for their jobs and where people are explicitly trying to decide based on economic criteria.

Procedure

Subjects were 194 MBAs enrolled in a corporate finance class at the University of Chicago. Subjects participated in a computer investment exercise as part of their class requirements. During the computer exercise, each subject made decisions about 18 cases similar to the investment described above. For each case, subjects decided whether to invest a cost (C) in exchange for a probability (p) of receiving a payoff (S). Some cases also involved a purchase fee, which subjects had to pay to buy the "rights to invest."

After the computer described each case, subjects decided whether to invest for the first time. If the subject chose to invest, the computer program generated a random number, compared it with subject's probability of success, and if the investment succeeded, awarded the payoff to the subject. If the investment failed, the subject could invest again or quit. Once the subject succeeded or quit, the round ended and the computer presented the subject with the next case. Each subject saw the cases in a different random order.

Materials

We constructed investment cases to have a structure like the examples discussed above. Each investment was described by an investment cost (C), a payoff (S), a purchase fee, and a probability distribution over the true probability of success (p). Investment costs were constant across cases at 10 dollars. We varied payoffs (either 67, 80, or 100 dollars), purchase fee (0, 3, or 6 dollars) and probability distributions. See Table 1 for the complete design.

Insert Table 1 about here.

Because we wanted to give people an opportunity to learn from their investments, we did not tell them the exact probability of success. Instead, for each case, we described one of two probability distributions, and told them that their probability of success was drawn from that distribution. In the first distribution, the true probability of success had an equal probability of being either 5, 15, or 25 percent. In the other distribution, the true probability was either 9, 20, or 31 percent. When a case appeared, the computer randomly chose one of the three probabilities as the true probability.

Because the probability of success was unknown, failures informed the subject about the true probability. We constructed each case so that it was optimal to quit investing after a finite number of failed investments. Using Bayes' rule, subjects could derive the expected probability of success after N failures, and compute when the expected payoff from a 10 dollar investment dropped below 10 dollars. This happened after one failure when the payoff was 67 dollars (when the probability drops below 10/67=.15), but after eight failures when the payoff was 100 dollars (when the probability drops below 10/100=.10). Thus, it was optimal to stop investing after one investment when the payoff was 67 dollars and after eight investments when the payoff was 100 dollars. Table 1 contains the optimal stopping points for each project, along with the project's net present value (NPV) if subjects stopped at this point. For comparison, Table 1 also reports the upper bounds on number of investments predicted by mental budgeting. To calculate these upper bounds, we subtract the purchase price from the payoff and divide by the 10 dollar cost of investment. This figure represents the number of investments that can be made before the total costs of investments exceed the total benefits. Mental budgeting predicts that people will stop at or before this point, depending on where they set their individual budget.

Combining the three payoffs with three purchase fees and two probability distributions produced 18 cases. Of the eighteen cases, subjects should have purchased 15. The other three had negative NPVs.

Results

Table 1 records the number of people who purchase a project[2] and the average number of investments they made. Table 2 presents a logistic regression to understand how long subjects invested. Each observation represents an individual's decision about whether or not to invest another time on a particular case. The investment decision on each trial was coded as a "1" if the subject invested, and a "0" if the subject quit. Thus an individual who invested three times before quitting would add four observations to our analysis: 1, 1, 1, 0. When an investment succeeded and the subject received the payoff they left the analysis.

Insert Table 2 about here.

The results in Table 2 indicate that subjects responded in a reasonable way to the experimental manipulations. For example, subjects were more likely to continue as the payoff for success increased from 80 to 100 (b=.415, z=4.60, p<.001) and less likely to invest when the average probability of success was lower (b=-.214, z=-2.77, p<.01).

This evidence indicates that subjects responded to some of the information they needed to calculate expected payoffs. This provides some support for the economic model, but the responses to payoff information is inconsistent with learning stories that assume that decision makers focus purely on historical information (e.g., stories that depend on reinforcement history, Goltz, 1992, 1993). It could be consistent with mental budgeting if people use this information to set their budgets or with theories of escalation that leave some room for economic variables to affect decisions (e.g., Brockner, 1992; Staw, 1996).

Consistent with the predictions of marginal decision making and mental budgeting, as failures accumulated, subjects were less likely to invest ($b=-.126$, $z=-5.06$, $p<.001$). However, consistent with mental budgeting but not marginal decision making, increasing the purchase fee from zero to 3 dollars made people quit earlier ($b=-.203$, $z=-2.25$, $p<.05$). Increasing the purchase fee from 3 to 6 dollars had a marginally significant effect in the same direction ($b=-.162$, $z=-1.63$, $p=.10$). Combined, these results provide some support for mental budgeting over marginal decision making or escalation.

At this point, the main analyses indicate little evidence of escalation. However, the quadratic term in the regression indicates that subjects decreased investment at a decreasing rate ($b=.008$, $z=5.66$, $p<.001$). This could be interpreted as evidence that some subjects escalated late in the process. However, it could also provide evidence that the composition of the sample changed over time. For example, if some subjects always invested 15 trials, and some subjects always invested 4 trials, then the only subjects who remained in the sample after 10 trials would be those that initially planned to invest 15 trials. This pattern would produce a positive quadratic term even if no one escalated investment.

One way to distinguish escalation from "long investment" is to isolate reactions to purchase fees. Escalators should have increased investment in response to an increase in the purchase fee. On the other hand, people that habitually invested a long time should invest a long time whether the purchase fee was high or low. To test whether there were subpopulations of escalators, we tested whether we could better predict a subject's response by knowing how that subject responded to the change in purchase fee on an earlier case. Our dependent variable was the difference between the number of investments a subject made when the purchase fee was 0 dollars and when it was 6 dollars; the independent variable was the same difference score calculated for another case. The regression yielded a non significant coefficient on the subject's previous response ($b=-.048$, $t=-1.35$, $p=.18$; R-squared of .00), providing no evidence of escalation.

The results of this experiment generally support mental budgeting. People become less likely to invest as investments continue to fail to produce the payoff, consistent with both budgeting and marginal decision making, but not with escalation. However, people also become less likely to invest in response to the initial sunk cost, consistent with budgeting but not marginal decision making.

However, there may be stories other than budgeting that might predict similar patterns of behavior. For example, disappointment or punishment stories would predict that people would reduce investment over time even if they did not set a mental budget.[3] To provide additional evidence for budgeting, we attempt to predict the 18 average stopping times recorded in Table 1 using the upper bounds set by mental budgeting and the optimal stop times specified by the economic model. The resulting regression predicts a reasonable amount of the variance (Adjusted $R2 = .45$), and the standardized coefficient on mental budgeting is quite large ($B = .989$, $p < .002$), while the standardized coefficient on the economic prediction is smaller and in the wrong direction ($B = -.517$, $p = .06$). This result indicates that mental budgeting predicts features of the overall data that could not be predicted with a simpler disappointment or reinforcement story. Furthermore, it provides strong evidence that decisions deviate from correct economic decisions in the direction predicted by mental budgeting.

STUDY 2

In this study we wanted to examine whether people converge toward correct economic behavior after receiving repeated experience in making decisions. A common feature of economic arguments is that people benefit greatly from experience with a task (Davis & Holt, 1993). After they perform the task a few times, people may converge on the correct decision, either because they are better able to think through the task or because low-level reinforcement and punishment pushes them towards the correct decision. To test the impact of experience, in this experiment we required subjects to perform the investment simulation five times. By contrasting their decisions on the first two rounds with their decisions on the last two rounds, we will be able to determine whether experience with the task significantly alters investment decisions. In addition, we simplified the decision by reducing the potential probability distribution to two values. This should make it even easier for subjects to apply their knowledge of the normative model.

Procedure

The procedure for this study was the same as in Study 1. Students enrolled in a corporate finance class at the University of Chicago (N=110) participated in a computer investment exercise as part of a class homework assignment.

As in the previous study, subjects sat at a computer screen and made decisions about investment cases (15 cases in this study). After they completed the exercise once, subjects repeated the exercise four additional times with the same 15 cases. As in the first round, the true probability for a case and the results of each investment were randomly determined. Also, the order of the 15 cases was randomly determine each time they repeated the exercise.

Materials

Investment cost in this experiment was again constant at 10 dollars, and we developed cases by manipulating payoffs (70, 90, or 150 dollars), purchase fee (0, 10, or 20 dollars), and probability distribution. However, we simplified the probability distribution to two values (5% or 45% for some problems, 10% or 40% for others). Table 3 reports the parameters, optimal stopping points, and NPV for each case.

Insert Table 3 about here.

Results

As in the previous experiment, we used a logistic regression to estimate the probability that a subject will continue to invest.[4]

Insert Table 4 about here.

The first column in Table 4 contains the results of this basic regression. As in Study 1 there was a linear and quadratic trend in continuation based on the number of previous investments. The probability of continuing decreased over investments (b= -.069, z=-4.37, p<.000) but at a decreasing rate (b=.007, z=7.82, p<.000). Again, the linear trend is consistent with mental budgeting and marginal decision making but not escalating commitment.

To understand the positive coefficient on the quadratic term, we conducted the same analysis of heterogeneity that we conducted in Study 1. This regression yielded a non-significant coefficient on the subject's previous response (b=-.005, t=-.301, n.s.; R-squared for the regression is .00). Thus, all subjects seem to have responded to an increase in purchase fee by quitting sooner, and there is no evidence that subpopulations of subjects differed in their response to sunk costs. Because there was no evidence of heterogeneity across subjects in responses to sunk costs, the positive sign on the quadratic term is best explained by heterogeneity across subjects in base level of investment (e.g., some subjects habitually invested 4 times and others invested 15).

To further distinguish among theories, we can look at the sign on the sunk cost terms. Increasing the purchase fee from 0 to 10 dollars significantly decreased investment (b=-.180, z=-3.46, p<.000) but increasing the purchase fee from 10 to 20 dollars did not (b=-.072, z=-1.12, p=.263). This pattern generally supports mental budgeting (since marginal decision making cannot explain the significant negative effect on the 0 to 10 dollar increase). However, increasing sunk costs does not seem to have a negative effect after the first 10 dollars.

As in Experiment 1 we can attempt to predict the 15 average stopping times in Table 4 using the upper bounds set by mental budgeting and the optimal stop times specified by the economic model. The resulting regression predicts quite well (Adjusted R2=.85). Again, the standardized coefficient on mental budgeting is quite large (B = .978, p < .0001), while the standardized coefficient on the economic prediction is smaller and in the wrong direction (B = -.280, p < .05).

We can now investigate whether this pattern of results is altered by experience and knowledge (see Table 4). To explore how experience affects investment decisions, we contrast investment decisions on the first two trials with responses on the last two trials (recall that each subject performed the exercise five times). At both times, people respond to the purchase fee by decreasing investment, and if anything, they are more likely to do so on late versus early trials. Thus the general pattern of responses that supports mental budgeting does not seem to be altered by experience.

To explore how knowledge affects decisions, we contrast investment decisions by those who made the smallest number of correct purchase decisions and those who made the largest number (i.e., who purchased positive NPV projects and declined negative NPV projects). The median number of correct purchases was 73.4%. Subjects who were above the median on this variable presumably understood the exercise better and were working harder at it. As can be seen in Table 4, when we focus on this group of highly competent subjects, the pattern of responses to most variables is more correct (e.g., subjects invest more as the payoff moves from 90 to 150 dollars and less when the probability distribution has the lower mean). However, the key evidence for mental budgeting remains significant--high knowledge subjects still quit earlier in response to the purchase fee.

To summarize, this study replicates and extends Study 1 using a different (simplified) set of investment decisions. As in Study 1, people stop investing earlier as investments mount or when they pay an up front purchase fee. This pattern is consistent with mental budgeting, but inconsistent with theories of escalation or marginal decision making. The subsequent analyses indicate that neither experience nor knowledge alters this basic pattern. Responses are similar when we look at subjects early or late in their experience. Responses are also similar for high and low knowledge subjects. Again, the upper bounds predicted by mental budgeting do a much better job predicting average stopping points than the optimal stops specified by economic theory.

GENERAL DISCUSSION

The two studies in this paper simultaneously tested three theories of decision making in response to negative feedback. The basic pattern of responses supports mental budgeting over marginal decision making and escalating commitment. We document this pattern using sophisticated subjects and in an environment that makes the norms of economic action highly salient. Study 2 indicates that the pattern of results is not dramatically altered by experience or knowledge. Furthermore, the upper bounds suggested by mental budgeting predict average stopping times in both experiments.

Although there are clearly many different factors that may prompt someone to escalate or de-escalate commitment (Staw & Ross, 1989), researchers have historically focused on individual, psychological factors. The current results suggest an important psychological limit on escalation. Escalation pits two kinds of losses against each other: In order to avoid "losing" an initial investment, people may desire to invest additional resources; however, when they invest additional resources, they must risk losing those resources as well. Even if people want to escalate commitment to redeem a monetary sunk cost, mental budgeting argues that they are unlikely to do so by increasing monetary investment. Although losing the initial investment is painful, so is the threat of losing future investments by exceeding a budget and sacrificing too much to achieve a goal.

When possible, people may adopt secondary strategies to redeem the initial sunk cost while avoiding the additional losses that would result from exceeding a budget. For example, when people are seeking to redeem a monetary sunk cost, people often invest <u>non-monetary</u> resources to make the initial monetary investment pay off. Research suggests that people may exert greater efforts to use a costly item (Arkes & Blumer, 1985). Staw & Hoang (1995) argued that NBA teams play high draft picks more than is warranted by their performance, presumably to justify the high salaries commanded by a high draft pick. Here, teams improperly escalate commitment in response to a monetary sunk cost, but they do not make additional monetary investments to do so. Heath (1995) showed that people would invest time (but not money) to salvage a monetary sunk cost, and money (but not time) to salvage a sunk cost of time. According to this logic, NBA teams may over-play high-priced draft picks but they are unlikely to give them higher raises. Thus, mental budgeting suggests that it is unlikely that decision makers will "throw good money after bad." On the other hand, they may throw good time after bad money.

Following the lead of Northcraft and Wolf (1984), we began the paper by exploring a precise normative model for how people should respond to negative feedback. Here, we again emphasize the importance of this kind of approach for researchers who study resource allocation. Research on escalation has persisted because researchers wanted to understand what seemed to be an irrational tendency to pursue losing courses of action. By having a clearer normative model, we are in a better position to assess when people should and should not pursue losing courses of action and to clarify what errors people actually make.

The importance of using clear normative models is even more apparent if we consider accumulating evidence that escalation is not the only reaction to sunk costs. Even when the literature agreed that escalation was the phenomenon to be explained, there were a number of disagreements about how to interpret it. The theoretical story becomes even more complex when we consider recent evidence that people may also de-escalate in response to sunk costs (Garland, Sandefur, and Rogers, 1990; McCain, 1986), even when they should not (Heath, 1995 and the current studies). These results indicate that we still have more to understand about how people allocate resources and respond to losses. However, they complicate the task of researchers by pointing out that people may make mistakes in either direction.

In his recent review of the literature, Staw (1996) cautioned against privileging economic models in our study of behavior: "by requiring a repeated test of behavioral versus economic criteria, one is implicitly assuming that economic explanations are closer to the truth than behavioral approaches." We argue that there has actually been far too little conversation among researchers who use economic and behavioral approaches. First, at the level of theory tests, competitive theory tests that contrast the predictions of economic theory with the predictions of behavioral approaches can only sharpen our theoretical tests (Platt, 1964).

Second, if we take the economic model not as a description of behavior, but as a normative model that clarifies how decisions should be made, this too puts us in a better position to understand behavior. By proposing a clearer model of how a decision should be made, normative theories direct our attention as researchers. They point out where errors may occur and they help us understand the direction and magnitude of problems. When we neglect the appropriate normative model, we may miss interesting behavioral phenomenon. For example, there are examples in the literature on "escalation" where people systematically quit investing <u>prematurely</u>, but researchers missed this phenomenon because they were not using the appropriate normative model to evaluate people's decisions (see Heath, 1995, for more on this point). In the research in this paper, distinguishing between informative and uninformative investments helped us distinguish between mental budgeting and marginal decision making.

Finally, using normative models helps us think about how organizations might effectively improve decisions. Each potential cause of escalation recommends a different way to solve problems of resource allocation in organizations. For example, rational choice researchers have borrowed social psychological insights that decision makers are often punished for inconsistency and have modeled escalation as an agency problem (Kanodia, Bushman, & Dickhaut, 1989; Prendergast & Stole, 1996). If decision makers are penalized for being inconsistent or for experiencing bad results during their tenure on a job, then they may escalate commitment because they want to keep their job or their reputation. Organizations can solve agency problems by altering incentives. However, to the extent that problems of escalation or de-escalation reside in how individuals approach decisions, such incentive structures may be at best misplaced, and at worst, counterproductive.

As an example of how organizational incentives might be counterproductive, consider the elaborate budgeting procedures that organizations use to predict and control resource outlays. Previous researchers in escalation have observed that escalation may be effectively prevented by "limit setting" procedures like those used in organizational budgeting procedures (e.g., Brockner, Shaw and Rubin, 1979; Simonson and Staw, 1992). However, by giving people additional incentives to adhere to a budget, organizational budgets may solve one problem (escalation) at the expense of another (de-escalation). Organizational participants often complain that opportunities were pursued ineffectively or missed completely because money "wasn't in the budget." Although it is clear that budgets do not always constrain behavior (Ross & Staw, 1993) they may do so often enough that organizations fail to take advantage of many valuable opportunities. By using a normative model to locate the key errors decision makers are making, we can better evaluate what organizational interventions might promote better decisions.

In sum, using a normative model of decision making as a lens to view the world, rather than conceding the superiority of rational models of behavior, actually may sharpen our insights about both rational and irrational behavior.

REFERENCES

Anderson, D., Sweeney, D., and Williams, T. 1993. Statistics for Business and Economics (Edition 5). New York: West Publishing Company.

Arkes, H. R., & Blumer, C. 1985. The psychology of sunk costs. Organizational Behavior and Human Decision Processes, 35: 124-140.

Armstrong, J. S., Coviello, N., & Safranek, B. 1993. Escalation bias: Does it extend to marketing? Journal of the Academy of Marketing Science, 21: 247-253.

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. & Rubin, J. Z. 1985. Entrapment in Escalating Conflicts: A Social Psychological Analysis. New York: Springer- Verlag.

Brockner, J., Shaw, M. C., & Rubin, J. Z. 1979. Factors affecting withdrawal from an escalating conflict: Quitting before it's too late. Journal of Experimental Social Psychology, 15: 492-503.

Davis, D. D. & Holt, C. A. 1993. Experimental economics. Princeton, NJ: Princeton University Press.

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., Sandefur, C. A., & Rogers, A. C. 1990. De-escalation of commitment in oil exploration: When sunk costs and negative feedback coincide. Journal of Applied Psychology, 75: 721-727.

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.

Goltz, S. M. 1992. A sequential analysis of continued investment of organizational resources in non-performing courses of action. Journal of Applied Behavioral Analysis, 25, 561-574.

Goltz, S. M. 1993. Examining the joint roles of responsibility and reinforcement history in recommitment. Decision Sciences, 24, 977-994.

Heath, C. (1995). Escalation and de-escalation of commitment in response to sunk costs: The role of budgeting in mental accounting. Organizational Behavior and Human Decision Processes, 62 (1), 38-54.

Heath, C. & Soll, J. (in press). Mental accounting for costs: The budgeting process for consumer expenses. Journal of Consumer Research.

Kanodia, C., Bushman, R., & Dickhaut, J. (1989). Escalation errors and the sunk cost effect: An explanation based on reputation and information assymetries. Journal of Accounting Research, 27, 59-77.

Larrick, R. P., Morgan, J. N., Nisbett, R. E. 1990. Teaching the use of cost-benefit reasoning in everyday life. Psychological Science, 1 (6): 362-370.

McCain, B. E. 1986. Continuing invesment under conditions of failure: A laboratory study of limits to escalation. Journal of Applied Psychology, 71: 280-284.

Northcraft, G. B., & Wolf, G. 1984. Dollars, sense, and sunk costs: A life-cycle model of resource allocation decisions. Academy of Management Review, 9: 225-234.

Platt, J. R. 1964. Strong inference. Science, 146: 347-353.

Prendergast, C. & Stole, L. 1996. Impetuous youngsters and jaded oldtimers. Working paper, University of Chicago.

Simonson, I. & Staw, B. M. 1992. De-escalation strategies: A comparison of techniques for reducing commitment to losing courses of action. Journal of Applied Psychology, 77: 419-426.

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. 1996. The escalation of commitment: An update and appraisal. In Z. Shapira (Ed.) Organizational Decision Making. Cambridge, England: Cambridge University Press.

Staw, B. M., & Fox, F. 1977. Escalation: Some determinants of commitment to a personally chosen course of action. Human Relations, 30: 431-450.

Staw, B. M., & Hoang, H. 1995. Sunk costs in the NBA: Why draft order affects playing time and survival in professional basketball. Administrative Science Quarterly, 40: 474-494.

Staw, B. M., & Ross, J. 1980. Commitment in an experimenting society: A study of the attribution of leadership from administrative scenarios. Journal of Applied Psychology, 65: 249.

Staw, B. M., & Ross, J. 1989. Understanding behavior in escalation situations. Science, 246: 216-220.

Whyte, G. 1986. Escalating commitment to a course of action: A reinterpretation. Academy of Management Review, 11: 311-321.

Zelizer, V. A. 1993. Making multiple monies. In Richard Swedberg (Ed.) Explorations in Economic Sociology: 193-212. New York: Russell Sage Foundation.

Footnotes

TABLE 1

Study 1: Experimental stimuli and results.

Experimental Design					Predictions		Results	
Case	Probability Distribution	Payoff	Purchase Price	NPV [project]	Optimal stop	MB upper bound	% purchase	Avg. stop time
1	5, 15, 25	67	0	0.1	1	6.7	*	3.9
2		67	3	-2.9	1	6.4	.26	5.5
3		67	6	-5.9	1	6.1	.30	4.4
4		80	0	3.8	4	8.0	*	4.9
5		80	3	0.8	4	7.7	.56	3.6
6		80	6	-2.2	4	7.4	.50	4.1
7		100	0	14.9	8	10.0	*	5.8
8		100	3	11.9	8	9.7	.95	5.9
9		100	6	8.9	8	9.4	.56	5.0
10	9, 20, 31	67	0	8.2	6	6.7	*	3.5
11		67	3	5.2	6	6.4	.83	4.0
12		67	6	2.2	6	6.1	.38	3.5
13		80	0	18.1	10	8.0	*	4.5
14		80	3	15.1	10	7.7	.96	4.2
15		80	6	12.1	10	7.4	.56	4.6
16		100	0	35.9	19	10.0	*	4.9

17	100	3	32.6	19	9.7	.97	4.5
18	100	6	29.9	19	9.4	.97	4.8

Notes. "*" Subjects did not make a purchase decision when the purchase price was zero. The Optimal Stop column lists the first time the estimated probability of success falls below cost/payoff. At this point the NPV of an additional investment becomes negative (see appendix for sample calculations). The NPV column lists the average payoff produced by each case if the subject correctly stops investing after they reach the optimal stopping point.

TABLE 2

Study 1: Predicting investment decisions

Variable	Coef.
Number of failures	-.126**
	(.025)
Number of failures^2	.008**
	(.001)
Purchase price (0->3)	-.203*
	(.090)
Purchase price (3->6)	-.162^
	(.099)
Payoff (67->80)	.112
	(.097)
Payoff (80->100)	.415**
	(.090)
Probability = 5, 15, or 25%	-.214**
	(.077)
Constant	2.84**
	(.118)
Pseudo R2	.017
Number of observations	10419
Chi2(7)	92.87

<u>Note</u>. ^$p \leq .10$, * $p<.05$, ** $p<.01$

TABLE 3

Study 2: Experimental stimuli and results

Experimental Design					Predictions		Results	
Case	Probability Distribution	Payoff	Purchase Price	NPV [project]	Optimal stop	MB upper bound	% purchase	Avg. stop time
1	5 or 45	70	0	10.6	3	7	*	3.6
2		70	10	0.6	3	6	.41	2.9
3		70	20	-9.4	3	5	.20	3.2
4		90	0	20.6	4	9	*	4.0
5		90	10	10.6	4	8	.87	4.1
6		90	20	0.6	4	7	.34	3.5
7		150	0	55.5	6	15	*	4.9
8		150	10	45.5	6	14	.95	4.6
9		150	20	35.5	6	13	.89	4.5
10	10 or 40	70	0	14.6	5	7	*	3.0
11		70	10	4.6	5	6	.43	3.1
12		70	20	-5.4	5	5	.21	3.2
13		90	0	29.1	9	9	*	3.7
14		90	10	19.1	9	8	.90	3.2
15		90	20	9.1	9	7	.37	2.8

Notes. "*" Subjects did not make a purchase decision when the purchase price was zero. The Optimal Stop column lists the first time the estimated probability of success falls below cost/payoff. At this point the NPV of an additional investment becomes negative (see appendix for sample calculations). The NPV column lists the average payoff produced by each case if the subject correctly stops investing after they reach the optimal stopping point.

TABLE 4

Study 2: Predicting investment decisions

Variable	All Trials	Early vs. Late Trials: Early (1 & 2)	Late (4 & 5)	Below vs. Above Median Correct Purchase Decisions: Below Med.	Above Med.
Number of failures	-.069**	-.021	-.177**	-.081**	-.074**
	(.016)	(.022)	(.032)	(.027)	(.020)
Number of failures^2	.007**	.004**	.015**	.010**	.006**
	(.001)	(.001)	(.002)	(.002)	(.001)
Purchase price (0->10)	-.180**	-.213**	-.157^	-.169*	-.228**
	(.052)	(.080)	(.086)	(.082)	(.070)
Purchase price (10->20)	-.072	-.039	-.282**	-.113	-.031
	(.065)	(.096)	(.108)	(.120)	(.078)
Payoff (70->90)	.245**	.295**	.216*	.369**	.162*
	(.053)	(.082)	(.086)	(.087)	(.067)
Payoff (90->150)	.447**	.329**	.636**	.306**	.573**
	(.064)	(.097)	(.106)	(.100)	(.086)
Probability=5 or 45%	-.125*	-.200*	-.054	-.067	-.161*
	(.053)	(.082)	(.085)	(.084)	(.068)
Constant	1.96**	2.00**	1.05**	1.82**	2.10**
	(.065)	(.101)	(.109)	(.100)	(.087)
Pseudo R2	.029	.023	.047	.038	.025
Number of obs.	19802	8896	7197	8489	11313
Chi2(7)	399.19	135.80	242.53	222.58	199.50

Note. ^ p≤.10, * p<.05, ** p<.01

[1] Subjects who did not wish to participate in the exercise or who were unable to log onto the computer for the computer exercise were given an alternate homework assignment.

[2] To decide whether to purchase the project, subjects should calculate the Net Present Value of the project by drawing out the decision tree for several steps. Although the expected probability of success drops after one failure, the expected value of the second investment is typically still positive. The NPV is determined by summing the value of the option to invest at each stage until the probability drops so low it no longer makes sense to invest. To model purchase decisions, we compared the multiple-step NPV rule against a simpler rule that assumes people only play out the decision one step (OS)--i.e., they assume that the probability falls in the middle of the distribution. In a logistic regression, OS predicted purchase decisions ($b=.408$, $z=6.96$, $p<.000$) but NPV does not ($b=.001$, $z=0.07$, ns). Thus, subjects' decisions tend to be extremely myopic. Because they downplay the value of options to invest after the first investment, subjects refuse to purchase many projects they should.

[3] A reinforcement story fails to account for two features of this data--one is the evidence of forward-looking calculation, the other is the sign on the sunk cost term. Strictly speaking, a reinforcement story should predict that people would escalate commitment in response to the sunk cost because this investment "succeeds" in moving them to the next stage of the process.

[4] By simplifying the probability distribution we made it easier to calculate the NPV for the purchase decision. However, in predicting whether people invest the first time, although

the coefficient on NPV is now positive ($b=.028$, $z=8.31$, $p<.000$), it is still several times smaller than the OS coefficient ($b=.139$, $z=23.47$, $p<.000$).