

Stuck in the Past: Why Managers Persist with New Product Failures

In this research, the authors examine the phenomenon of escalation bias in the context of managing new product introductions. In particular, they identify three general paths—Decision Involvement Inertia, Decision Involvement Distortion, and Belief Inertia Distortion—that can lead managers to escalate their commitments. The authors test the relative strength of these paths in driving observed escalation behavior. The results show that involvement with the initial decision, a key construct in numerous explanations for escalation behavior (e.g., agency theory, self-justification), is not a necessary condition to induce commitment to a losing course of action (i.e., escalation bias). Rather, the authors find that the driving force behind escalation behavior is improper use of initial positive beliefs in the face of negative new information. This insight has implications for the groundwork necessary for organizations to design systems, policies, and procedures to help them avoid the trap of escalation bias that is often associated with major strategic decisions.

Marketing managers are often asked to make major decisions, even though the information available to them at the time of the decision is incomplete and uncertain. Typical examples of such decision situations are launching a new product, initiating an everyday-low-pricing or a “no-promotions” strategy, opening a new channel of distribution, and initiating a major new media campaign. Given that the decision environment is often highly uncertain and the financial stakes large, it is not surprising that firms often revisit such decisions after they obtain new information to determine whether the firm should continue with, modify, or terminate the initially chosen course of action.

It is well documented that when managers who are publicly committed to a course of action are asked to reevaluate that action, they tend to remain committed to it even when the new information indicates that the action should be terminated. See, for example, studies of the Vietnam War and Desert Storm (Lipshitz 1991), Expo 86 (Ross and Staw 1986), the Apollo moon missions (Mitroff 1974), the Campeau–Federated merger, the coffee wars between Philip Morris and Procter & Gamble (Bazerman and Neale 1992), the National Basketball Association draft (Staw and Hoang 1995), information technology projects (Keil 1995; Keil, Mann, and Rai 2000), and a plethora of experimental investigations (Arkes and Blumer 1985; Boulding, Morgan, and Staelin 1997; Brockner and Rubin 1985; Staw 1976). This stylized phenomenon, often referred to as “escalation of commitment,” can be disastrous for firms, especially in today’s hypercompetitive markets, which require fast and accurate reactions and adaptability on the part of companies.

Eyal Biyalogorsky is Associate Professor of Management, Graduate School of Management, University of California, Davis (e-mail: eyalog@ucdavis.edu). William Boulding is Professor of Marketing (e-mail: bb1@duke.edu), and Richard Staelin is Edward and Rose Donnell Professor of Business Administration (e-mail: rstaelin@duke.edu), Fuqua School of Business, Duke University.

Decisions about new products seem particularly susceptible to the problem of escalation of commitment. Schmidt and Calantone (1998) show that innovative projects tend to generate a high level of commitment from managers, often leading to ongoing investment in failing projects. Industry practitioners echo this sentiment. The director of the New Products Showcase and Learning Center stated, “It sometimes takes more courage to kill a product that’s going nowhere than to sustain it” (Lukas 1998, p. 44). Still, killing a new product can be the best course of action, as suggested by comments from a former chief executive officer of Eastman Kodak when speaking to participants in a senior executive education program. He indicated that he considered a new product team a “success” when members appropriately decided to *stop* investing in a new product idea and labeled the project “the world’s fastest failure.”

We use the new product introduction setting to explore the reasons managers tend to stay with losing courses of action and thus forgo the best course of action. We begin by developing a model of how managers combine specific information about the potential product and their prior business experience to decide whether to launch a new product and, if they choose to launch the product, how they decide whether to continue with the venture after receiving negative new information. We then estimate the parameters of this model using data that we obtain from an experiment in which managers are asked to make an initial product-launch decision and a subsequent reevaluation decision. Using the resulting estimates, we provide empirical support for and against numerous previous explanations of escalation of commitment. In turn, this enables us to address the important question of how firms can reduce managers’ tendency to stay with losing courses of action.

Paths to Escalation

To make the preceding discussion more specific, imagine that a manager is provided with a detailed analysis of the potential viability of a new product. This analysis contains

information about specific aspects of the market—the distribution of possible initial market shares, industry growth rates, likely responses by competitors, and summary financial data, such as the anticipated net present values (NPVs) of the venture. The manager uses all of this information to form beliefs about what is likely to happen and, on the basis of these beliefs, comes to an initial decision about whether to launch the new product.

Assume that the manager in our example decides to introduce the product. Moreover, assume that this manager is assigned to lead the product launch so that he or she is publicly connected to the initial launch decision. Two years later, the firm decides to reevaluate the viability of the product. The analyst again provides the manager with information, this time based on new data collected since the launch decision. Among other things, the new information contains updated figures for realized market share, new estimates of industry growth rates, and newly anticipated competitive responses. This information is also summarized in terms of the anticipated distribution of NPVs, assuming that the firm continues with the product. In addition, the analyst provides information about the expected NPV of exiting the market and using the assets associated with the venture in another application. Finally, assume that the “normative” decision (i.e., the decision based on the expected outcomes associated with sticking with the product relative to the expected alternative use of funds) is to pull the plug on the product and reinvest the remaining assets. When the manager is asked either to stay the course or to reinvest the assets in the new application, he or she decides to stay the course, demonstrating escalation of commitment.

Why might escalation bias occur? The extant literature can be organized around three general paths that lead to this occurrence. The first path implicitly or explicitly assumes that because the decision maker was publicly involved with the first decision, this involvement directly influences the second decision. Returning to our example, this means that the manager continues with the product launch despite being fully aware of the negative new information because of his or her involvement with the initial decision. This may happen because the manager fears “losing face” if the initial decision is changed (Brockner, Rubin, and Lang 1981) or because of the effort required to justify the change (Fox and Staw 1979). Another possibility is that the manager uses the NPV assessment associated with the first decision as a reference point and consequently perceives a lower assessment as a loss. Prospect theory posits that people tend to become more risk seeking in the domain of losses (Kahneman and Tversky 1979). Framing the issue as a loss may lead the manager to accept more risks than normal and therefore stick with the initial launch decision (Arkes and Blumer 1985; Bazerman 1984; Whyte 1986). In all of these explanations, the manager’s involvement with the initial decision leads to inertia and reluctance to change, even though the negative new information may be correctly perceived. Thus, we refer to this path as “Decision Involvement Inertia.”

The manager’s involvement with the initial decision again plays a critical role for the second major escalation path. Here, involvement with the first decision does not affect the second decision directly; instead, it improperly


influences the belief structure underlying the second decision (Bacharach, Bamberger, and Mundell 1995; Conlon and Garland 1993; Festinger 1957; Samuelson and Zeckhauser 1988; Silver and Mitchell 1990; Wicklund and Brehm 1976). The manager, who was involved with the first decision, misinterprets the new information so that it is in concert with his or her initial decision, resulting in biased beliefs that, in turn, lead to the decision to stay the course. Therefore, we call this path “Decision Involvement Distortion” because the manager’s involvement with the initial decision leads to a distortion in the perception of the new information.

The third path also concerns biased perceptions of the new information. However, in this case, the bias comes from the person distorting and/or improperly weighting the new information to conform to his or her initial beliefs, not from his or her involvement with the initial decision. Thus, returning to our example, this mechanism suggests that it is not necessary for the manager to have made the initial product introduction decision to exhibit commitment escalation subsequently. Rather, it is necessary only that the manager has formed positive beliefs about the viability of the initial product launch, even with no public involvement in the launch decision, to induce escalation. In other words, this bias does not occur because the manager was involved in the initial decision but rather because he or she does not sufficiently adjust prior beliefs and thus chooses to continue the original course of action. Note that the mechanism for these biased beliefs is similar to that associated with a confirmatory bias or sticky prior beliefs (Bolton 2003; Jonas et al. 2001; Klayman 1995; Schwenk 1986). Thus, we refer to this path of bias as “Belief Inertia Distortion.”

In summary, the literature on escalation can be organized around three basic mechanisms (paths) that lead to escalation of commitment. We summarize these mechanisms in Figure 1. Two of the mechanisms rely on the maker of the second decision being involved in the first decision. The third relies only on the decision maker initially having positive beliefs about the factors that led to the first decision. All three mechanisms can cause managers to exhibit escalation behavior, but they differ with respect to the process that causes the behavior. Moreover, they differ in terms of their implications regarding how to prevent escalation of commitment.

Determining which paths are operative is not an easy task. All the paths involve complex relationships between second-stage outcomes and first-stage explanatory variables. It is not enough to demonstrate simple correlations between first- and second-stage variables. For example, to test Decision Involvement Inertia, it is necessary to show that involvement with the initial decision influences the second decision, after we control for the effects of the updated second-stage belief structure. Likewise, testing Decision Involvement Distortion requires us to show that the manager’s updated belief structure affects the continue/stop decision and then that involvement with the initial go/no-go decision affects the manager’s updated belief structure, after we control for the possible normative and nonnormative effects of prior beliefs on updated ones. Finally, to test Belief Inertia Distortion, we must show that the manager’s

FIGURE 1
Paths Leading to Escalation



updated belief structure affects the continue/stop decision and that the manager's prior beliefs have a nonnormative effect on his or her updated beliefs beyond any predicted normative effect on the updated beliefs, independent of any involvement with the initial decision.

The preceding discussion makes it clear that if we want to disentangle the potential influences of the first decision from the influence of prior beliefs, we must explicitly model and measure the belief-updating process. In Heath's (1995) review of the literature, he expresses concern that most escalation studies fail to control for initial beliefs. Consequently, we take extra care to address this issue.

Given the complexities associated with disentangling our effects of interest, we chose to conduct an experiment in which we manipulate the manager's involvement with the initial decision, the level of the manager's initial (prior) beliefs, and the degree of (lack of) market success subsequent to the initial product launch. We then measure initial and updated beliefs and observe the second decision. This enables us to estimate the parameters of our belief-updating model and, on the basis of those parameters, to assess the impact of each of the three potential paths on escalation of commitment. Formally, we test the following three hypotheses:

H_1 (Decision Involvement Inertia): When the decision maker's updated belief structure is held fixed, being highly involved with the decision to choose an initial course of action has a direct, positive effect on the next decision, thus producing escalation behavior.

H_2 (Decision Involvement Distortion): Being highly involved with the decision to choose an initial course of action biases the decision maker's updated belief structure in the direction of supporting the initial course of action. In turn, this biased updating process increases the likelihood that

the decision maker will stay with the initial course of action.

H_3 (Belief Inertia Distortion): Independent of any involvement with the initial decision, a decision maker's initial positive belief biases his or her updated belief structure in the direction of the initial belief structure, thus producing escalation behavior.

Note that these three hypotheses are not mutually exclusive and therefore can operate simultaneously. Furthermore, any empirical tests of these hypotheses must distinguish between normative and nonnormative updating of beliefs. Consequently, in the next section, we describe a model for belief updating and, in the process, describe how we can distinguish between normative and nonnormative updating effects in the model.

Belief-Updating Model

Our model of the belief-updating process acknowledges that a marketing manager normally receives different and, perhaps, conflicting pieces of information about many aspects of the environment. We assume that the manager integrates this information into an overall assessment that guides a decision. Therefore, we model not only the manager's overall assessment and updating process but also the process of updating beliefs about specific underlying aspects of the environment that lead to the overall assessment.

We begin with the conceptualization of managerial decision making that Boulding and colleagues (1994) propose. Their model assumes that all beliefs can be represented by probability distributions in which the mean of the distribution is the person's belief about the most likely occurrence and the variance of the distribution captures the person's uncertainty about the particular belief. Boulding and colleagues posit two different types of underlying beliefs and one summary belief. The underlying beliefs refer to the specific levels of the market factors that influence performance of the product (also referred to as "what" knowledge) and beliefs about the market structure (i.e., how each market-factor belief is weighted; also referred to as "how" knowledge). This conceptualization has the manager combine the two components (factor levels and weights) to form "summary" knowledge, that is, a summary belief or overall evaluation. For example, for a new product decision, market-factor beliefs would include all the factors that are believed to have an influence on the project's overall success. These market factors could include items such as the obtained market share, the price of the new product offering, market growth, competitors' responses, and the cost of capital. The manager then uses knowledge of the relevant market structure to establish weights for how each market-factor belief affects the project's overall success, which is captured by the manager's summary belief. Boulding and colleagues assume that it is this summary belief that ultimately determines the manager's decision.

We formalize this process as follows: Let SB_{it} be the mean of manager i 's summary belief about the course of action at time t , and let FB_{ikt} be the mean of the manager's belief about the level of the k th market factor that affects

overall market performance. Let $g_{it}(\cdot)$ be the manager's mental mapping of how these market factors combine to result in overall performance:

$$(1) \quad SB_{it} = g_{it}(FB_{ikt}).$$

We assume that the manager's mental mapping, $g_{it}(\cdot)$, can be closely approximated by a Fishbein-style linear formulation (Ajzen and Fishbein 1977). Thus, in Equation 2, we replace the function $g_{it}(\cdot)$ with a vector of importance weights (w_{ikt}) that captures the manager's perception of the importance of each market factor:

$$(2) \quad SB_{it} = \sum_k w_{ikt} FB_{ikt}.$$

Conceptually, the assumption of a linear mapping is supported by numerous studies that have shown that a linear model provides a robust, parsimonious, and paramorphic representation of how people combine multiple stimuli when forming an overall evaluation (e.g., Ajzen and Fishbein 1977; Bettman 1979; Einhorn 1970; Johnson and Meyer 1984; Keeney and Raiffa 1976). Equally important, the data-generating process we used to create the information presented to participants was linear. Thus, participants who correctly perceived the relationships in the data should have used a linear mapping. Empirically, we tested several nonlinear specifications for the participants in the study and found no indication that a nonlinear specification was necessary or appropriate.

Consider now SB_{i2} , which represents the manager's summary belief at the continue/stop decision point. Equation 2 suggests that there are three potential sources of bias of this summary belief: (1) The manager's market-factor beliefs, FB_{ik2} , could differ from normatively correct levels; (2) the manager's importance weights, w_{ik2} , could be distorted compared with normatively correct weights; and (3) when forming a summary belief, the manager could inject extraneous, irrelevant information that Equation 2 does not capture. For example, if the manager had an argument with his or her children over breakfast, an ensuing bad mood could affect his or her beliefs. Theoretically, we can capture this effect in Equation 2 by augmenting the market-factor-beliefs vector to include the manager's mood. As researchers, we have access to and can measure only a limited set of possible factors, and thus any estimation must try to control for unobserved factors.

Next, we describe how to determine whether any of these sources of bias affect the manager's updated beliefs at the continue/stop decision stage. In doing so, we are interested not in just any difference between the manager's beliefs and the normatively correct beliefs but rather in the possible systematic biasing effects of the manager's first-stage (go/no-go) beliefs (SB_{i1} , FB_{ik1} , w_{ik1}) and the initial decision on his or her second-stage beliefs.

Updating Market-Factor Beliefs

We begin by modeling how a manager updates the mean of his or her market-factor beliefs (see Equation 3). The model includes three terms: (1) the effect of any relevant new information that the manager receives about market factors

(NI_k), (2) the effect of the manager's prior market-factor belief (FB_{ik1}), and (3) the effect of the manager's prior summary belief (SB_{i1}).

$$(3) \quad FB_{ik2} = \alpha_{1k} NI_k + \alpha_{2k} FB_{ik1} + \alpha_{3k} SB_{i1}.$$

Note that we do not explicitly model the effects of the initial decision on the updating process. Rather, we hypothesize that if the person's involvement with the initial decision affects the updating process, the coefficients in Equation 3 will systematically vary depending on the level of involvement. In other words, involvement interacts with the right-hand side of Equation 3. In addition, note that if we assume that the manager's beliefs can be described by a normal distribution, the first two terms are consistent with normative updating, where the coefficients α_{1k} and α_{2k} reflect the relative reliabilities of the new information and the initial belief, respectively. Conversely, the third term represents the hypothesized nonnormative effect of the initial summary belief on the updating of market-factor beliefs. The rationale behind this last term is the well-documented phenomenon that people tend to perceive new information as compatible with their prior beliefs, exhibiting a confirmatory bias (Hoch and Ha 1986). A significant, positive coefficient on the initial summary belief is evidence of biased updating of the market-factor belief.¹

Updating Importance Weights

Normatively, importance weights should change only if the manager receives new information that specifically indicates that the model of how market factors interact has changed. However, we conjecture that new information that pertains only to the levels of market factors may still induce a change in importance weights. Specifically, if the manager's assessment of a market-factor level increases (decreases), we suggest that there will be an associated increase (decrease) in the importance that he or she accords to that market factor. Such a pattern would reflect "elastic" importance weights (Boulding, Kalra, and Staelin 1999; Hsee 1995). Specifically, the weights would increase in value when new information supports a positive initial belief and decrease in value when new information contradicts the positive initial belief. In Equation 4, the third term captures this reasoning; if the difference between the manager's new market-factor belief and prior market-factor belief is positive (i.e., the new information is perceived as more positive), the importance weight for this factor increases. If the difference is negative, the reverse occurs. As in Equation 3, the first two terms in Equation 4 represent the expected normative effects.

$$(4) \quad w_{ik2} = \beta_{1k} NI_k + \beta_{2k} w_{ik1} + \beta_{3k} (FB_{ik2} - FB_{ik1}).$$

¹Failing to find a significant effect on the initial summary belief does not mean that the updating is normative, because the other two coefficients may still be different from their normative values. Here, however, we restrict our attention to the more stringent test of the significance of the prior summary-belief coefficient.

We estimate a slightly different version of Equation 4. It will become clearer when we describe the experiment that we do not manipulate or provide any new information that indicates a change in the market structure (the importance weights). Thus, β_{1k} should be zero, and we estimate only the last two parameters in Equation 4:

$$(5) \quad w_{ik2} = \beta_{2k}w_{ik1} + \beta_{3k}(FB_{ik2} - FB_{ik1}).$$

In this “special case updating” given in Equation 5, $\beta_{2k} = 1$ and $\beta_{3k} = 0$ would be consistent with normative updating, and deviations from these parameter values would reflect biased updating.

Updating the Summary Belief

If the manager’s market-factor beliefs and/or importance weights at the time of the continue/stop decision are biased, then according to Equation 2, the summary belief that the manager constructs from those underlying beliefs will be biased as well. However, it is possible that the manager’s updated summary belief is not constructed in this manner. Rather, the manager might update the former summary belief on the basis of some macrolevel evaluation of the veracity of the new information. Such a model is consistent with normative Bayes-rule updating, but at a higher level of aggregation. Because the two possible updating models are based on the same information, it follows that the two should agree in a well-calibrated model. On the basis of this principle, we can test for the existence of bias at the summary-belief level. Consider a general model that nests our two alternative models of summary-belief updating:

$$(6) \quad SB_{i2} = \gamma_1NI + \gamma_2SB_{i1} + \gamma_3\Sigma w_{ik2}FB_{ik2}.$$

If the mean of the summary belief is updated by means of Bayes’ rule with the prior summary belief and the new information, the γ_3 term in Equation 6 should provide no additional explanatory power. Conversely, if the updated belief is constructed on the basis of the updated underlying factor beliefs, the first two terms in Equation 6 should provide no additional explanatory power. If all three terms simultaneously provide explanatory power for the updated summary-belief measures, bias is entering the summary belief.

In summary, estimation of Equations 3, 5, and 6 enables us to test directly for the existence of bias in the updating process and, therefore, the mechanisms underlying escalation behavior (H_1 – H_3).

Method

Before describing our study, we briefly review the requirements needed to test adequately for the occurrence of the three proposed mechanisms underlying escalation behavior. First, as Northcraft and Wolfe (1984, p. 227) point out, we need to “examine decision making situations in which commitment of further resources is explicitly economically inadvisable.” Heath (1995, p. 52) also mentions this need and further claims that “previous demonstrations of ‘escalating commitment’ with explicit investments did not provide relevant information about marginal costs and bene-

fits.” To address these concerns, we provide the decision makers with explicit information about the marginal costs and benefits of both continuing with and discontinuing the course of action. Furthermore, the information shows that to continue is economically inadvisable. Second, because involvement with the initial decision and initial beliefs are key constructs in the three paths, it is necessary to manipulate both involvement with the first decision and initial beliefs. This is not straightforward, because measuring beliefs requires the person to be somewhat involved with the initial decision. However, using several different manipulation checks, we show that even though all our participants are somewhat involved with the initial decision, our involvement manipulation leads to significant differences in their involvement level. Finally, there is the issue of disentangling normative belief updating from biased updating. Here, we use our theoretical model to partition the two effects.

More specifically, we conducted an experiment in which we asked participants to take the role of management in a new product introduction and reevaluation situation. The study involved manipulating three variables: the degree to which participants were involved with the initial decision (high involvement versus low involvement); the extent to which the initial information provided a positive initial indication about future profits and, thus, positive NPV (very positive versus positive); and the degree to which subsequent information was nonconfirmatory (very negative versus negative). Together, these three manipulations enable us to tease apart the various effects related to H_1 – H_3 .

We used a new product setting similar to our initial example that appears in a previous study (Boulding, Morgan, and Staelin 1997) in which senior-level managers exhibited a marked tendency to continue with the product. Specifically, between 45% and 100% of the managers continued with the product, depending on the particular condition. In contrast, when Boulding, Morgan, and Staelin (1997) gave the financial information associated with the reevaluation decision to a group of 20 academic experts, all chose to discontinue the project. This latter finding provides an important normative benchmark, suggesting that continued commitment to the project in our setting is indeed evidence of escalation bias.

Participants

We obtained responses from 142 participants who completed the experiment. Approximately two-thirds of these participants were MBA students from four different full-time programs, two on the East Coast and two on the West Coast. The remaining participants were midlevel managers who were enrolled in an executive MBA program at one of the East Coast schools. All participants had completed at least one full semester of studies that included instruction in decision making. Because our sample used MBA students, it is logical to ask if this sample can be generalized to the population of senior-level managers who normally make major investment decisions. To address this issue, we replicated the control condition used in Boulding, Morgan, and Staelin’s (1997) study, which used senior-level managers as

participants, with a sample of our participants. In our study, 60% of the participants (12 of 20) in the control condition discontinued the product, whereas 66% (27 of 41) discontinued it in Boulding, Morgan, and Staelin's study.² The difference in stopping rates between the two studies is not significant ($t = .45, p > .30$). This point of comparison suggests that there is no substantial difference in the way participants in the two studies responded to the information provided, increasing our confidence that we can generalize our findings to the population of senior-level managers sampled in Boulding, Morgan, and Staelin's study.

The Case

Participants were asked to complete the Quality Valve Company case.³ The first two pages of the case describe the new product, the history of the market, the company, the primary competitor, and the investment associated with an improved valve targeted at the large-truck emissions-control market. The specific initial decision problem was whether the organization should invest \$2.5 million in new machinery and launch the improved valve.

As part of the case description, participants learned that the financial success of the new product venture depended on three key uncertainties: the share of the market the new product garners after launch, the industry's growth rate, and the competitive-entry response of the major competitor to the launch of the new product (which was captured by the timing and effectiveness of an innovative competitive offering in the future).

Participants initially received one of two different analyst's reports (positive or very positive) that included projections of all the key market factors along with the distribution of estimated NPVs, assuming that the product is launched. They also were given the expected NPV and could easily determine from the NPV distribution the probability that the project would yield a positive NPV. In addition, the analyst's report provided estimates of the relative importance of each of the key uncertainties in determining the product's financial success; these estimates are 70% for market share, 22% for competitive-entry response, and 8% for market growth. After reading the case description and the analyst's report, participants made a launch/no-launch decision (high-involvement condition only) and described their beliefs about the expected levels of relevant market factors and their mean overall summary beliefs.

Next, participants in the high-involvement group who suggested launching the product and all the participants in the low-involvement group were given the second part of the case. Specifically, they received negative new information about the product's performance during its first two

years on the market. The information included realized values for the three key market uncertainties and was presented in a balanced fashion. Thus, the realized market share was significantly less than anticipated, the realized market growth rate was greater than anticipated, and the information about the major competitor's entry response was neutral in that no new information was given on either the timing or the effectiveness of its future offering. In addition, participants were told that an uncorrectable production problem explained the lower-than-expected market share but that the sales force and research-and-development groups were very happy about the product introduction. Nonetheless, because market share was the most important factor in determining profitability and the production problem was not correctable, on balance, the information was negative. Moreover, the analyst's revised (normatively updated) financial report predicted an expected NPV from future cash flows that was much lower than initially projected. More important, this NPV forecast was substantially lower than the guaranteed cash value associated with the alternative of dropping the product and selling the equipment. Thus, as Heath (1995) and Camerer and Weber (1999) stress, the marginal costs and benefits of the alternative courses of action were made explicit to participants before they made the second decision. Specifically, the normative decision for the organization based solely on the new financial information and the concept of expected values was to withdraw the product and sell the equipment. This choice was given to participants. After receiving the new information, participants were asked to select one of two alternatives: continue with the product or withdraw it and sell the equipment.

Experimental Manipulations

Initial information. The initial information provided to participants was either positive or very positive. In the very positive condition, the expected market share following the launch was higher, as was the analyst's assessment of the expected NPV from the product. In both conditions, the information was sufficiently positive to result in the vast majority of participants choosing to launch the product.

New information. The new information provided to participants was either negative or very negative. In the very negative condition, participants were given a realized market share after the launch that was lower than the market share given to participants in the negative condition. In addition, the analyst's assessment of the expected NPV from the product was lower in the very negative condition. In both conditions, the information about realized market share, the uncorrectable production problem, and the analyst's revised NPV assessment declined sufficiently compared with the guaranteed cash-out value of stopping the project to make stopping the project the normatively correct decision.

Involvement manipulation. The theoretical explanations underlying the two decision involvement paths suggest that the influence of the initial decision on the subsequent continue/stop decision is the result of the decision maker's per-

²Although this control cell replicates Boulding, Morgan, and Staelin's (1997) study, the levels of information that they provide differ slightly from those used in this study. Consequently, these data cannot be used for further comparisons with other data collected in this research.

³The case is available on request.

sonal involvement in the initial decision (see, e.g., Staw and Ross 1987). The literature on escalation bias has manipulated involvement with a decision by

1. Manipulating the level of participation the second-stage decision maker had with the initial decision by having him or her participate (or not) in making the initial decision (Barton, Duchon, and Dunegan 1989; Staw 1976), making him or her primarily responsible for the decision or responsible only for making a recommendation to be reviewed by others (McAllister, Mitchell, and Beach 1979), and making him or her the company president or someone with no connection to the company (Arkes and Blumer 1985);
2. Manipulating the personal stake the decision maker had in the success of the initial course of action (Arkes and Blumer 1985) by, for example, implying that there were or were not job security implications (Fox and Staw 1979); and
3. Manipulating how much the decision maker expected to need to justify the decision to others (Bobocel and Meyer 1994; Simonson and Staw 1992).

We used the first two methods to manipulate involvement with the initial decision. Specifically, high-involvement participants were put in the position of making the initial decision to launch or not to launch the product after receiving the previously discussed information. These participants were instructed that if they launched the product and it succeeded, they would likely be on the fast track to a senior management position. They were told that choosing not to launch the product would not hurt their careers but that they would lose this particular opportunity to be on the fast track. Those who chose to launch the product were subsequently told that they were appointed as the product's brand manager. Consequently, they had overall responsibility for everything that happened, beginning with the decision to launch the product.

Conversely, low-involvement participants were told that the boss, who had a highly centralized decision-making style, would make the launch decision. However, the boss wanted to know the manager's beliefs about the key uncertainties associated with the launch. Furthermore, participants in this condition were told that the boss wanted unbiased opinions, and thus all evaluations about these key factors were to be provided anonymously. After they provided the requested belief information, these low-involvement participants were told that during the two years the product was on the market they were assigned to an entirely different project and that they were being called back to help make the reevaluation decision. Thus, the high- and low-involvement conditions differed in terms of (1) the level of the decision maker's involvement in the initial launch decision (similar to the method that McAllister, Mitchell, and Beach [1979] used), (2) the level of the decision maker's involvement with the product's performance after launch (Arkes and Blumer 1985), and (3) the potential career implications of the success of the initial decision for the decision maker (Arkes and Blumer 1985).

To ensure that we successfully manipulated involvement with the initial decision, we conducted two independent manipulation checks. We provide details of these two checks in the Appendix. In summary, for all manipulation checks, there were significant differences in the expected directions between the two groups with respect to their

involvement with the initial decision. Still, we acknowledge that the low-involvement situation required the decision maker to think hard about the different aspects of the decision and, in this way, differed from some prior studies in which the decision maker was not provided any information about the first decision. However, as Heath (1995) points out, it is important to capture the participants' beliefs before the first decision is made. This required the participants to know about the decision and form beliefs about the elements of the decision that were relevant at the time. Consequently, all participants were at least somewhat "involved" with the decision. In the "Results" section, we discuss a third "manipulation check" that is based on the participants' actual actions. To foreshadow these results, again we found that the two groups differed, but this time with respect to their decision behavior. This provides additional support for our contention that there are significant differences in the two involvement groups.

Design

We used an incomplete $2 \times 2 \times 2$ design for the study. The low-involvement condition was fixed at the positive level of the initial information (as opposed to very positive), but it was crossed with both levels of new information. Thus, the design included only six of the eight possible cells. We used this design for parsimony because there are no theoretically meaningful predictions about the interaction of involvement with the initial decision and the level of new information.

Measures

Our theory testing revolved around three different measured constructs: market-factor beliefs, importance weights, and summary beliefs. After each decision, we took two summary-belief measures. They were the manager's assessment of expected NPV and the probability of success (the probability that NPV would be positive). We combined these measures after normalizing each measure over the sample to have a mean of zero and a standard deviation of one. We then summed the two measures for each participant to form the summary-belief measure used in our analysis.

We took measures for each of the three market-factor beliefs after both decisions. We measured the market growth rate and market-share beliefs using a nine-point scale anchored between "very negative" and "very positive" (the midpoint was designated as "neutral"). Our use of affective evaluations rather than actual expectations ensured that all participants were on the same "goodness" metric. In this way, a higher rating reflects a higher evaluation even if, for example, a particular manager believes that a 20% market share is a great success whereas another considers a 40% market share a relative failure. For the market-factor belief associated with the major competitor's entry response, we used two measures: the likelihood of competitive launch and the effect of a competitive launch. We measured both on nine-point scales, normalized them, and combined them into a single variable that, as with our summary-belief measure, captured both the probability of the event occurring and the value of the event.

Finally, we measured importance weights for the components relevant to the second decision by asking partici-

pants to specify the relative importance of the three market factors in determining the project's financial success. In particular, participants divided 100 points between the three factors.

Results

Because we are interested in situations that involve two decisions, we discarded data for the 7 high-involvement participants who chose not to launch and for the 5 low-involvement participants who, given their beliefs, presumably would have chosen not to launch.⁴ Of the remaining 130 participants, 68 (52%) decided to stay with the losing course of action, exhibiting behavior that is compatible with an escalation bias. Given this evidence of nonnormative behavior, we next report findings that bear on the strength of the three nonnormative paths based on our framework.

Decision Involvement Inertia Results

The Decision Involvement Inertia path implies that involvement should directly influence the second decision. To test this, we heeded the warnings of Heath (1995) and controlled for summary beliefs. Table 1 presents the results of our logistic regression for three different models. In Model 1, the only independent variables are two of our experimental manipulations—initial information and new information—and our measure of the participant's summary beliefs. Note that summary beliefs are highly significant and completely mediate any effects of the initial and new information. In Model 2, we add a variable that captures the experimental manipulation of involvement with the initial

⁴All 7 of the 98 participants in the high-involvement conditions who chose not to launch the product estimated probability of success to be less than 50%. The average for these seven participants was 46%, as opposed to 73% for those who continued. In addition, 5 of the 44 participants in low-involvement conditions forecasted a negative NPV and/or a less than 50% chance of a positive NPV. Given this initial negative forecast, we assumed that they would not have gone forward with the initial launch and thus eliminated them from further analysis. Furthermore, when we compared this initial decision not to launch for the two involvement groups, holding fixed the level of initial information, we observe that 5 of 59 did not launch in the high-involvement group, and we assume that 5 of 44 would not have launched in the low-involvement group. These proportions are not significantly different.

decision. As is immediately evident, involvement does not reach traditional levels of significance ($p = .12$). Moreover, there are no significant differences between Model 2 and Model 1 in the log-likelihood (122.135 versus 124.609) or in the predicative accuracy (84.3% versus 82.6%). Thus, we find little statistical justification for the Decision Involvement Inertia path.

A question might arise as to why our involvement measure did not affect the second decision, given that numerous previous studies have shown an effect. One explanation is that our involvement manipulation did not provide enough difference between the two groups in terms of their perceptions of being involved with the first decision. However, when we ran a model without the summary-belief measure but with the involvement measure (Model 3), we found that involvement was positively associated with continuing with the initial course of action and that this effect was significant at the .05 level. This result has three major implications: First, it provides one more piece of evidence that our involvement manipulation worked. Second, it reproduces many previous findings on involvement in which researchers did not control for summary beliefs. Third, a comparison of this result with the insignificant involvement effect in Model 2 shows that even though involvement may be associated with escalation of commitment, being psychologically and/or socially involved with the initial decision is not a necessary condition for escalation to occur. Instead, it seems that escalation bias enters the decision through the updated belief structure. Consequently, we turn our attention to the two paths that pertain to distortion of information and, thus, biased updating.

Decision Involvement Distortion Results

Before we present the Decision Involvement Distortion results, note that we tested both distortion paths using some variables that we measured instead of manipulated. This required us to control for background factors, such as participants' individual differences, that could be correlated with both our independent variables and our dependent measure of interest (and thus produce misleading estimates). We did this using an instrumental variable (IV) procedure, in which the selected instruments are known to be free from any individual background factors. The IV procedure replaces the actual value of a measure with its pre-

TABLE 1
Results of Logistic Regression on the Continue/Stop Decision (N = 126)

Variable	Model 1	Model 2	Model 3
Intercept	-4.94* (1.041)	-5.358* (1.09)	-.81* (.395)
Initial information	-.11 (.532)	-.458 (.578)	.388 (.442)
New information	.124 (.439)	.201 (.449)	.41 (.366)
Involvement with the initial decision		.804 (.517)	.85* (.437)
Updated summary belief	.0936* (.02)	.092* (.0193)	

* $p < .05$.

Notes: Standard errors are in parentheses.

dicted value using the selected instruments, thus purging unobserved background factors from the measures. Note that the presence of correlation of our independent variables with unobserved individual differences is testable. Thus, in each equation, we used Hausman's (1978) test to identify whether the IV estimation procedure was required to avoid misleading estimates due to unobserved background factors that differed across participants and were correlated with both our dependent and independent measures. If the test indicated the presence of correlation, we used the IV procedure (this only happened for the summary-belief equation). As a result, our findings based on the use of measured variables can be interpreted as having the same validity as traditional analyses of variance based on manipulated variables.

Decision Involvement Distortion (H_2) suggests that involvement with the initial decision indirectly affects managers' decisions by altering the way they update their beliefs. If this is true, we would observe different patterns for the updating process for the two involvement groups. Consequently, we tested this hypothesis by determining whether the coefficient vectors, both slope and intercept, varied significantly between the two groups for each of the belief-structure updating equations.

We began by testing the summary-belief estimation equation (Equation 6). The appropriate F test indicated that the estimated intercept and slope coefficients did not vary between the two involvement groups for this equation ($F_{4, 119} = .76$). Next, we estimated the updating equations for the three market-factor beliefs for the same two involvement groups (Equation 3). Again, we failed to find any evidence of a difference in updating between the two groups (market share: $F_{4, 122} = 1.51$; competitive-entry response: $F_{4, 127} = 1.19$; and market growth: $F_{4, 122} = 1.53$). Finally, we estimated the updating equation for the importance weights associated with the market-factor beliefs. Here, we estimated a single updating equation (Equation 5) for the importance weights because the weights were constrained

by a constant sum scale and therefore were interrelated. Again, we failed to find significant differences in updating for the two groups ($F_{3, 362} = 1.14$). In summary, we failed to find any difference in updating of beliefs due to involvement with the initial decision, and therefore we found no support for the path associated with Decision Involvement Distortion.

Belief Inertia Distortion Results

Our previously reported logistic analysis suggests that updated summary beliefs are the driving forces behind the second continue/stop decision. Because we found that involvement does not affect the coefficients in the updating equations, we used the combined sample to estimate the three market-factor-belief-updating equations, the importance-weight-updating equation, and the summary-belief-updating equation. In each case, we tested H_3 by examining the nonnormative aspects of the equations.

Market-factor beliefs. We report the results for the market-factor-belief-updating equations in Table 2. These results indicate support for Belief Inertia Distortion. In support of H_3 , we find that the effect of the initial summary belief is positive and significant in both the market-share-updating ($p < .05$) and the market-growth-updating ($p < .01$) equations. In addition, as we expected, the initial summary-belief effect is not significant in the competitive-entry-response-updating equation. Because there was no new information provided about the major competitor's entry response, participants would not be expected to engage in an updating process (biased or unbiased) regarding the competitive response.

Importance weights. We report the results from the importance-weight-updating equation in Table 3. Again, these results demonstrate support for Belief Inertia Distortion. Consistent with H_3 , we found a significant ($p < .001$) increase (decrease) in importance weights for market factors when participants perceived the new information as

TABLE 2
Results of the Market-Factor-Belief-Updating Equations

Variable	Market-Factor Belief		
	Market Share (N = 129)	Market Growth (N = 129)	Competitive-Entry Response (N = 128)
Consistent with Normative Updating			
Intercept	3.047*** (.592)	5.647*** (.646)	3.032*** (.464)
New information	.665** (.244)	-.088 (.258)	-.169 (.227)
Initial market-factor belief	.108 (.075)	.149* (.088)	.348*** (.078)
Consistent with Biased Updating			
Initial summary belief	.72* (.386)	.265** (.102)	-.084 (.361)

* $p < .05$.

** $p < .01$.

*** $p < .001$.

Notes: Standard errors are in parentheses.

TABLE 3
Results of the Importance-Weights Equation
(N = 366)

Variable	
Consistent with Normative Updating	
Initial importance weight	.764* (.029)
Consistent with Biased Updating	
Relative positive or negative new component information	2.234* (.644)

* $p < .001$.

Notes: Standard errors are in parentheses; no intercept regression is present in the model.

more (less) positive than their perception of the initial market-factor information.

Summary beliefs. We present the estimation results of the summary-belief-updating equation in Table 4. These results show that the effect of the initial summary belief is positive and significant ($p < .001$), as is the effect of integration of the participants' updated market-factor beliefs and importance weights ($p < .001$). As we noted in our discussion of Equation 6, for a well-calibrated manager, one but not both of these effects should be significant. Both terms are highly significant, which is a clear indication of bias entering the updated summary belief.

In summary, we find strong support for Belief Inertia Distortion. Moreover, the distortion appears concurrently at multiple levels—that is, at the aggregate level of summary beliefs and at the more disaggregated level of both component beliefs and importance weights.

Alternative Explanations for the Data

Previous theories of escalation, such as self-justification, face saving, and framing of options, have pointed to involvement with the initial decision as the main driver of escalation behavior. Our results are compatible with these previous theories in that we find that high-involvement decision makers are more likely to stay committed to a losing course of action than low-involvement decision makers. However, our results suggest that involvement with the ini-

TABLE 4
Results of the Summary-Belief-Updating Equation
(N = 124)

Variable	
Intercept	34.664* (12.162)
New information	2.871 (30.151)
Initial summary belief ^a	.251** (.066)
$\Sigma W_{ik2}FB_{ik2}^a$.217** (.041)

* $p < .01$.

** $p < .001$.

^aVariable is instrumented.

Notes: Standard errors are in parentheses.

tial decision is not a necessary condition for escalation behavior. Instead, we find that biased belief updating is the driving force behind escalation behavior. Indeed, we find that participants in the low-involvement conditions who formed initially favorable opinions of the new products exhibited the same behavior as participants in the high-involvement condition.

Given the limited role of biased belief updating in previous studies, it might be expected that at least some, if not most, of the escalation effects would operate through the paths associated with decision involvement. However, our results counter much of the extant literature. Consequently, we explore possible alternative explanations for our findings.

Perhaps the simplest explanation for our finding of no effects from Decision Involvement Inertia and Decision Involvement Distortion is that our involvement manipulation did not provide meaningful separation between the two groups. This separation is critical because the explanations associated with these two paths rest on the assumption that the reevaluation decision makers differ in terms of involvement with the initial decision. Importantly, our three independent manipulation checks indicated significant differences between the two groups in terms of both the measured level of involvement and behaviors associated with different levels of involvement. Thus, we reject the idea that the lack of support for the explanations associated with these two involvement paths is due to a failed involvement manipulation.

Another possible alternative explanation for the lack of support for the decision involvement paths is the possibility of reverse causality in some of the paths. In this scenario, involvement/responsibility for the initial decision still acts as the “trigger” for escalation bias. The logic is as follows: Responsibility for the initial decision causes the decision maker to stay the course. Having made this decision, the decision maker modifies his or her beliefs to bring them into accord with the decision. That is, the causal arrow runs from the second decision to the updated beliefs.

To explore the feasibility of this account, we searched for evidence that would be compatible with this reverse causality argument. Specifically, the reverse causality interpretation implies that initial involvement with the decision must be correlated with the updated beliefs after we remove the mediating role of the subsequent decision. However, as we show in our analysis of the Decision Involvement Distortion path, the degree of involvement with the initial decision is independent of updated summary beliefs, market-factor beliefs, and weights. Therefore, the conjecture of reverse causality between the second decision and updated beliefs is incompatible with our data.

Implications

In this article, we provide an integrative framework that organizes the existing literature around three possible paths to escalation behavior. Although our methodology allows for the simultaneous conceptual and empirical existence of all three paths, our results suggest that the driving force behind escalation behavior in our study was not involve-

ment with the initial decision but rather biased belief updating that outweighs initial positive beliefs.

In considering the theoretical implications of our findings, we acknowledge that these findings are based on one laboratory study, albeit with a diverse set of participants. Thus, although we believe that these findings generalize to other populations and other organizational settings (e.g., the National Basketball Association draft; hiring decisions; other strategic marketing decisions, such as selecting an ad campaign or choosing a new channel of distribution), we caution readers to draw their own conclusions about the extent to which our findings generalize. In addition, we acknowledge that all our participants had some involvement with the initial decision, and thus our study differed from some prior studies in which the noninvolved group had no contact whatsoever with the initial decision. That said, our findings suggest that Decision Involvement Inertia (involvement with the initial decision directly affecting subsequent decisions) is not the major driving force behind escalation. This casts doubt on several commonly invoked explanations for escalation, such as a framing effect, agency theory, and public or self-justification.

We acknowledge that this is a highly unexpected conclusion. Because the lack of a Decision Involvement Inertia effect was so surprising, we took extra care in our analysis to ensure that no such effect was being masked by either our method or our analysis. Having taken these steps, we point to two factors in our methodological approach that could account for why our findings differ from much of the existing escalation literature. First, with regard to escalation research, Northcraft and Wolfe (1984), Heath (1995), and Camerer and Weber (1999) point to the importance of having an explicit normative benchmark and note the absence of such benchmarks in most escalation studies. From this starting point, they conclude that though there are a large number of prior studies that examine escalation of commitment, these studies provide little empirical evidence in support of the existence of escalation bias. Therefore, we took great pains to provide information about marginal costs and benefits for each option, and in our setting, it is clear that the continuation decision is nonnormative. Thus, we believe that our study provides clear evidence of nonnormative escalation behavior. Second, Heath (1995) argues for the importance of disentangling the role and measurement of beliefs from the role and measurement of making the initial decision. Our study is the first to do this, enabling us to control for beliefs when testing for the effects of involvement with the initial decision.

We also find no support for the hypothesis of Decision Involvement Distortion (involvement with the initial decision leads to biased belief updating). Although this path has received much less attention in the literature, this result is important because it again shows that involvement with the initial decision is not the root cause for escalation behavior.

Instead, our results point to the powerful influence and complexity of biased belief updating, which weights initial positive beliefs too heavily (the Belief Inertia Distortion path). Note that the biasing influences on belief updating enter the process at many different levels. Specifically, we find that prior beliefs nonnormatively affect not only the

updated market-factor beliefs but also the importance weights placed on those factors and the summary beliefs. Thus, all these sources of bias enter into the updating process concurrently.

Note also that these nonnormative updating results are compatible with Boulding, Kalra, and Staelin's (1999) results. This compatibility is notable because their setting was very different. In that study, consumers were the information processors, and longitudinal evaluations of service quality were used instead of updated evaluations of a new product introduction. Consequently, we conclude that the model of belief updating presented in our article should be useful in understanding belief updating in a wide range of situations. Our findings should be of interest to researchers who study belief formation and updating as well as those who are interested in understanding escalation. In general, the results point out that people have difficulty recognizing redundancies in the information they use to form overall beliefs. This is of particular importance to both marketing scientists who design marketing decision support systems and marketing researchers who are interested in understanding how consumers make significant decisions over time.

From a management perspective, the real issue is how to prevent escalation behavior in situations such as the introduction of a new product or the start of a new strategic alliance. Our findings should be helpful in determining which approaches have the most promise in preventing escalation behavior. However, because escalation was caused by a biased belief-updating process, our findings highlight the inherent difficulty in trying to eradicate escalation behavior completely. This may provide some explanation for why previous studies examining deescalation strategies produced mixed results (Boulding, Morgan, and Staelin 1997; Keil and Robey 1999; McNamara, Moon, and Bromiley 2002; Simonson and Staw 1992). The findings from those studies, in conjunction with our findings, lead us to suggest the following methods for preventing escalation behavior: (1) Change the organizational structure such that continue/stop decisions are made by someone with no prior beliefs about the project; (2) use stopping rules that are based on objective data; and (3) accept that decision makers hold biased beliefs, and institute policies and procedures that minimize the adverse effects of these biased beliefs. Such a policy could be an educational program that stresses the idea that managers tend to examine data with the goal of making the world appear consistent with their own views of reality (Carlson and Russo 2001). Specifically, it could stress the need for a decision maker to document how he or she might be wrong, thus forcing him or her to attend to negative information.

However, even these suggestions are not foolproof. For example, Bolton (2003) shows that such analytic reasoning is unlikely to eliminate bias from sticky priors if these priors were generated by nonanalytic thinking. Likewise, when McNamara, Moon, and Bromiley (2002) tried to use the approach of changing decision makers, they were able to attenuate but not eliminate escalation in a commercial lending situation. Similarly, Boulding, Morgan, and Staelin (1997) found that stopping rules adopted at the time of the initial decision also attenuated escalation but that some

decision makers simply ignored their own stated stopping rule when making subsequent decisions. This willingness to change the rules is similar to what Keil, Mann, and Rai (2000) found in the domain of information technology projects. They found that the best explanation for escalation in such projects was what they referred to as the “completion effect”—a tendency to substitute the goal of completing the project for the goal of maximizing economic returns. Together, these studies indicate that implementing organizational solutions will not be easy and that managers need to be doubly alert to the possibility of escalation bias. Moreover, research is needed to understand the underlying drivers of biased belief updating to help identify better ways to reduce escalation behavior in managerial settings. For example, research on decision processes at the team or organization level may point to additional factors that might help individuals and, therefore, firms avoid the trap of escalation bias.

Appendix

Manipulation Check 1⁵

In the first manipulation check, we collected data from 59 participants from the same MBA student pool previously described, varying only the high-involvement (29 participants) and low-involvement (30 participants) conditions. Participants completed the same initial decision/evaluation part as in the main study. After reading the part in the case that informed them that the company decided to proceed with the product introduction, they were asked to indicate how well ten different questions (see the “Manipulation Check Questions”) reflected their beliefs. The questions assessed (on a nine-point scale ranging from “low” to “high”) how much respondents felt publicly committed, accountable, and responsible for the product launch.⁶

Each of the ten items is significantly different ($p < .001$ or better) between the low- and the high-involvement groups and in the expected direction. If we sum the items and divide by ten, the mean of the low-involvement group is 5.1 and the mean of the high-involvement group is 7.2. These means are significantly different ($t = 6.84, p < .001$), and these results indicate significantly different involvement in the low- and high-involvement groups.

Manipulation Check 2

A total of 44 participants taken from the same MBA student pool and evenly divided between the high- and the low-involvement conditions completed the same initial decision/evaluation part as in the main study. After finishing this decision/evaluation task, they were asked to fill in a debriefing sheet that instructed them to “describe why you think the Quality Valve Company should or should not launch the new product.”

⁵In both manipulation check studies, we collected data in the positive prior information/negative new information conditions.

⁶Cronbach’s alpha for the ten items is .94. Principal component factor analysis indicates a single factor solution.

Prior research indicates that people who feel accountable for (i.e., psychologically and/or socially involved with) a decision have a need to justify their position. Therefore, they generate attitude-consistent thoughts, which lead to more (less) thoughts that support (oppose) their initial attitude (Huber and Seiser 2001; Lambert et al. 1996; Schlenker 1980; Tetlock 1985; Tetlock, Skitka, and Boettger 1989). Thus, we expect the high-involvement respondents to mention more arguments in support of and fewer arguments opposed to their decision than the low-involvement respondents. In addition, evidence in the literature suggests that people who are accountable for their decisions write more thorough justifications (Koonce, Anderson, and Marchant 1995), are more likely to qualify their opinions, and exhibit more recognition of the trade-offs involved (Lord 1992). Thus, we expect the respondents in the high-involvement condition to report relatively more arguments that make explicit trade-offs between the pluses and minuses of the decision.

Two coders, who were blind to the hypotheses, placed participants’ arguments into the following three categories: (1) arguments that support the initial decision,⁷ (2) arguments that do not support the initial decision, and (3) arguments that make explicit trade-offs between supporting and nonsupporting aspects. The percentage of agreement between the coders was 84%, with a Cohen’s (1960) kappa of .73 and a Perreault and Leigh (1989) reliability index of .87 (95% confidence interval of .81–.93). All measures indicate a reasonable level of agreement between the two coders.

We used the F test for differences in count data (Kanji 1993, p. 51) to test for differences between our manipulated conditions. As we expected, the number of supporting arguments made by the high-involvement group was significantly greater than the number made by the low-involvement group ($F_{67, 110} = 1.62, p < .05$), and the pattern was reversed for nonsupporting arguments ($F_{29, 54} = 1.77, p < .05$).⁸ As we also expected, the high-involvement group produced significantly more explicit trade-off arguments than the low-involvement group ($F_{3, 33} = 8.00, p < .001$).

Manipulation Check 1 Questions

Respondents were asked to assess “How well do the following statements reflect your beliefs” on a nine-point scale ranging from “completely disagree” to “completely agree” for the following ten statements:

1. If the project succeeds, I will be given substantial credit.
2. People will associate the success or failure of the project with me.

⁷The coders actually coded arguments as for or against the launch. For obvious reasons, we did not want the coders to be aware of the respondent’s actual decision. By examining participants’ launch decisions later, we were able to ascertain whether the arguments were supporting or not.

⁸All reported results are based on the average of the two coders’ values. Analysis based on each coder separately yields the same results.

3. I am largely responsible for the decision to launch the product.
4. The outcome of this project will have a major impact on my career.
5. If the project fails, I will be given substantial blame.
6. People will attribute the decision to me.

7. I feel responsible for the chosen course of action.
8. I feel attached to the decision to launch the product.
9. My input into the decision process was public.
10. If the project fails, top management will believe that I gave poor advice.

REFERENCES

- Ajzen, I. and M. Fishbein (1977), "Attitude-Behavior Relations: A Theoretical Analysis and Review of Empirical Research," *Psychological Bulletin*, 84 (5), 233–50.
- Arkes, H.R. and C. Blumer (1985), "The Psychology of Sunk Costs," *Organizational Behavior and Human Decision Processes*, 35 (1), 124–40.
- Bacharach, S.B., P. Bamberger, and B. Mundell (1995), "Strategic and Tactical Logics of Decision Justification: Power and Decision Criteria in Organizations," *Human Relations*, 48 (5), 467–88.
- Barton, S.L., D. Duchon, and K.J. Dunegan (1989), "An Empirical Test of Staw and Ross's Prescriptions for the Management of Escalation of Commitment Behavior in Organizations," *Decision Sciences*, 20 (3), 532–44.
- Bazerman, M. (1984), "The Relevance of Kahneman and Tversky's Concept of Framing to Organizational Behavior," *Journal of Management*, 10 (3), 333–43.
- and M. Neale (1992), "Nonrational Escalation of Commitment in Negotiation," *European Management Journal*, 10 (2), 163–68.
- Bettman, J.R. (1979), *An Information Processing Theory of Consumer Choice*. Reading, MA: Addison-Wesley.
- Bobocel, D.R. and J.P. Meyer (1994), "Escalating Commitment to a Failing Course of Action: Separating the Roles of Choice and Justification," *Journal of Applied Psychology*, 79 (3), 360–33.
- Bolton, Lisa E. (2003), "Stickier Priors: The Effects of Nonanalytic Versus Analytic Thinking in New Product Forecasting," *Journal of Marketing Research*, 40 (February), 65–79.
- Boulding, W., A. Kalra, and R. Staelin (1999), "The Quality Double Whammy," *Marketing Science*, 18 (4), 463–84.
- , Marian Chapman Moore, Richard Staelin, Kim P. Corfman, Peter Reid Dickson, Gavan Fitzsimons, Sunil Gupta, Donald R. Lehmann, Deborah J. Mitchell, Joel E. Urbany, and Barton A. Weitz (1994), "Understanding Managers' Strategic Decision-Making Process," *Marketing Letters*, 5 (4), 413–26.
- , Ruskin Morgan, and Richard Staelin (1997), "Pulling the Plug to Stop the New Product Drain," *Journal of Marketing Research*, 34 (February), 164–76.
- Brockner, J. and J.Z. Rubin (1985), *Entrapment in Escalating Conflicts*. New York: Springer-Verlag.
- , ———, and E. Lang (1981), "Face-Saving and Entrapment," *Journal of Experimental Social Psychology*, 17 (1), 68–79.
- Camerer, C.F. and R.A. Weber (1999), "The Econometrics and Behavioral Economics of Escalation of Commitment: A Re-examination of Staw and Hoang's NBA Data," *Journal of Economic Behavior and Organization*, 39 (1), 59–82.
- Carlson, K.A. and J.E. Russo (2001), "Biased Interpretation of Evidence by Mock Jurors," *Journal of Experimental Psychology—Applied*, 7 (2), 91–103.
- Cohen, J. (1960), "A Coefficient of Agreement for Nominal Scales," *Educational and Psychological Measurement*, 20 (1), 37–46.
- Conlon, D.E. and H. Garland (1993), "The Role of Project Completion Information in Resource Allocation Decisions," *Academy of Management Journal*, 36 (2), 402–413.
- Einhorn, H.J. (1970), "The Use of Nonlinear, Noncompensatory Models in Decision Making," *Psychological Bulletin*, 73 (2), 221–30.
- Festinger, L. (1957), *A Theory of Cognitive Dissonance*. Stanford, CA: Stanford University Press.
- Fox, F.V. and B.M. Staw (1979), "The Trapped Administrator: Effects of Job Insecurity and Policy Resistance upon Commitment to a Course of Action," *Administrative Science Quarterly*, 24 (3), 449–71.
- Hausman, J.A. (1978), "Specification Tests in Econometrics," *Econometrica*, 46 (6), 1251–71.
- 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.
- Hoch, S.J. and Y. Ha (1986), "Consumer Learning: Advertising and the Ambiguity of Product Experience," *Journal of Consumer Research*, 13 (2), 221–33.
- Hsee, C.K. (1995), "Elastic Justification: How Tempting but Task-Irrelevant Factors Influence Decisions," *Organizational Behavior and Human Decision Processes*, 62 (3), 330–37.
- Huber, O. and G. Seiser (2001), "Accounting and Convincing: The Effect of Two Types of Justification on the Decision Process," *Journal of Behavioral Decision Making*, 14 (1), 69–85.
- Johnson, E. and R. Meyer (1984), "Compensatory Models of Non-Compensatory Choice Processes: The Effect of Varying Context," *Journal of Consumer Research*, 11 (1), 528–41.
- Jonas, E., S. Schulz-Hardt, D. Frey, and N. Thelen (2001), "Confirmation Bias in Sequential Information Search After Preliminary Decisions: An Expansion of Dissonance Theoretical Research on Selective Exposure to Information," *Journal of Personality and Social Psychology*, 80 (4), 557–71.
- Kahneman, D. and A. Tversky (1979), "Prospect Theory: An Analysis of Decisions Under Risk," *Econometrica*, 47 (2), 263–91.
- Kanji, G.K. (1993), *100 Statistical Tests*. London: Sage Publications.
- Keeney, R.L. and H. Raiffa (1976), *Decision with Multiple Objectives: Preferences and Value Tradeoffs*. New York: John Wiley & Sons.
- Keil, M. (1995), "Escalation of Commitment in Information Systems Development: A Comparison of Three Theories," *Academy of Management Journal*, 38 (2), 348–52.
- , J. Mann, and A. Rai (2000), "Why Software Projects Escalate: An Empirical Analysis and Test of Four Theoretical Models," *MIS Quarterly*, 24 (4), 631–64.
- and D. Robey (1999), "Turning Around Troubled Software Projects: An Exploratory Study of the Deescalation of Commitment to Failing Courses of Action," *Journal of Management Information Systems*, 15 (1), 63–87.
- Klayman, J. (1995), "Varieties of Confirmation Bias," in *Decision Making from a Cognitive Perspective: The Psychology of Learning and Motivation: Advances in Research and Theory*, Vol. 32, J. Busemeyer, R. Hastie, and D.C. Medin, eds. San Diego: Academic Press.

- Koonce, L., U. Anderson, and G. Marchant (1995), "Justification of Decisions in Auditing," *Journal of Accounting Research*, 33 (Autumn), 369–84.
- Lambert, A.J., S. Cronen, A.L. Chasteen, and B. Lickel (1996), "Private Versus Public Expressions of Racial Prejudice," *Journal of Experimental Social Psychology*, 32 (5), 437–59.
- Lipshitz, R. (1991), "The Road to Desert Storm," *Organization Studies*, 16 (2), 242–64.
- Lord, A.T. (1992), "Pressure: A Methodological Consideration for Behavioral Research in Auditing," *Auditing: A Journal of Practice and Theory*, 11 (2), 90–108.
- Lukas, P. (1998), "The Ghastliest Product Launches," *Fortune*, (March 16), 44.
- McAllister, D.W., T.R. Mitchell, and L.R. Beach (1979), "The Contingency Model for the Selection of Decision Strategies: An Empirical Test of the Effects of Significance, Accountability, and Reversibility," *Organizational Behavior and Human Performance*, 24 (2), 228–244.
- McNamara, G., H. Moon, and P. Bromiley (2002), "Banking on Commitment: Intended and Unintended Consequences of an Organization's Attempt to Attenuate Escalation of Commitment," *Academy of Management Journal*, 45 (2), 443–52.
- Mitroff, I. (1974), *The Subjective Side of Science*. Amsterdam: Elsevier.
- Northcraft, G.B. and G. Wolf (1984), "Dollars, Sense, and Sunk Costs: A Life Cycle Model of Resource Allocation Models," *Academy of Management Review*, 9 (2), 225–34.
- Perreault, William D., Jr., and Laurence E. Leigh (1989), "Reliability of Nominal Data Based on Qualitative Judgments," *Journal of Marketing Research*, 26 (May), 135–48.
- Ross, J. and B.M. Staw (1986), "Expo 86: An Escalation Prototype," *Administrative Science Quarterly*, 31 (2), 274–97.
- Samuelson, W. and R. Zeckhauser (1988), "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, 1 (1), 7–59.
- Schlenker, B.R. (1980), *Impression Management: The Self-Concept, Social Identity, and Interpersonal Relations*. Monterey, CA: Brooks/Cole Publishing.
- Schmidt, J.B. and R.J. Calantone (1998), "Are Really New Product Development Projects Harder to Shut Down?" *Journal of Product Innovation Management*, 15 (2), 111–23.
- Schwenk, C.R. (1986), "Information, Cognitive Biases, and Commitment to a Course of Action," *Academy of Management Review*, 11 (2), 298–310.
- Silver, W.S. and T.R. Mitchell (1990), "The Status Quo Tendency in Decision Making," *Organizational Dynamics*, 18 (4), 34–46.
- Simonson, I. and B.M. Staw (1992), "Deescalation Strategies: A Comparison of Techniques for Reducing Commitment to Losing Courses of Action," *Journal of Applied Psychology*, 77 (4), 419–26.
- 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 (1), 27–44.
- and H. Hoang (1995), "Sunk Costs in the NBA: Why Draft Order Affects Playing Time and Survival in Professional Basketball," *Administrative Science Quarterly*, 40 (3), 474–94.
- and J. Ross (1987), "Behavior in Escalation Situations: Antecedents, Prototypes, and Solutions," in *Research in Organizational Behavior*, Vol. 7, B.M. Staw and L.L. Cummings, eds. Greenwich, CT: JAI Press.
- Tetlock, P.E. (1985), "Accountability: The Neglected Social Context of Judgment and Choice," *Research in Organizational Behavior*, 7, 297–332.
- , L. Skitka, and R. Boettger (1989), "Social and Cognitive Strategies for Coping with Accountability: Conformity, Complexity, and Bolstering," *Journal of Personality and Social Psychology*, 57 (4), 632–40.
- Whyte, G. (1986), "Escalating Commitment to a Course of Action: A Reinterpretation," *Academy of Management Review*, 11 (2), 311–21.
- Wicklund, R.A. and J.W. Brehm (1976), *Perspectives on Cognitive Dissonance*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Copyright of *Journal of Marketing* is the property of American Marketing Association and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.