Open access • Journal Article • DOI:10.1509/JMKG.70.2.108

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

Eyal Biyalogorsky, William Boulding, Richard Staelin

Institutions: University of California, Davis, Duke University

Published on: 01 Apr 2006 - Journal of Marketing (American Marketing Association)

Topics: Escalation of commitment and Confirmation bias

STUCK IN THE PAST:
WHY MANAGERS PERSIST WITH NEW PRODUCT FAILURES

April 24, 2000

Eyal Biyalagorsky, William Boulding, and Richard Staelin*

*Assistant Professor, UC Davis, and Professor and Edward and Rose Donnell Professor, respectively, at the Fuqua School of Business, Duke University. The authors thank Jim Bettman, Chip Heath, John Lynch, John Payne, Wes Hutchinson, Darryl Banks, and seminar participants at the 1996 Marketing Science Conference, the 1997 Haring Symposium, the 1999 MSI Trustee's Meeting, the Wharton School, and Duke University for their helpful comments. We also thank Kevin Keller and David Bell for their help in obtaining subjects.

STUCK IN THE PAST:
WHY MANAGERS PERSIST WITH NEW PRODUCT FAILURES

ABSTRACT

In this research we examine the phenomenon of escalation bias in the context of managing new product introductions. In particular, we propose a formal descriptive model that captures the potential effects of confirmatory bias. We do this by having the initial belief structure (i.e., the set of beliefs leading to the first decision) non-normatively affecting the manager's perceptions of new information along with its normative role in the updating of beliefs. This over-weighting of the initial (positive) belief structure can lead the manager to continue (improperly) with a (losing) course of action.

We investigate this new model, and others, via a controlled experiment where participants are asked to act as managers involved with a new product introduction and re-evaluation decision. Specifically, we manipulate prior belief structures, whether or not the person made the initial decision to launch the new product (i.e., was personally and publicly committed to the initial decision), and the nature of the new information about the success of the new product launch.

Our results show that making the initial decision is not a necessary condition to induce commitment to a losing course of action, i.e., escalation bias. Even if a person is not responsible for the initial decision, and thus is not "committed" to this decision, this person can exhibit escalation bias. In particular, a person's initial positive beliefs about a situation non-normatively affect the subsequent data interpretation and integration process, which in turn alters the updated belief structure of the decision maker and the subsequent decision to stay with a losing course of action. Thus, we believe attention should shift from a narrow focus on the commitment construct to a broader perspective that accounts for the role of the initial belief structure in explaining subsequent data acquisition, interpretation, and choice processes. By providing an understanding of how incorrect beliefs are formed and cause escalation, our formal descriptive model provides the groundwork necessary for designing systems to help managers avoid the trap of escalation bias.

This paper explores why managers tend to stay committed to a course of action despite strong negative feedback with respect to the advisability of this action. This tendency, often referred to as escalation of commitment, has received extensive attention in the literature. Just a few examples of this tendency to stay the course of action are as follows: the Vietnam War and Desert Storm (Lipshitz 1991), case studies of Expo 86 (Ross and Staw 1986), the Apollo moon missions (Mitroff 1974), the Campeau-Federated merger, the coffee wars between Philip Morris and P&G (Bazerman and Neale 1992), overplaying high paid NBA players (Staw and Ha, 1995), and a plethora of experimental investigations (see for example Arkes and Blumer 1985; Boulding, Morgan, and Staelin 1997; Brockner and Rubin 1985; Schmidt and Calantone 1998; and Staw 1976). The implications of this tendency for an organization are potentially substantial and include areas such as the management of R & D, choice of supplier relationships, choice of compensation schemes, and making bank loans. In this paper, we will focus on the context of new product introductions to illustrate and understand this phenomenon.

Two streams of literature purport to explain the escalation phenomenon. The first of these is based on the concept of self-justification, i.e., a person stays committed to justify a past decision (e.g., Staw 1976; Brockner and Rubin 1985, etc.). The second is based on either the psychology of sunk costs (e.g., Arkes and Blumer 1985, etc.) or framing of values given an S-shaped utility function for losses and gains (e.g., Kahneman and Tversky 1979; Bazerman 1984; and Whyte 1986). Based on these literatures, escalation of commitment was considered a well-documented and well-understood phenomenon. However, a provocative paper by Heath (1995) challenges the premise that people have shown the existence of escalation bias.[1] Instead, he argues that existing evidence of escalation bias is suspect for one of two reasons. First, some studies misinterpret what constitutes bias, since in these studies continued commitment to a course of action was actually the rational (normative) course of action. Second, in the remaining studies it was unclear that staying the course of action was non-normative since the expected future returns relative to an alternative course of action were not clearly indicated. In this latter instance Heath argues that since information was unavailable to made a normative decision, individual's unobserved beliefs about marginal future returns may have been positive enough to

[1] Armstrong (1996) cites other evidence challenging the generalizability of escalation bias.

support continued commitment to a course of action. For example, coaches may play higher paid NBA basketball players more, all else equal, because they believe that this action will yield higher long term returns because these players have more long run potential than some other (lower paid) player.

One implication of Heath's arguments is that to positively establish the existence of escalation bias it is not only necessary to show that managers choose to remain with the original course of action, but also that there is a financially normative alternative clearly available to the decision maker. A second implication is that beliefs are an important determinant of continued commitment to a course of action. Although Heath points to the importance of these beliefs, he is agnostic as to how these beliefs are formed or whether they themselves are normative. It is these two implications that guide the research herein.

Specifically, the overarching purpose of this paper is to gain a better understanding of the mechanisms that underlie escalation bias. In order to do this, we must first establish that escalation bias exists. This leads us to investigate escalation bias in an experimental setting for new product introduction decisions where we provide clear financial information that indicates the normative business model choice is to discontinue the new product introduction in favor of the clearly superior alternative of cashing out the investment.

The second implication from Heath leads us to center our attention on the role of beliefs and how they affect the decision to continue commitment to a course of action. We start with suggestive evidence provided by Boulding et al. (1997) that beliefs underlying the decision to stay the course of action are biased. If these beliefs are biased in the direction associated with the initial chosen action, then managers may remain inappropriately committed to a course of action, i.e., exhibit escalation bias. Since previous escalation explanations have not relied on beliefs but instead concentrated on the effect of commitment to an initial course of action, we compare and contrast these two different explanations as they relate to escalation bias.

In focusing on the role of beliefs, we note the surprising independence of two extensive literatures in organizational behavior and psychology—escalation bias and confirmatory bias. If, as conjectured in Heath and Boulding et al., underlying beliefs drive continued commitment to a course of action, it seems reasonable to explore whether the underlying process of belief updating

is itself normative. Specifically, is the underlying belief updating subject to confirmatory bias? If so, very positive initial beliefs about an investment could explain continued commitment in the face of negative new information that normatively should result in discontinuing a course of action. We note that such a result would perhaps be unsurprising when viewed through the lens of confirmatory bias. What we find surprising is that such a lens has not been applied to the phenomenon of escalation of commitment. Rather, as noted, other theoretical explanations dominate this literature.

In addition to providing this new lens to examine escalation bias, we provide a more precise conceptual model of confirmatory bias. The idea of confirmatory bias is certainly not new. What we believe is new is that we suggest and demonstrate that *confirmatory processing takes place concurrently at multiple levels in the updating process*. In particular, we propose and empirically test a three-part structure of how confirmatory bias affects a person's updated beliefs. First, we show that a person's prior overall belief about the viability of a project directly biases the person's perception of the updated levels of the factors (e.g., market share, competitive response, etc.) felt to influence the ultimate success or failure of the project. We label this effect "factor level bias." Second, we show this prior overall belief indirectly biases the weights given to each factor as it determines the project's final outcome. We label this effect "factor weight bias." Third, we find the person's prior overall belief also directly biases the person's updated overall belief about the viability of the project above and beyond the biases associated with the biased factor levels and factor weights. We label this effect "initial belief stickiness."

In sum, this research examines the underlying process that leads to escalation of commitment. By examining this underlying process we are able to disentangle the role of belief updating from traditional explanations of escalation bias. Moreover, by modeling this underlying process we are able to provide greater insights into the specific mechanisms of how confirmatory bias works. Finally, in studying this process we emphasize the importance of a normative benchmark when one studies belief updating and decisions that exemplify escalation of commitment.

MECHANISMS UNDERLYING ESCALATION OF COMMITMENT

A typical escalation setting starts when a manager decides to take an initial course of action after forming some overall beliefs about an idea under consideration. Once the decision is made in accordance with these beliefs, the manager becomes "committed" to the course of action. Subsequently, the manager is exposed to additional, negative, information. If the manager decides to stay with what is now revealed to be a losing course of action, the manager is classified as exhibiting an escalation bias.

In our research we examine this sequence of events and identify which of, and to what degree, each of these events (i.e., the initial beliefs, the initial decision, and the new information) drive escalation bias. Excellent reviews and syntheses of the escalation bias literature can be found in Staw and Ross (1987) and Brockner (1992). Very briefly, prior explanations of escalation bias typically center on how an *initial decision* (and not the belief structure) affects the decision rules used in the second (stay the course) decision. One explanation posits that managers maintain commitment to an initial decision to preserve their self, or public, image (Rubin and Brockner 1975; Staw 1981) A second explanation attributes escalation to framing effects. Specifically, a manager frames the second decision leading to escalation relative to the losing outcome of the initial decision (which is a sunk cost loss), rather than from a neutral starting position. Thus, the bias here results from a valuation process that deviates from the normative (i.e., marginal benefits versus marginal costs) decision making process.

More recently, a third explanation proposes that commitment to the initial decision does not change the decision criteria for the second decision, but instead causes selective attention to, and manipulation of, new information, which in turn leads to escalation (Boulding et al. 1997). Although commitment to the decision, per se, remains the driving force in escalation, this proposed mechanism leading to escalation is somewhat different in that it centers on how commitment affects a person's updated beliefs.[2]

We propose a mechanism that differs from these prior explanations. In particular, we

[2] A fourth explanation, due to Samuelson and Zeckhauser (1988), has less to do with one's actual commitment to the first choice than the fact that the status quo created by an initial (i.e., anyone's) commitment is highly salient. Samuelson and Zeckhauser support this conjecture with a series of experiments that show individuals disproportionately choose to stay with the status quo. We control for this possible explanation in our study.

believe escalation bias results from a confirmatory bias that affects the belief updating process. This is in contrast to commitment to an initial decision causing people to either justify their past decision or improperly frame their new decision. We start by hypothesizing that managers overweigh (relative to a normative Bayesian updating model) their prior (positive) beliefs when they acquire new information about factors that should affect their new overall belief about the continued viability of the current course of action. This overweighing is due to three different processes. First, individuals perceive the level of the new information more positively compared to a normative bystander. Second, we postulate that managers overweigh (relative to a normative financial model based on these same factors) factors found to be more positive than expected and under-weigh factors found to be less positive than expected when integrating the new information. Third, we posit that managers, even after adjusting for the fact that they non-normatively perceive and weight new information, also non-normatively consider their prior overall belief when forming their new updated belief. In combination, these three behaviors imply that the manager's initial belief structure (i.e., the set of beliefs leading to the first decision) non-normatively affects the manager's subsequent information search and integration process and thus the final decision to continue (improperly) with a (losing) course of action. Importantly, this occurs independently of the act of making the initial commitment decision.

Since our interest is, in part, on how managers non-normatively incorporate new information into their initial belief structures we need to specify a normative benchmark. This leads us to start with a standard normative Bayesian updating model. We then augment this model to reflect the fact that descriptively individuals often behave non-normatively. Also, since we are dealing with a business decision, our development includes a normative financial model that determines the weights that a manager should place on each factor assuming the financial model accurately captures the situation. Finally, we acknowledge that managers often bring in private information not captured by the financial model. Therefore, we control for the possible role of private information in our subsequent empirical analysis, allowing us to test various theoretical accounts of escalation bias without fear of contamination due to this issue.

We investigate our augmented descriptive model, and others, via a set of controlled experiments in the setting of a new product launch. We choose this setting because it involves

sequential decisions, each based on new, and perhaps negative information. It is also the case that, anecdotally, stopping a new product venture is not easy. For example, the director of the New Products Showcase and Learning Center states that "It sometimes takes more courage to kill a product that's going nowhere than to sustain it" (Lukas 1998). Compatible with this premise, the CEO of Eastman Kodak considered a new product team a "success" when they appropriately decided to stop investing in a new product idea and labeled the project "the world's fastest failure."[3] Thus, decisions about new products are important to the organization and are potentially susceptible to escalation bias.

In our experimental setting we manipulate prior belief structures, commitment to the initial decision, and the new information about the performance of the new product launch. This allows us to estimate the impact of these three factors on the subsequent information search and integration process and the second decision to stay with, or terminate, a (failing) product launch. Also, we control for potential status quo bias by holding fixed the status quo position across all experimental conditions. Finally, we use a within/between subjects design that allows us to hold fixed any impact of private information on the part of the subjects.

We next present a model of the managerial decision-making process that develops the role of prior beliefs and new information. We then describe our experiment and the subsequent data analysis. We conclude by discussing the implications of our findings.

BELIEF UPDATING IN ESCALATION SITUATIONS

Consider the following scenario. A manager acquires information on a number of factors concerning a new product concept and its market potential. Based on this information and the manager's prior knowledge of how these factors influence product success, the manager (or some analyst) develops a business model that captures how the key factors influence the financial viability of the project. The manager then combines this financial model with any other information the manager believes is relevant to form some overall summary belief (evaluation). Based on this (positive) summary belief, the manager decides to launch the product. Subsequently, the manager updates his/her knowledge of the situation by monitoring the

[4]From a speech given to participants in a senior executive management program.

(objectively negative) performance of the product based on new information concerning many of the same factors used to form the initial overall evaluation. When asked if the firm should continue to market the product, or whether resources supporting the product should be transferred to another (certain) venture, the manager forms a new summary belief about the advisability of continuing the project. Based on this new belief, the manager decides to stay with the original course of action.

This scenario illustrates the potential influence of beliefs in making the two different decisions, i.e., the initial launch decision and the decision to stay the course after receiving negative feedback. The questions then become:

1. Do these beliefs influence the two decisions?
2. If so, what factors influence these beliefs, especially the belief held just prior to the second decision? Is this second decision belief influenced by the initial belief, the new data, and/or commitment to the initial decision?
3. If the initial belief and the new information affect the second decision belief, is the updating process normative, or is there another process that better describes how managers update their belief structure?
4. What is the role of commitment? Does it affect beliefs, the process by which beliefs are updated, or does it affect the decision directly after adjusting for the updating process?

To further motivate these research questions and to provide some data on the relative impact of commitment and initial beliefs we briefly present some empirical results from an experiment we more fully describe later in this paper. In this experiment we provide managers with some initial information about the forecasted success of a particular product launch, measure their initial beliefs about specific aspects of this potential launch, and ask them if they want to launch the product. If they say yes, we then provide them new (negative) information about the actual current status of the product after launch and a new forecast of future profitability. We then ask them a) if they want to stay the course or terminate the project and obtain a for-certain dollar value from "cashing out" (where the for-certain amount exceeds the expected return associated with continuing the project) and b) measure their beliefs. In this basic

setting we manipulate three factors: the manager's initial beliefs (positive or very positive); the nature of the new information (negative or very negative); and the degree to which the person is publicly responsible for and committed to the initial launch decision (high or low commitment to the initial decision).

In analyzing the data associated with the second decision we found we could correctly predict this decision 77% of the time simply by using the person's stated summary beliefs associated with the second decision. This 77% should be compared to the naïve modal response prediction percentage of 52%, since 52% of the individuals in our sample decided to stick with the project. From these results we infer that individuals' updated beliefs strongly influence the second decision. Perhaps this is unsurprising. However, more surprisingly, our commitment manipulation had no influence on either the person's initial belief structure or the final belief structure.[4] Instead, the only factors that influenced the person's updated beliefs were the person's initial beliefs and the new information. Simply put, our data indicate that beliefs strongly influence the second decision, but commitment does not influence these beliefs.

Although provocative, these results leave a number of unanswered questions. For example, what is the process by which a manager updates his/her beliefs? Is it normative or does commitment or some other bias, e.g., confirmatory bias, affect the updating process? Does commitment enter directly into the decision after one adjusts for the updated beliefs, or, is the effect of commitment insignificant after adjusting for beliefs? To answer these questions we develop two sets of models that specify how managers form beliefs and then update these beliefs after receiving new information. The first set of models is rational (normative) in that it reflects an underlying financial model associated with the business decision and a Bayesian updating model that reflects how new information should be processed. The second set is descriptive in that it acknowledges managers may follow a non-normative decision making process. We then augment these two sets of models to allow commitment to affect the updating process and/or directly affect the second decision.

In developing these models, we start with the conceptualization that when faced with a

[4] As we later detail, the manipulation check presents strong evidence that the commitment manipulation was successful.

decision concerning a new business venture, managers form mental models that include two different underlying beliefs and one summary belief (Boulding et al. 1994). The first underlying belief is about the specific *factor levels* that influence the performance of the business venture. The second underlying belief is about the *factor weights*, i.e., how and to what degree each of the factor level beliefs affects the overall summary belief. For example, for our new product decision, the factor level beliefs might include beliefs about what will be the resulting market share, price, overall market growth, competitors' responses, and the cost of capital. The factor weight beliefs are used by the manager to build a model of how factor levels combine into some overall profitability measure (e.g., probability of success, expected profitability, etc.). This profitability measure could be augmented by other private information. This augmented summary belief (sometimes referred to the person's evaluation of the project) is used to make the decision.

Normative Model of Belief Formation and Updating

We start the formal presentation by specifying how a manager forms his/her summary belief for the proposed project. We next discuss how the underlying normative components of this summary belief, i.e., factor levels and factor weights, are normatively updated when the manager is exposed to new information. These updated values are then used in forming the updated normative summary evaluation, i.e., the evaluation used to make the second decision of whether to stick with, or cancel, the project. For simplification in the ensuing presentation, we specify general functional relationships. We refer the reader to Appendix A for more explicit mathematical development.

We begin by noting that for a typical business decision there is a normative financial model that maps factor levels and weights into a measure of financial performance such as net present value (NPV). Often, this model is reflected in a spread sheet that has inputs like market share, market growth, prices, costs, etc., and outputs of forecasted profits or losses by year. We represent this normative business model as follows:

(1) $NPV_t = f_1(\text{Factor Level}_{kt}, \text{Factor Weight}_{kt})$,

where the "k" subscript refers to the kth factor and the "t" subscript refers to the time, or period, t when the information is available.

Note that equation 1 is normative in that it reflects an underlying business model. However, it does not necessarily describe how the manager forms a summary belief (i.e., project evaluation), or even represent the "correct" summary measure if the manager has "private" information that he/she believes should be factored into the decision. We capture this as follows. Let PI_i represent private, person specific information, where the i subscript is for manager i, and the manager's summary belief about the project's performance be SB_{it}. Then assume the following relationship:

(2) $SB_{it} = f_2(NPV_t, PI_i)$,

i.e., a manager's summary belief is composed of a normative business model and some private information that the manager believes should influence his/her overall evaluation. It is this summary belief that the manager uses to make his/her decision about the project at time t. Note that while private information often plays a key role in a business decision, it is typically not observable to the researcher. Since this is the case herein, in our analysis we control for this person specific fixed effect during estimation. We discuss this issue in more detail in Appendix A and in the section on estimation.

Equations 1 and 2 make it clear that current factor levels and weights play a major role in determining the person's overall assessment of the viability of a project. Consequently, we next develop models of how a manager *should* update these types of knowledge after receiving new information. We then augment these normative updating models to provide descriptive models of how managers actually update these two types of knowledge.

We start with updating of factor levels. We assume the manager's uncertainty about the true level of any given market factor can be represented by a normal distribution with mean equal to the true mean. Moreover, we assume that the manager's belief about the reliability of any new information about a market factor can be captured by a normal distribution. It then follows from

standard Bayesian analysis (e.g., see Pratt et al. 1995) that the manager's best guess (i.e., the mean) of the updated factor level is a linear combination of the manager's initial best guess (mean) of the market factor level and the new objective information. The weights in this linear combination reflect the relative reliabilities of the initial beliefs and the new information. Specifically,

(3) Factor $\text{Level}_{kt+1} = f_3(\text{Factor Level}_{kt}, \text{Objective New Information}_{kt+1})$,

where Factor Level_{t+1} represents the updated belief and Factor Level_t represents the prior belief.

In modeling factor weight updating we again start with the assumption that managers update their knowledge whenever they obtain new information about the underlying structure of the business model as it maps to the market. Again assuming normality, a normative model has the manager's best guess (i.e., the mean) of the updated factor weight be a linear combination of the manager's initial best guess (mean) of his/her factor weight belief and the new information about factor weights. However, in our empirical setting, as in many real world settings, there is no new information suggesting that the factor weights used in the normative business model should be updated. Without any new information, the Bayesian updating model predicts no updating. Thus, *in our setting* we have the following special case of normative updating of weights,

(4) Factor $\text{Weight}_{kt+1} = \text{Factor Weight}_{kt}$.

Confirmatory Bias Updating Models

Equations 1-4 capture our normative updating process (assuming one controls for the role of private information in equation 2, which may or may not be normative). However, much research calls into question the descriptive validity of these normative updating models. For example, research shows that in numerous settings people exhibit a confirmatory bias. People look for, attend to, and weight new information depending on their prior belief structure (e.g., Boulding et al. 1999; Einhorn and Hogarth 1978; Mahoney and Demonbreun 1978; Mynatt et al.

1977, 1978; Russo et al. 1996, 1998; see Klayman 1995 for a review of this extensive literature). To reflect these stylized facts, we next augment our normative updating models by developing descriptive models of the updating process for factor levels, factor weights, and summary beliefs. In doing so we rely on the observation that people tend to perceive new data to confirm their currently held beliefs.

Updating Factor Level Beliefs

Our descriptive updating model for factor levels still has the manager blending his/her prior beliefs with the new information. However, we now allow the manager's *perception* of the objective information to differ from the objective information in a systematic fashion. Our rational for such an assumption is based on the considerable empirical evidence that people tend to perceive new information to be compatible with their prior beliefs, i.e., they exhibit a confirmatory bias. Specifically, in recent work Russo et al. (1998) show that new attribute evidence is distorted in the favor of the initially preferred brand. Similarly, Hsee (1995) provides evidence of what he labels "elasticity" in shifting perceived factor levels in the direction of a subject's more favored option.

In our setting, we believe two managers with different prior opinions about the likelihood of a project's success will *perceive* specific objective factor level information differently. Specifically, analogous to a model proposed and tested by Boulding et al. (1999), we assume a confirmatory process wherein a person's perception of the objective new information is a combination of the new information and the person's prior summary belief, i.e.,

(5) Biased Perception of New Information$_{ikt+1}$ = f_5 (Objective New Information$_{kt+1}$, SB_{it}),

where manager i's summary belief, SB_{it}, is hypothesized to have a positive influence. Thus, an individual's perception of new information is pushed in the direction of, i.e., confirms, his/her prior summary belief.

Next, our descriptive model assumes an individual uses this biased perception in updating his/her beliefs about factor levels. In particular, we propose that the individual's updated belief

about a factor level is a blend of the individual's prior factor level and the person's biased perceptions of the new information about the factor level. This yields our descriptive model of how a manager updates a factor level,

(6a) Biased Factor Level$_{ikt+1}$ = f_6(Factor Level$_{kt}$, Biased Perception of New Information$_{ikt+1}$),

or, using equation 5,

(6b) Biased Factor Level$_{ikt+1}$ =f_6(Factor Level$_{kt}$, Objective New Information$_k$, SB$_{it}$).

Comparing the descriptive model predicated on confirmatory bias, i.e., equation 6b, to the normative equation 3 provides a clear test of whether or not managers exhibit a confirmatory bias when updating their factor level beliefs. Specifically, these two equations differ by the SB$_{it}$ term in equation 6b. If this term is significant, managers display a confirmatory bias, and consequently the updated factor levels are biased. Thus, we offer the following hypothesis:

H1 (Factor Level Bias): *Updating of factor levels is biased, i.e., the effect of SB_{it} in equation 6b will be positive and significant.*

Updating Factor Weight Beliefs

Our normative updating model for factor weights assumes that a manager should not alter this belief *unless there is new information pertaining to how the market factors should be weighted*. Given our particular empirical setting, there is no such information and thus, normatively, the weights should not change. However, prior research suggests that individuals attend to more and/or weight more heavily data that are most compatible with (i.e., confirm) their initial belief structure (e.g., Hoch and Ha 1986, Boulding et al. 1999). In a similar vein, Hsee (1995) infers the existence of "elastic" factor weights that shift in the direction that supports choice of a subject's initially favored option.

These findings lead us to posit a descriptive model that has the weights changing over

time in a predictable fashion in situations where there is no new information on how the factor levels should be weighted. More specifically, managers will give more weight to market factors that confirm their prior (positive) summary beliefs and less weight to those that are not supportive of these summary prior beliefs. We model this by allowing the importance weight for those factors perceived to be more positive than last time to increase and those factors perceived to be more negative than last time to decrease. For example, say the new information leads the manager to believe the obtained market share is higher than s/he originally expected. Since higher market share is positively related to a positive evaluation, the manager will update his or her beliefs by increasing the weight of market share. More formally, we posit the following process:

(7) Biased Factor $\text{Weight}_{ikt+1} = f_7(\text{Factor Weight}_{kt}, [\text{Biased Factor Level}_{ikt+1} - \text{Factor Level}_{kt}])$.

Note that the normative model for updating of factor weights, i.e., equation 4, is a subset of equation 7. In the normative model, the change in factor level does not enter into the updating process. Therefore, the nested nature of the normative and descriptive models allows us to discriminate between the two models. Specifically, our belief that the factor weights will change depending on the manager's perceptions of the new factor level relative to his or her prior beliefs about the factor levels leads to the following hypothesis:

H2 (Factor Weight Bias): *Updating of factor weights is biased. The factor weights will increase with increases in the factor levels, decrease with decreases in factor levels, and remain constant with no changes in factor levels, i.e., the effect of the difference in factor levels term in equation 7 should be positive and significant.*

Updating Summary Beliefs

Our descriptive model of the summary belief process augments our position that a person's summary evaluation is a function of the financial model and private information by again incorporating a confirmatory bias. Specifically, similar to the notion of anchoring and adjustment articulated by Hogarth and Einhorn (1992), we propose that managers may inappropriately over-use the prior evaluation as a reference point when forming their new overall

belief. We posit this because we believe managers may not have a good grasp on the information dependencies between the various beliefs and thus fail to recognize that the effect of their prior beliefs has already been accounted for in updating the market factor level and weight beliefs. Moreover, we propose that these managers do not use the normative financial model but their own business (mental) model which is based upon their perceptions of market factor levels and weights. In addition, they incorporate any private information they might have to arrive at their summary belief evaluation.

Note that this descriptive model has *three* sources of bias. The first two sources come from the biases associated with the updated factor levels and weights, i.e., equations 6 and 7. This yields a biased perception of NPV_{t+1}. On top of this bias, we propose an incremental biasing effect due to stickiness in the initial prior. Capturing these effects leads to the following descriptive model:

(8) Biased $NPV_{it+1} = f_8(\text{Biased Factor Weights}_{ikt+1}, \text{Biased Factor Levels}_{ikt+1})$

where Biased NPV_{it+1} is the managers perception of the new normative business model, and

(9) $SB_{it+1} = f_9(\text{Biased } NPV_{it+1}, PI_i, SB_{it})$.

We next decompose the biased NPV perception into the true NPV value and the biased component caused by the biased factor weights and factor levels, i.e.,

(10) Biased $NPV_{it+1} = NPV_{t+1} + \text{Business Model Bias}_{it+1}$.

Substituting equation 10 into equation 9 yields the following relationship:

(11) $SB_{it+1} = f_{11}(NPV_{t+1}, \text{Business Model Bias}_{it+1}, PI_i, SB_{it})$.

Note that equation 11 is now directly comparable to its normative counterpart, i.e., equation

2. Specifically, equation 11 differs from equation 2 by the Business Model $Bias_{it+1}$ and SB_{it} terms. Thus, estimation of equation 11 provides a direct mechanism for assessing the degree to which, and how, managers behave non-normatively when updating their summary beliefs about a project.

More formally, after controlling for NPV_{t+1} and PI_i, we can reject the normative benchmark in favor of our descriptive model if we find support for the following hypotheses:

H3 (Initial Belief Stickiness): *a) The effect of the prior summary evaluation on the current summary evaluations will be positive and significant. b) The effect of business model bias (found by removing NPV_{it+1} from of Biased NPV_{it+1}) will be positive and significant.*

Note that H3 allows us to quantify if, and the degree to which, managers are non-normative in their decision-making. This approach differs from prior work that compares decision-makers' subjectively determined weights to weights based on the relative reliabilities of the priors and the new information. We did not take this approach since (as researchers) we did not know managers' assessments of these two reliabilities and thus the objective weights the managers should have placed on the their prior beliefs and the new information. Still we are able to compare managers' observed behavior with normative behavior by explicitly modeling how these subjective weights might be a function of factors other than the relative strength of the prior and the reliability of the new information. Moreover, since the normative model is a subset of our descriptive model, we are able to test the veracity of both models.

Implications of Biased Updating for the Second Decision

Our descriptive model predicts that if a person has positive initial beliefs, these prior beliefs, independent of commitment to an initial course of action, will have more influence on the new decision than they should. Since a person with a positive prior will have updated beliefs that are positively biased, this person will exhibit a tendency to inappropriately stick with the current course of action in the face of negative new information, i.e., demonstrate escalation bias. Consequently, instead of emphasizing the importance of an initial *decision*, we emphasize the possibility that managers are unduly influenced by an initial *belief*.

We acknowledge that the undue influence of a positive initial belief is often empirically

confounded with an initial decision, i.e., commitment to a course of action. However, our model predicts it can also arise with no commitment to an initial course of action. This leads to the following hypothesis:

H4 (Initial Belief Bias)*: Independent of any initial commitment to a course of action, an initial (positive) summary belief overly influences a person's updated summary belief. It is this non-normative (biased) summary belief that influences the subsequent decision.*

Thus, H4 represents a summary of our descriptive model. We present this full model graphically in Figure 1.

Insert Figure 1 about here

In the introduction to this paper we presented data showing that commitment to an initial decision does not directly influence initial or updated summary beliefs. Thus, if commitment is to have any effect on the second decision it must do this in one of the following two ways. The first way, suggested by Boulding et al. 1997, has commitment to an initial decision indirectly affecting the second decision by influencing the *process* by which beliefs are updated. We test this indirect path by allowing the effects of the variables included in equations 6b, 7, and 11 to vary according to the degree of commitment to the initial decision. This conjecture of an "indirect effects model" due to initial commitment is captured by the following hypothesis:

H5 (Initial Commitment Process Bias): *Commitment to an initial course of action indirectly influences the likelihood that a manager will stick with the initial course of action by affecting the belief updating process as specified in equations 6b, 7 and 11. Ceteris paribus, the stronger the initial commitment to a course of action, the higher the impact of the initial summary evaluation on the updated market factor evaluations, the updated weights, and the posterior summary evaluation.*

A second alternative model has the degree of initial commitment directly affecting the subsequent d*ecision*, above and beyond the effect due to summary beliefs. One can test this alternative model by seeing if initial commitment to a course of action has a significant main effect on the continuation decision after controlling for summary beliefs. Specifically, this "direct effects" model is captured by the following hypothesis:

H6 (Initial Commitment Bias): *Commitment to an initial course of action has a direct*

effect on managers' decisions; ceteris paribus, the stronger the initial commitment to a course of action, the more likely the manager is to stick with the initial course of action.

In sum, evidence in favor of hypotheses 1-4 indicates support for our proposed descriptive model, i.e., that escalation bias is due to bias in the updating of beliefs. In contrast, evidence in favor of hypotheses 5 and 6 indicates support for more traditional views that escalation bias is due to an initially chosen course of action.

STUDY

To test our descriptive model of how managers form and update beliefs and assess the degree to which beliefs and initial commitment to a course of action affect the continuation decision, we conducted an experiment in which participants faced a typical escalation situation. The context was that of a new product introduction and evaluation task. Participants were asked to read information (a "case") about a new product introduction in which they took on the role of the manager responsible for the new product. This information, which included the results from a financial model, was generally positive, and indicated that the new product developed by their R & D department would, in all likelihood, be a success on the market, i.e., generate a positive net present value (NPV). Based on this information, participants were asked to make a launch/no-launch decision. Next, those participants who decided to launch the product received information about the product's actual performance during its first two years on the market. Included in this information were projections from a financial model about future performance of the product. They were also told they could take an alternative course of action, namely, terminate the project by "cashing out" for a certain dollar amount. More specifically, this new information indicated that the product was not performing as well as expected (in some cases the actual market share was less than half that expected) and that the expected value of the updated NPV forecast was less than the for-certain dollar value of terminating the project. Thus, the normative decision based solely on the financial information available in the case was to discontinue the project. Participants at this stage were then asked to decide to either continue or discontinue offering the product on the market. In this way we emulated a typical escalation situation. Participants in the first stage commit to a course of action and/or develop positive

beliefs about the new venture, only to receive negative feedback in the second stage. They then must decide to either continue with the original course of action or to stop.

We use this particular setting because in a previously published study (Boulding et al. 1997) employing the same methodology and only slight differences in the given information, senior level managers exhibited a marked tendency to continue with the product. Specifically, between 45% to 100% of managers in different conditions continued with the product, though the normative decision based on the new financial information was to discontinue the product.[5] Consequently, we were confident that participants in our study would exhibit the same tendency toward escalation, thereby allowing us to test the hypotheses about the underlying process of escalation.

Subjects

We obtained usable responses from 142 participants that completed the experiment. Approximately two thirds of these subjects were MBA's from four different full time programs, two on the East Coast and two on the West Coast. The remaining subjects were mid-level managers enrolled in a weekend MBA program at one of the East Coast schools. All participants completed at least one full semester of studies, which included instruction in decision making. Since our sample uses MBA students it is logical to ask if this sample can generalize to the population of senior level managers who normally make major investment decisions. To address this issue, we replicated the control condition used in the Boulding et al. (1997) study (which utilized senior level managers as subjects), using a sample of our subjects. In our study 60% of the participants (12 of 20) in this control condition discontinued the product versus 66% (27 of 41) that discontinued in the Boulding et al. study.[6] The difference in stopping rates between the two studies is not significant ($t=.45$, $p>.30$). This one point of comparison suggests that no substantial differences exist in the way participants in the two studies responded to the provided

[5] For comparison, Boulding et al. reported that no one in the group of academic experts decided to continue with the project when asked to make this decision based on the financial information made available.

[6] Although this control cell replicates the Boulding et al. study, the levels of provided information differ slightly from those used in the current study. Consequently, these data cannot be used for further comparisons with the other data collected in the research herein.

information and increases our confidence that we can generalize our findings to the population of senior managers represented in the Boulding et al. study.

In addition to the subjects described above, we utilized 59 subjects from the same MBA subject population for purposes of a check on our manipulation of commitment. We report these results below after providing a more detailed description of the stimuli and our measurements.

The Case

Participants in our study were asked to complete the Quality Valve Company Case. The first two pages of the case described a new product launch and the associated investment decision for an improved valve targeted at the large truck emissions control market. The case provided information on the history of the market, the company, and the primary competitor. The initial decision problem was whether to invest $2.5M in new machinery and launch the improved valve.

As part of the case description participants learned that financial success of the new product venture depended on three key uncertainties. These were the forecasted share of market the new product garners after launch, the industry growth rate, and the competitive response to the new product launch (which is comprised of the timing and effectiveness of a future innovative competitive offering).

The participants received an analyst's report that included projections of all the key market factors (in the form of a distribution of estimates), and the output of a risk analysis based on these projections which resulted in a distribution of estimated NPVs assuming the product is launched. From this distribution of outcomes, subjects could determine the expected NPV and the probability that the project yields a positive NPV. The analyst's report also provided estimates of the relative importance of each of the key uncertainties in determining the products financial success, i.e., the coefficients from a regression analysis where the forecasted NPV from the risk analysis was the dependent variable.[7] Specifically, the relative importance weights were stated to be 70% for market share, 22% for competitive response, and 8% for market growth. In all, the analyst's report was positive enough to make the decision to launch highly attractive.

[7] The R^2 of this model was .98 indicating that these three factors almost perfectly linearly determine the resulting NPV. It is for this reason that we can subsequently assume that factor weights are independent of factor levels.

After receiving this information, participants in the "high initial commitment" condition were asked to make a launch/no-launch decision (we later describe the "low initial commitment" condition). If participants initially chose not to launch the product, they were told they were done with the task. If they chose to launch, the case continued with information about the product's performance after two years of being on the market. The information included the realized values of the three key factors. This information was presented in a balanced fashion such that even though the project was not doing as well as expected, there was some positive news. Thus, though the realized market share was significantly less than anticipated, the realized market growth rate was greater than anticipated, and the information regarding competitive reaction was neutral in that no new information was given. In addition, participants were told that an uncorrectable production problem explained the lower than expected market share, but that the sales-force and R&D groups were very happy about the product introduction. Nonetheless, since market share was the most important factor in determining profitability, and the production problem was uncorrectable, the balance of the information was negative. This fact was underscored by the analyst's revised (i.e., normatively updated) financial report that predicted a much lower expected NPV from future cash flows. More importantly, this NPV forecast was substantially lower than the certain cash value associated with the alternative of selling off the equipment. Thus, the normative decision *based solely on the new financial information* was to cash-out. After receiving the market feedback, participants were asked whether they wanted to continue or to withdraw the product.

Design

Given that our hypotheses concern how initial beliefs, commitment to an initial decision, and the introduction of new (negative) information affect the updated beliefs and thus the second decision, we manipulate each of these factors. This results in a two (level of initial information--positive or very positive) by two (level of new information--negative or more negative) by two (high initial commitment or low initial commitment) incomplete between subjects design. The low initial commitment level was crossed only with the level of new information, while the level of initial information was kept fixed at the low condition. Thus, the design included only 6 out

of the possible 8 cells. We use this incomplete design since there are no theoretically meaningful predictions about the interaction of commitment and the level of initial information.

We manipulated initial commitment to the product launch as follows. As described previously, in the high initial commitment condition, participants were put in the position of making the decision to launch or not launch the proposed product. 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. If they chose not to launch the product they were told it would not hurt their careers, but they would lose this particular opportunity to be on the fast track. Those who chose to launch the product were also told they were appointed to the brand manager position for the product. Thus, they had responsibility for all that happened with the product starting with the decision to launch.

In the low initial commitment condition, participants were not responsible for the launch decision, and were not asked to make a recommendation whether or not to launch. Instead they were told that their boss, who had a very centralized decision making style and was responsible for making the decision to launch or not, had asked them for their beliefs about the three key market factors and two summary measures of the potential financial success of the product. Furthermore, participants in this condition were told that their boss wanted unbiased opinions and thus all evaluations were to be provided to him anonymously. Therefore, even though they a) were exposed to all the information, b) had to evaluate it carefully, and c) had to provide private statements about their beliefs, they were not put in a position where they had to publicly commit to a course of action (or even a belief), or feel responsible for its outcome. In the second part of the case, these low initial commitment subjects were told that during the two years the product was on the market they were involved with a different project, and that they now were being recalled to review the product performance and decide whether to continue or cash-out.

We note that our commitment manipulation also implicitly manipulates the rewards for pursuing the project since if the launch succeeded the manager "would likely be on the fast track to a senior management position." This increases the likelihood that commitment, per se (H6), will explain escalation bias and thus decreases the chance that our proposed mechanism of biased belief updating (H1-H4) will entirely explain escalation bias. We choose to take this

conservative approach since it mimics typical organizational reward structures. Further, as discussed in the literature (Northcraft and Wolfe 1984; Bowen 1987; Heath 1995; Boulding et al. 1997), escalation bias might be an individually rational response to organizational reward systems. Thus, in our setting, support for H6 could be considered compatible with normative individual behavior, whereas support for H1-H4 would not be compatible with our normative benchmark.

Commitment Manipulation Check

To check on our initial commitment manipulation we collected separate data in two of the cells of our design. We did this as a separate experiment to avoid contaminating the second decision data due to collecting measures on degree of initial commitment.

Specifically, we obtained responses from 30 subjects in the low initial commitment, low prior information, low new information cell and 29 subjects in the high initial commitment, low prior information, low new information cell, thereby holding fixed prior and new information. In both cells we collected the same initial decision/evaluation data from the subjects as we did in the main study. Also, as in the reported experiment, subjects were informed that the company had decided to proceed with the product introduction. At this point we collected data designed to measure degree of initial (public) commitment. We asked subjects to indicate how well 10 different questions (reported in Appendix B) reflected their beliefs using a nine point scale anchored by "completely disagree" and "completely agree." Examples of these questions include "I am largely responsible for the decision to launch the product," "I feel attached to the decision to launch the product," and "My input into the decision process was public."

Before reporting these manipulation check results we note that the Cronbach alpha for the 10 items is .94. Further, principal components factor analysis indicates a single factor solution. We take these results as evidence that our 10 items are tapping the same underlying construct, which we denote as commitment to the initial decision.

Each of the 10 items is significantly different ($p<.001$ or better) between the low and high initial commitment groups and in the expected direction. If we sum the items and divide by 10, the mean of the low commitment group is 5.1 and the mean of the high commitment group is 7.2.

These means are significantly different ($t=6.84$, $p<.001$). We take these results as evidence of significantly different levels of commitment in our low and high initial commitment conditions.

Measures

Our theory revolves around three different measured constructs: factors levels, factor weights, and overall summary beliefs. After each decision (or evaluation, for the low initial commitment group) we took multiple measures of the participants' summary beliefs and individual measures for each of the factors levels.[8] We also took measures of the factor weights after the second stage of the experiment (recall that the normative weights for the three key uncertainties were given to the managers in the first stage of the experiment).

The two summary belief measures were the manager's assessment of expected NPV and probability of success, i.e., the probability that NPV would be positive. We combined these measures after first normalizing each measure (over the sample) to have a mean of zero, and standard deviation of one. We then summed for an individual the two measures to form the summary belief measure used in our analysis.

We had one measure each for market growth rate and market share factor levels. Both variables were measured using a 9-point scale anchored between "very negative" and "very positive" with the middle point designated as "neutral." The use of affective evaluations instead of actual expectations ensures that all participants are on the same “goodness” metric. This avoids the problem that one manager's view of a 20% market share may be better than another’s evaluation of a 40% market share. For the competitive response factor level we used two measures, the likelihood of competitive launch and the effect of a competitive launch. Both were measured on 9-point scales. Thus, as with our summary belief measure, we captured both the probability of the event occurring and the value of the event.

Finally, we measured the factor weights for the period two decision by asking participants

[8]We tested whether asking respondents to make the decision before giving their evaluations of the market factors affected their evaluations. We did this by comparing responses from the first stage of the experiment for the committed (made decision before providing evaluations) and non-committed (made no decision before providing evaluations) groups. We found no significant difference ($p>.70$) in the evaluations between these groups.

to specify the relative importance of the three market factors in determining the project's financial success. This importance was determined by asking subjects to divide a hundred points between the three factors.

ANALYSIS

Since we are interested in managers who make two decisions, we discard data from the 12 subjects who chose not to launch (or, presumably would have chosen not to launch in the low initial commitment condition).[9] Of these remaining 130 subjects, 68 (52%) decided to stay with the (losing) course of action, i.e., they exhibited a behavior compatible with an escalation bias. Given this behavior, we next analyze the process by which each manager evaluated the new information and updated his/her beliefs. We then analyze how their beliefs and degree of commitment affect their second decision.

Belief Updating

We start by examining the updating process for factor levels. In H1 we hypothesized that the updated factor level belief will be positively affected by the prior belief about the market factor level, the new information received about the market factor, and the prior summary belief (see equation 6b). To test H1 we use as the dependent variable our updated (time = 2) measure of the factor levels. The independent variables consist of the initial measured factor levels, dummy variables capturing the level of the new information provided for market share and market growth (no new information was provided for competitor response), and the initial

[9]Seven of the 98 participants in the high initial commitment conditions chose not to initially launch the product. All these subjects estimated probability of success to be less than 50%. The average for these seven subjects was 46%, versus 73% for those that continued. In addition, five of the 44 subjects in the low initial commitment conditions forecasted a negative NPV and/or a below 50% chance of a positive NPV. Given this initial negative forecast we assumed they would not have gone forward with the initial launch and thus eliminated them from further analysis. Further when we compare this initial decision not to launch for the two commitment groups, holding fixed the level of initial information, we note that 5/59 did not launch in the high initial commitment group while we assume 5/44 would not have launched in the low commitment group. These proportions are not significantly different.

summary belief.

Before reporting results of these estimates, we briefly summarize our discussion in Appendix A concerning the need to control for any individual specific unobserved effects that may correlate with our independent and dependent measures and therefore bias our obtained results. This need to control for unobserved effects is necessary because our analysis uses individual measures in addition to manipulated variables used in traditional experimental (ANOVA) analyses. Consequently, we use an instrumental variable (IV) procedure, where the selected instruments are known to be free from any individual specific effect. We then employ the Hausman (1978) test to identify whether this two-stage estimation procedure is required to ensure consistent estimates of our parameters of theoretical interest. The Hausman specification test result for our estimation of equation 6b is distributed $F_{2,119}$ for each of the three factor level updating equations.[10] The calculated values are 1.38, .024, and .85 for the market share, competitive response, and market growth equations, respectively. All three F values are insignificant and indicate that the ordinary least squares (OLS) results are free from bias due to omitted variables, i.e., unobserved subject specific differences. Thus, one can comfortably interpret the results as having the same validity as traditional ANOVA analysis.

We conducted our analyses as follows. First, we ran separate OLS equations for each market factor. Note that there is no new information about competitor response. Since there is no new information, our managers could not form a biased perception of the new information (i.e., equation 5). Therefore, there is no biased new information to substitute into equation 6a. Thus, for the competitive response market factor we would predict no effect for "new information" or the initial summary belief. As predicted, we found our estimates for the effects of the objective new information and the initial summary belief were not significantly different from zero for the competitive response market factor. Moreover, after removing these two insignificant variables, the coefficient on the initial level of competitive response, as predicted, was not significantly different from one.

We then we ran separate OLS analyses for updating the market share and market growth

---

10 The two degrees of freedom come from the fact that we instrumented the two measured independent variables, i.e., Factor Level and SB. The 119 degrees of freedom acknowledges that we had to drop some observation due to

factors and determined that the coefficients for the prior factor belief and prior summary belief did not differ significantly across the in the two different factor level updating equations. Given this empirical result and because there is no theoretical reason to expect these updating coefficients to differ, we present the pooled estimates for these two variables in Table 1.

Insert Table 1 about here

For purposes of testing H1, i.e., factor level bias, the key variable is the prior summary belief, SB_{it}, since normative factor level updating (equation 3) and our descriptive updating model (equation 6b) differ by this term. In support of this descriptive model and the underlying confirmatory bias process suggested by H1, we find that the effect of the prior summary belief is positive and significant.

H2 implies a bias in the factor weight updating process. Since, there is no reason to believe that the descriptive model, as specified in equation 7, varies across the different market factors, we again test whether we can pool estimation across the three market factors. We fail to find significant differences in the updating coefficient across the factors, so we again report the pooled estimates. In addition, as before, we first instrument all measured variables on the right hand side of equation, i.e., the change in factor levels. In this instance, the Hausman test statistic, which is distributed $F_{1, 364}$, equals 11.94, which is significant. Thus, the OLS estimates suffer from omitted variable bias. Consequently, we employ an IV estimation procedure to purge the independent variables of any unobserved individual level variation thereby ensuring consistent estimates and a valid test of our hypothesis.

The key variable for purposes of testing H2, i.e., updating of factor weights, is the change in factor levels. This term appears in the descriptive model (equation 7), but not in the normative model (equation 4). Our hypothesis of an underlying confirmatory bias process is supported by the fact that the coefficient on the change in factor levels variable is positive and significant ($t = 3.5$, $p < .01$). This result is compatible with the premise that subjects modify their weights upward or downward depending on how they perceived the new information. If subjects perceive the new information for a market factor to have a value more positive than their initial

---

missing variables.

value, subjects increase their weight on that factor. If the new value is less than the initial value, subjects decrease the weight of that market factor.

The two H3 hypotheses imply that initial belief stickiness and a biased perception of the business model lead to a bias in the updating of summary beliefs. Once again we test this hypothesis via regression analysis. We use the updated summary belief as our dependent measure and the independent variables consist of the prior summary belief, a measure of the business model bias (see equation 10A in Appendix A for details), and dummy variables representing the initial and new objective NPV levels. Since we do not observe (measure) the manager's private information we instrument all our measured variables thereby controlling for any omitted variable bias (i.e., any unobserved differences between subjects).

The key variables in this analysis are the prior summary belief and the measure of business model bias. These variables appear in our descriptive updating model (equation 11), but not the normative updating model of summary beliefs (equation 2). Hypothesis 3a suggests stickiness in the initial prior, and is tested by the coefficient on the initial summary belief. Hypothesis 3b suggests the biased perceptions of factor levels and factor weights enter into the updating process and is tested by the coefficient on the Business Model Bias measure. We report our estimates in Table 2. As seen from this table, the effects of both the prior Summary Belief and the Business Model Bias measure are positive and significant. Thus, we find support for our H3a and H3b, suggesting multiple, concurrent, sources of confirmatory bias in updating of summary beliefs.

Insert Table 2 about here

Summary of Our Updating Analyses

Estimation of our descriptive model (as stated in equations 6, 7, and 11) provides strong evidence for our proposed hypotheses. The initial summary belief has a direct, non-normative, impact on the updated summary belief (H3a). In addition, it non-normatively impacts the summary updated belief via the Business Model Bias term (H3b). This bias term is a function of the factor levels bias (H1) and the factor weights bias (H2). Further, in both instances the non-normative impact observed in the summary belief updating process is positive. Thus, these non-

normative impacts do not cancel out.

This updating process does not assume any role for commitment by the individual to an initial decision. Instead it centers on the premise that a person's prior summary belief non-normatively affects the updating process. Moreover, as reported earlier, we find that this non-normative updated summary belief affects the person's decision to continue or discontinue the project. Consequently, it appears that the person's positive prior belief non-normatively affects the second decision. Before further analyzing this link between prior beliefs and subjects' second decisions, however, we test whether initial commitment, per se, modifies the updating process by changing the strength of the non-normative aspect of the updating process, i.e., we test H5. We then further investigate how the person's summary belief affects choice and if initial commitment has a direct effect on the stop/continue decision, i.e., we test H4 and H6, respectively, in terms of the second decision.

The Effects of Initial Commitment on the Belief Updating Process:

H5 states that initial commitment indirectly affects managers' decisions by altering the way people update their beliefs about market factors and their overall summary evaluation, i.e., commitment induces bias in the updating process. To test this hypothesis we re-estimate equations 6b, 7, and 11 for two different samples: those that made an initial commitment to introduce the product, and those who only evaluated the initial entry possibility. We then ask whether the coefficient vectors (both slope and intercept) vary between the two groups for each of the three equations. The appropriate F-tests indicate that the estimated coefficients do not vary between groups for any of the estimation equations: for equation 6b the three values are $F_{4,122}$=1.51 for market share, $F_{1,127}$=1.19 for competitive response, and $F_{4,122}$=1.53 for market growth; for equation 7 $F_{3,362}$=1.14; and for equation 11 $F_{4,119}$=.76.

Thus, counter to the conjecture of Boulding et al. (1997), we fail to find support for H5. In other words, the act of making an initial commitment to a decision does not change the way people perceive and weight the new information or update their summary beliefs relative to people that made an initial evaluation but no commitment decision.[11]

[11] This is not surprising since we initially reported that the second (updated) summary belief was not affected by

The Effects of Initial Commitment and Summary Beliefs on Managers' Decisions:

We now turn our attention to the manager's stop/no stop decision. H6 states that initial commitment directly affects this second decision. H4 states that the biased summary belief drives the second decision. We examine the validity of these two hypotheses via a series of nested logit models where the second stop/continue decision serves as the dependent variable. We report these estimates in Table 3. In Model 1 we include only the two dummy variables representing the experimental manipulations for the initial and new (financial) information denoted NPV_t and NPV_{t+1}, respectively. In this instance only the effect of the initial information, NPV_t, is significant and we predict 58.5% of the choices correctly. The null model is the modal choice, which in our situation, is to continue and occurs 52% of the time. Thus, Model 1 increases prediction accuracy by 6.5 percentage points. In Model 2 we test for the importance of initial commitment, per se. Specifically, we add to Model 1 the dummy variable representing the experimental manipulation of initial commitment to a decision. In this model both the initial and new information conditions are insignificant while the effect of commitment is highly significant. Further, our ability to correctly predict choices increases to 60.8% and Model 2 fits the data significantly better ($p<.05$) than Model 1 according to the log likelihood criterion. This result is compatible with H6.

Insert Table 3 about here

To assess the impact of the updated summary beliefs (i.e., SB_{i2}), we estimate Model 3, which includes dummy variables for the initial and new NPV levels and our instrumented measure of the updated summary belief. Again, we note that NPV_t and NPV_{t+1} are insignificant while the effect of the updated summary belief is highly significant. Our ability to correctly predict choices increases to 82.6%, and the model fits the data significantly better ($p<.01$) than Model 1. Thus, these results indicate strong support for our descriptive model, i.e., H4.

Next, in Model 4 we pit H6 against H4 by jointly estimating the effects of both initial commitment and updated summary beliefs. As shown in Table 3, we are now able to correctly

---

commitment after adjusting for the initial information provided and the new information provided (t=1.13).

predict 84.3% of choices, though the model fits the data only marginally better ($p<.20$) than Model 3. The summary evaluation measure remains highly significant, while the effect of initial commitment is only marginally significant ($p<.15$). In this test we see only weak support for H6, i.e., a direct effect of commitment on subsequent decision choices, but strong support for H4.

In sum, analysis of the second decision suggests strong support for our descriptive model that proposes escalation bias due to belief updating bias (H4). We find no support for an alternative model that has an indirect effect of initial commitment on belief updating and therefore subsequent decisions (H5), and only marginal support for an alternative model that has a direct effect of initial commitment on subsequent decisions (H6). As a caveat, one explanation for the obtained results is the coarseness of our measure of initial commitment (i.e., a dummy variable) compared to the continuous measure of the summary belief. As a result, the power of the logit analysis regarding the initial commitment effect is lower.

DISCUSSION

In this research we study, within the context of a new product launch, the process underlying the phenomenon of escalation of commitment, i.e., the stylized fact that people tend to stick with a losing course of action. Our goal is to better understand the underlying process that leads to this observed outcome. However, before proceeding with our observations, we caution the reader that our results are based upon a case study of a hypothetical new product introduction. As always in these situations, one must be concerned with the external validity of the results. We leave this judgment up to the individual reader.

With this caveat in mind, our results are compatible with our proposed descriptive model of how managers update beliefs. Moreover, this model explains why a manager might stick to a course of action longer than one would predict using a normative updating process. Perhaps most interesting, we find initial commitment to a decision has no effect on the updating process once we controlled for the person's prior summary belief. Moreover, initial commitment does not have a significant effect on the second decision to stay the course after controlling for the summary belief. We take these findings as support for our conjecture that the escalation of commitment phenomenon is due to the person's initial summary belief and subsequent

confirmatory bias in updating these beliefs, rather than commitment to an initial decision.

This conclusion is strengthened when one considers that our commitment manipulation implicitly increased the reward for staying with the project. Thus, it appears that there is little support for the conjecture that managers are individually rational (normative) when they remain committed to a losing course of action. Instead, we find support for our premise that managers behave non-normatively because of "commitment" to their initial evaluation.

In particular, at the process level we find substantial evidence of a confirmatory bias that leads the manager to give too much weight to prior beliefs in three different updating processes. Specifically, we find undue weight given to the person's prior beliefs when managers update their beliefs about market factors, factor weights, and new summary beliefs. We also show empirically that these non-normative summary evaluations, in addition to the normative component of a person's summary evaluation, affect the actual decision to stay with the initial course of action. Moreover, all this is shown independent of whether or not the manager publicly commits to the initial decision to launch the product. Consequently, initial commitment to a decision and the attendant justification explanation is not needed to have a person stay with a (losing) course of action.

By providing a formal descriptive process model of how managers update their knowledge we are able to get new insights into the managerial decision making process and how confirmatory bias leads to distortions in information processing. We make three summary observations about the value of quantifying our descriptive and normative models. First, since our descriptive model is formalized into a mathematical system of equations we can subject it to empirical testing. Second, since the normative model nests within the descriptive model we can compare and contrast our descriptive model relative to a normative model benchmark. In this case we choose a Bayesian updating model and a standard NPV financial model as integral parts of our benchmark model. Third, this benchmarking allows us to formally identify (measure) the extent, and sources, of bias relative to normative behavior. This knowledge then becomes a starting point for identifying remedies to the observed problem.

Our model suggests that distortions (relative to our benchmark model) occur because managers fail to appropriately recognize information redundancies. Table 4 and Figure 1

summarize the three non-normative paths that we hypothesize are a part of the underlying process generating the stop/continue decision. Each of these paths is non-normative because they represent the use of redundant information. In our analysis, we find significant support for each of these non-normative paths. Thus, we conclude that updating is biased, relative to a normative Bayesian updating model. Managers appear to be "stuck in the past" in terms of their use of information.

Insert Table 4 about here

We believe that the model of belief updating presented in this paper should be useful in understanding belief updating in a wide range of situations. In addition, our model development and subsequent analyses highlight the inherent problems managers have dealing with information at different levels (disaggregate and aggregate), and with information redundancy. As these issues are part of most managerial decision making situations, further research on these issues has the potential to shed light on many behaviors.

We conclude by discussing implications of our model with respect to theorizing about possible remedies to the escalation phenomenon. For example, Boulding et al. (1997) found empirical evidence that using a different decision-maker to make the stop/continue decision and using predetermined stopping rules were effective in curtailing escalation. In contrast, providing more, and better, information seemed to have little, if any, effect. These findings are consistent with our model which implies that one must take prior information "out of play" to avoid over-weighting this information. Note, however, that decision de-coupling and pre-determined stopping rules are a necessary, but not sufficient approach for avoiding escalation bias. De-coupling eliminates escalation bias induced by prior commitment to a course of action. However, if the new decision maker previously formed initial beliefs, e.g., made an assessment of probability of success and/or came up with an initial belief about the expected financial performance of the project prior to receiving the new information, then the potential for escalation bias will remain. Our model suggests that to eliminate tendencies toward escalation bias, one must de-couple the new decision-maker not only from the original commitment decision, but also from the original evaluation process. Unfortunately, this implies that the

second decision is made without the personal experience gained from the first decision, which is, presumably, of value. However, this research highlights that past experience can be a double-edged sword in the context of managerial decision making.

References

Arkes, H. R., and C. Blumer
1985 "The Psychology of Sunk Costs," *Organizational Behavior and Human Decision Processes*, 35, 124-140.

Armstrong, J. Scott
1996 "Management Folklore and Management Science – On Portfolio Planning, Escalation Bias, and Such", INFORMS, *Interfaces* 24:4 (July-August).

Bazerman, Max
1984 "The Relevance of Kahneman and Tversky's Concept of Framing to Organizational Behavior," *Journal of Management*, 10, 333-343.

Bazerman, Max and Margaret Neale
1992 "Nonrational Escalation of Commitment in Negotiation," *European Management Journal*, 10(June), 163-168.

Boulding, William, Ruskin Morgan, and Richard Staelin
1997 "Pulling the Plug to Stop the New Product Drain," *Journal of Marketing Research*, 34 (February), 164-176.

Boulding, William, Ajay Kalra, and Richard Staelin
1999 "The Quality Double Whammy," *Marketing Science*, 18 (4), 463-484.

Boulding, William, et al.
1994 "Understanding Managers' Strategic Decision-Making Process," *Marketing Letters*, 5(4), 413-426.

Bowen, Michael G.
1987 "The Escalation Phenomenon Reconsidered: Decision Dilemmas or Decision Errors?," *Academy of Management Review*, 12(1), 52-66.

Brockner, Joel, and J. Z. Rubin
1985 *Entrapment in Escalating Conflicts*, New-York: Springer-Verlag.

Brocker, Joel
1992 "The Escalation of Commitment to a Failing Course of Action: Toward Theoretical Progress", *Academy of Management Review*, 17, 39-61.

Einhorn, Hillel J. and Robin M. Hogarth
1978 "Confidence in Judgement: Persistence of the Illusion of Validity," *Psychological Review*, 85, 396-416.

Hausman, Jerry A.
1978 "Specification Tests in Econometrics," *Econometrica*, 46, 1251-1271.

Heath, Chip
1995 "Escalation and De-escalation of Commitment in Response to Sunk Costs: The Role of Budgeting in Mental Accounting", *Organizational Behavior & Human Decision Processes*, Vol. 62, No. 1, April, pp. 38-54, 1995.

Hoch, Stephen J. and Young-Won Ha
1986 "Consumer Learning: Advertising and the Ambiguity of Product Experience," *Journal of Consumer Research*, 13 (2), 221-233.

Hogarth & Einhorn
1992 "Order Effects and Belief Updating: The Belief-Adjustment Model", *Cognitive Psychology*, 1992, Vol. 24, No. 1, January 1992, pp. 1-55.

Hsee, Christopher K.
1995 "Elastic Justification: How Tempting but Task-Irrelevant Factors Infuence Decisions", *Organizational Behavior & Human Decision Processes*, Vol. 62, No. 3, June, pp. 330-337, 1995.

Kahneman, Daniel and Amos Tversky
1979 "Prospect Theory: An Analysis of Decisions Under Risk," *Econometrica*, 47, 263-291.

Klayman, Joshua
1995 "Varieties of Confirmation Bias," in *Decision Making from a Cognitive Perspective*, Jerome Busemeyer, Reid Hastie, and Douglas C. Medin, eds. The Psychology of Learning and Motivation, Vol. 32, Advances in Research and Theory. San Diego: Academic Press.

Lukas, Paul
1998 "The Ghastliest Product Launches," *Fortune*, March 16, 44.

Mahoney, M. J. and B. G. Demonbreun
1978 "Psychology of the Scientist: An Analysis of Problem-Solving Bias," *Cognitive Therapy and Research*, 1, 229-238.

Mitroff, I.
1974 *The Subjective Side of Science*, Amsterdam: Elsevier.

Mynatt, C.R., M.E. Doherty, and R.D. Tweney

1977 "Confirmation Bias in a Simulated Research Environment: An Experimental Study of Scientific Inference," *Quarterly Journal of Experimental Psychology*, 29, 85-95.

Mynatt, C.R.
1978 "Consequences of Confirmation and Disconfirmation in a Simulated Research Environment," *Quarterly Journal of Experimental Psychology*, 30, 395-406.

Northcraft, G.B., & Wolf, G.
1984 "Dollars, Sense, and Sunk Costs: A Life-Cycle Model of Resource Allocation Decisions" *Academy of Management Review*, 9, pp. 225-234.

Pratt, John W., Howard Raiffa, and Robert Schlaifer
1995 *Introduction to Statistical Decision Theory*, Cambridge, MA: MIT Press.

Ross, Jerry, and Barry M. Staw
1986 "Expo 86: An Escalation Prototype," *Administrative Science Quarterly*, 31(June), 274-297.

Russo, J. Edward, Victoria H. Medvec, and Margaret G. Meloy
1996 "The Decision of Information During Decisions", *Organizational Behavior and Human Decision Processes*, 66 (1), 102-110.

Russo, J. Edward, Margaret G. Meloy, and Victoria Medvec,
1998 "Predecisional Distortion of Product Information", *Journal of Marketing Research*, November, pp. 438-452.

Samuelson, William and Richard Zeckhauser
1988 "Status Quo Bias in Decision Making," *Journal of Risk and Uncertainty*, 1, 7-59.

Schmidt, J.B. and Roger J. Calantone
1998 "Are *Really* New Product Development Projects Harder to Shut Down?," *Journal of Product Innovation Management*, 2, 111-123.

Staw, Barry 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, Barry M.
1981 "The Escalation of Commitment to a Course of Action," *Academy of Management Review*, 6, 577-587.

Staw, Barry M., and Jerry Ross

1986 "Behavior in Escalation Situations: Antecedents, Prototypes, and Solutions," in *Research in Organizational Behavior*, ed. B. M. Staw, and L. L. Cummings, Vol. 7, 297-332, Greenwich, CT:JAI Press.

Staw, Barry M., and Ha Hoang
1995 "Playing Time and Survival in Professional Basketball" in *Administrative Science Quarterly*, 40 (3), 473-494.

Whyte, Glen
1986 "Escalating Commitment to a Course of Action: A Reinterpretation," *Academy of Management Review*, 11(2), 311-321.

APPENDIX A

In this Appendix we provide a more detailed description of a) the structural models that are reflected in our presentation in the text, and b) the estimation approach we used to control for omitted variables. Our basic modeling approach is to assume that any factors affecting the variable of interest do so linearly and their effects are additive. During estimation, unobserved variables are controlled for using instrumental variables.

Structural Equations

Let the financial performance of the project as determined by the normative business model be represented by NPV_t where the subscript t captures the time when the business model is run. Also assume the underlying business model can be represented by a linear function of key factors that are believed to affect the business situation at hand. Denote the level of these factors as $FactorLevel_{kt}$ where the subscript k represents the kth factor. Also denote the weight of this factor as $FactorWeight_{kt}$. Then define NP_t to be a weighted linear combination of these key factors, i.e.,

(1A) $NPV_t = \Sigma_k (FactorWeigh_{kt})*(FactorLevel_t)$.

We acknowledge that a manager may use private information in addition to the normative business model to form some overall summary belief SB_{it},about the project. This leads to the following equation

(2A) $SB_{it} = NPV_t + PI_i$

where PI_i captures manager i's private information. We assume SB_{it} is the belief that determines the manager's decision to stay, or stop the project at time t.

When a manager receives new information, he/she must update his/her beliefs. We use a Bayesian updating model as our benchmark normative updating process. Under the assumption

of normal priors and a normal likelihood function, the normative model for the mean of the kth updated factor level, denoted as $UpdatedFactorLevel_{ikt+1}$, is linear combination of the old information about the factor level, $FactorLevel_{ikt}$, and the objective new information, denoted $ObjectiveNewInformation_{ikt+1}$, i.e.,

(3A) $$FactorLevel_{ikt+1} = \alpha_k FactorLevel_{ikt} + (1-\alpha_k) ObjectiveNewInformation_{ikt+1}$$

where α_k and $(1-\alpha_k)$ are the relative reliabilities of the two pieces of information.

In modeling how factor weights should be updated, we use the fact that in our situation no new information is given with respect to how the weights might change. Thus, our normative model for updated factor weight is

(4A) $$FactorWeight_{ikt+1} = FactorWeight_{ikt}.$$

Our descriptive models of the updating process reflect our assumption that managers exhibit a confirmatory bias. We use the current summary belief as our measure of a person's current beliefs, i.e., the belief he/she is trying to confirm.

Our model of how a person perceives new factor level information states that the person's perception is more in line with his/her prior beliefs than would be the perceptions of an objective observer. We represent this biased perception of data as $BiasedPerceptionOfNewInfo_{ikt+1}$ and assume

(5A) $$BiasedPerceptionOfNewInfo_{ikt+1} = (\delta) ObjectiveNewInfo_{ikt+1} + (1-\delta) SB_{it}$$

where $0 \le \delta \le 1$. When $\delta = 1$, there is no confirmatory bias with respect to the updating of the level of a factor. When $\delta < 1$ there is some confirmatory bias.

We next assume a person blends this biased perception of the data concerning a factor level with his/her prior mean belief about this factor level. This yields the following:

(6A) $BiasedFactorLevel_{ikt+1} = \alpha_{1k} FactorLevel_{ikt} +$

$(1-\alpha_{1k})\ BiasedPerceptionOfNew\ Info_{ikt+1}$

$= \alpha_{1k} FactorLevel_{ikt} + (1-\alpha_{1k})[\delta ObjectiveNewInfo_{ikt+1} + (1-\delta) SB_{it}]$

$= \alpha_{1k} FactorLevel_{ikt} + \alpha_{2k} ObjectiveNewInfo_{ikt+1} + (1-\alpha_{1k}-\alpha_{2k}) SB_{1t}$

where $\alpha_{2k} = (1-\alpha_{1k})\delta$.

Our model of how managers update the factor weights used in the prior period, $FactorWeight_{ikt}$, assumes these managers increase a factor weight if their updated perceptions of the new factor level is larger than before and they decrease the weight if converse holds. This leads us to postulate the following equation;

(7A) $BiasedFactorWeight_{ikt+1} = FactorWeight_{ikt} + \beta(BiasedFactorLevel_{kt+1} - FactorLevel_{ikt})$

where $\beta > 0$. If $\beta = 0$, there is no confirmatory bias.

We model the updating process for summary beliefs by first acknowledging that a manager's perception of the current business model is a function of his/her biased perceptions of factor levels and factor weights. Specifically, we assume this biased perception captures the true business model performance plus some bias, i.e.,

(8A) $BiasedNPV_{it+1} = NPV_{it+1} + BusinessModelBias_{it+1}$,

where $BiasedNPV_{it+1}$ is the manager's perception of the financial performance. We measure the business model bias by assuming a manager forms overall perceptions of the business model in the same functional form as the normative business model, i.e.,

(9A) $BiasedNPV_{it+1} = \Sigma(BiasedFactorWeight_{kt+1}) * (BiasedFactorLevel_{ikt+1})$.

We measure the business model bias by subtracting the business model NPV from the measured

managers perception of this NPV, i.e.,

(10A) $BusinessModelBias_{it+1} = BiasedNPV_{it+1} - NPV_{it+1}.$

We next assume the summary belief is a blend of a) the manager's perceptions of the current business model and any private information and b) the non-normative influence of prior summary beliefs, i.e.,

(11A) $SB_{it+1} = (1-\gamma)[Biased\ NPV_{ikt+1} + PI_i] + \gamma SB_{it}$

where $0 < \gamma < 1$. If $\gamma = 0$, there is no direct confirmatory bias effect.

Using equation 8A, we get

(12A) $SB_{it+1} = (1-\gamma)\ [NPV_{it+1} + BusinessModelBias_{it+1} + PI_i] + \gamma SB_{it}.$

Thus, we expect a manager's updated belief to be affected three different ways by this person's prior belief. The first two are indirect and occur because of SB's influence on the perception of factor levels and factor weights as captured by the *BusinessModelBias* term. (See equations 9A and 10A.) The latter occurs because of SB's direct influence.

Estimation

To test our hypotheses we estimate equations 6A, 7A and 12A. In doing so, we point out the following. First, each equation captures its normative counterpart with the appropriate parameter restrictions. Thus, when $\delta=1$ (and therefore $\alpha_{2k} = 1-\alpha_{1k}$) equation 6A becomes equivalent to equation 3A. Likewise when $\beta = 0$ in equation 7A, this equation is equivalent to equation 4A. Finally when $\gamma = 0$ in equation 12A (and $\delta = 1$ in equation 6A and $\beta = 0$ in equation 7A), we see that equation 12A is equivalent to equation 2A (after substituting in equation 1A).

Second, note that we use measured variables instead of manipulated variables to test our

hypotheses. In addition, equation 12A explicitly includes one individual level variable that is not measured, i.e., PI_i. These two facts lead to the possibility that each of the equations may contain individual specific unobserved effects that are correlated with both the independent and dependent measures. For example, a person might exhibit "optimism" and always use the high end of the scale for his/her reported measures. For these reasons, we use an instrumental variable (IV) procedure, where the selected instruments are known to be free from any individual specific effect.[12] We then employ the Hausman (1978) test to identify whether this two-stage estimation procedure is required to ensure consistent estimates of our parameters of theoretical interest.

Finally note that in many instances, the measurement scales used for the dependent and independent variables differs. Consequently, it is not possible to directly estimate many of our structured parameters, but instead only test for statistical significance. However, in those instances where the scale used for both the independent and dependent variables is the same we are able to get direct estimates of the structural parameter in addition to its statistical significance.

12 We use an instrumental variables procedure whenever our measures appear on the right hand side of an equation. We only use as instruments those variables that we know are free from an unobserved individual specific effect, i.e., the experimental treatments (commitment, initial prior, new information) and a vector of difference variables Z_{ikt} - $Z_{i.t}$, where $Z_{i.t}$ is the average of the three market factors within each individual.

APPENDIX B

MANIPULATION CHECK QUESTIONS

Respondents were asked to assess "How well do the following statements reflect your beliefs" on a 9-point scale ranging from "Completely disagree" to "Completely agree" for the following 10 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.
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.

Table 1

Pooled Market Factor Level Updating Model Estimation (Equation 6b)

Dependent Variable: Factor Level Belief$_{ikt+1}$

Variable	Parameter estimate	Standard Error
Market Share Intercept	3.114***	.467
Market Growth Intercept	5.570***	.704
Factor Level Belief$_{ikt}$	0.128**	.057
SB_{it} (Prior Overall Belief)	.0049**	.0028
New Information: Market Share	.650***	.250
New Information: Market Growth	-.075	.250

*** Significant at the .01 level.
** Significant at the .05 level.

Table 2

Overall Evaluation Model (Equation 11)

Dependent Variable: SB_{it+1} (Current Overall Belief)

Variable	Parameter estimate	Standard Error
Constant	33.55***	6.25
SB_{it} [a] (Prior Overall Belief)	0.25***	0.066
Business Model Bias[a]	0.285***	0.054
NPV_{t+1}	4.07	3.77

*** Significant at the .01 level.

[a] We instrument these variables in estimation.

Table 3

Logit Models of the Continue Decision[a]

Model	Intercept	Meas. Sum. (SB_{it+1})	NPV_t[b]	NPV_{t+1}[b]	Initial Commit[b]	-2 log L	% correct
1	-0.27 (0.3)	-	0.73* (0.07)	0.34 (0.34)	-	175.715	58.5
2	-0.81** (0.05)	-	0.39 (0.38)	0.41 (0.26)	0.85** (0.05)	171.816	60.8
3	-4.94*** (0.001)	0.094*** (0.001)	-0.11 (0.83)	0.12 (0.77)	-	124.609	82.6
4	-5.36*** (0.001)	0.092*** (0.001)	-0.46 (0.43)	0.21 (0.64)	0.80 (0.12)	122.135	84.3

[a] p-values in parentheses.
[b] Dummy variables capturing the experimental manipulations.
*** Significant at the .01 level.
** Significant at the .05 level.
* Significant at the .10 level.

Table 4

Summary of Tests For Non-Normative Updating Processes

Updating Process	Normative Equation	Descriptive Equation
Factor Levels[a]	3	6b
Factor Weights[b]	4	7
Summary Beliefs[c]	2	11

[a] Test of H1: Significance of SB_{it}

[b] Test of H2: Significance of (Biased Factor $Level_{t+1}$ – Factor $Level_t$)

[c] Test of H3: Significance of Business Model Bias and SB_{it}

EXHIBIT 1

The Quality Valve Company
Confidential Product Analysis

Analysts' Summary

The Investment

- The project will require the purchase of a special valve producing machine that will cost $2,500,000, which must be purchased before the launch.
- The machine is expected to have a 10 year useful life, but will be virtually worthless at the end of 10 years. The machine utilizes the new coated valve technology, but doesn't require any design modifications, which makes it easy to sell before the end of the useful life if the project is terminated. Recently, similar machines have sold quickly on the market for roughly the depreciated book value of the machines.

The Market

- Great Lakes Valve Company is the only major competitor to Quality and is the current market leader for emissions control valves for large trucks.
- The initial size of the market is expected to be about 63,200 units.

Key Assumptions

- Market Growth Rate (MGR):

The market growth rate is forecasted to be 3% per year. However, the analyst recognizes that he could be either high or low on this estimate. When pressed about his confidence on the 3% estimate he acknowledged that the highest he saw growth being was 7% per year. Moreover, he felt there was only a 25% chance that the growth would exceed 4%. More generally, he gave the following estimates:

0% chance the growth rate exceeds 7.0%
25% chance the growth rate exceeds 4.0%
50% chance the growth rate exceeds 3.0%
75% chance the growth rate exceeds 2.0%
100% chance the growth rate exceeds -1.0%

Quality's Initial Market Share:

Based upon Great Lakes' prior success with introducing their new product, the analyst believes there is at least a 50% chance of garnering a 42% market share by the end of the first year. As with the market growth estimate, there was some uncertainty about this 42% figure. Thus, it could be as low as 10% and as high as 65%. More generally, he gave the following estimates:

0% chance the initial market share exceeds 65.0%
25% chance the initial market share exceeds 48.5%
50% chance the initial market share exceeds 42.0%
75% chance the initial market share exceeds 33.0%
100% chance the initial market share exceeds 10.0%

- Loss of Market Share

The analyst also acknowledged that Quality probably could not keep this market share indefinitely. He saw two ways that Quality could lose share. First, it could lose share to some of the small firms that would try to compete on price. Such competition could easily result in a loss of 5% of Quality's business in the second year of operation and another 2.5% in the third year.

The second way Quality could lose market share is if Great Lakes re-enters the market with a yet to be discovered technological improvement. It is expected that Great Lakes will start such R&D, but that it will take several years before a product will be ready. The earliest a new product could be ready by Great Lakes is year 3, but it is expected that it will be year 6 when Great Lakes launches a new product.

After much discussion, it was decided that the probability that Great Lakes would enter in any given year was a function of Quality's initial success in re-entering the specialty valve market. Thus, if Quality is highly successful, Great Lakes will be more inclined to conduct research, increasing the probability of an entry by Great Lakes.

Figure 1

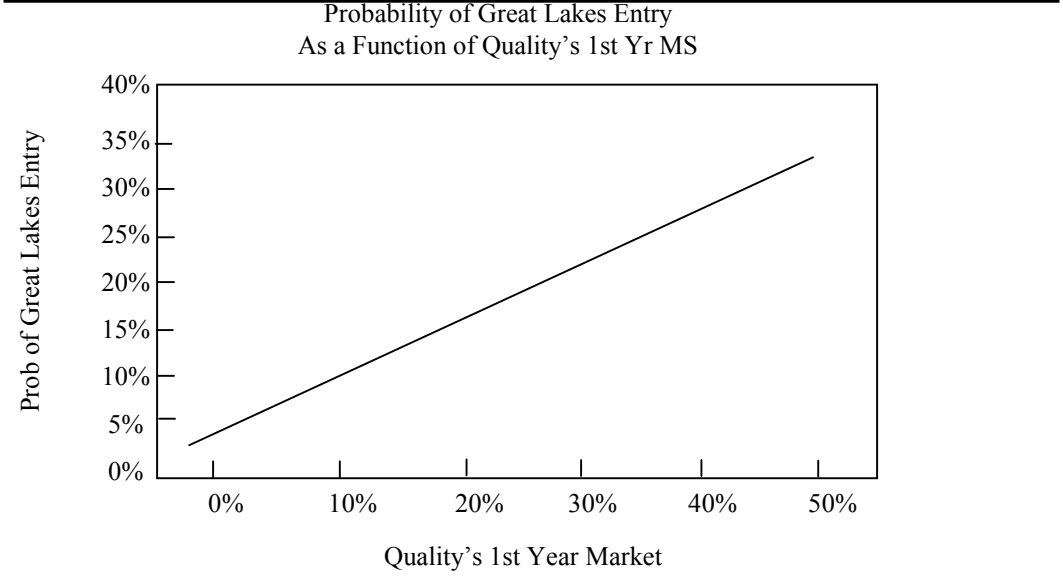

- Effect of Great Lakes' New Product on Quality's Market Share:

The last key uncertainty in the decision is the success of Great Lakes re-entry into the market assuming they are able to come up with a new technological discovery. The analyst's current belief is that Great Lakes would steal 70% of Quality's existing market (i.e. Quality would be able to keep 30% of its share). However, the analyst was very uncertain about this estimate. He quantified his beliefs in terms of the percentage of Quality's market share after Great Lakes' entry to its market share prior to Great Lakes entering. Below are his estimates.

0% chance	the percentage exceeds 80.0%
25% chance	the percentage exceeds 45.0%
50% chance	the percentage exceeds 30.0%
75% chance	the percentage exceeds 20.0%
100% chance	the percentage exceeds 0.0%

Financial Forecast

An analysis was performed acknowledging the above stated uncertainties about the four key assumptions. Specifically, the analyst used a risk analysis by running 500 different scenarios in accordance with the probability distributions for the four key inputs. Each scenario resulted in cash flows for the life of the project. These cash flows were summarized in terms of a NPV figure.

A number of observations were made from this analysis. First, it was possible to determine the impact of each of the key uncertainties on the NPV figure. Thus, the following weights were determined:

Market growth rate	-	8%
Initial market share	-	70%
Competitive reaction (including time of entry and effect of entry)	-	22%

Second, the 500 scenarios generated a distribution of possible outcomes as shown below.

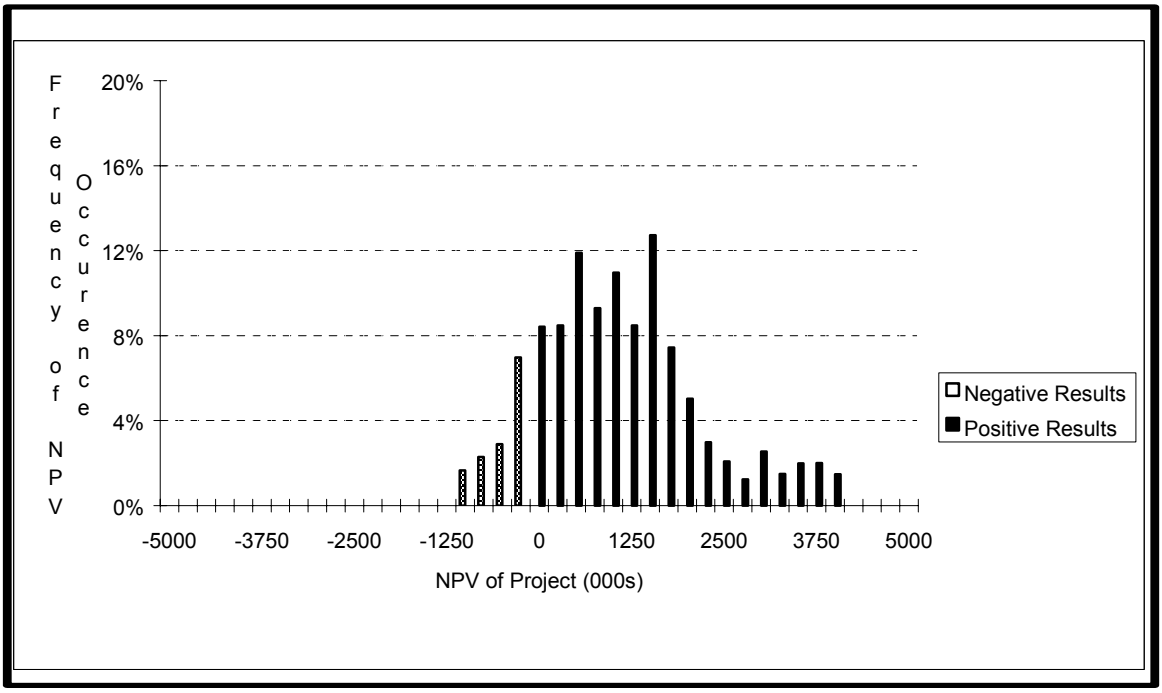

Simulation Results

Expected (Mean) Result	$1,117,408
Maximum Result	$4,373,696
Minimum Result	-$1,267,021

Overall the risk analysis shows that the average of all the scenarios yields a NPV over the 10-year horizon of roughly $1,117,000. The graph shows that most of the projected scenarios yield a positive NPV. As a reference point, McKinsey and Company reports that only 33% of innovative new products succeed. The graph also shows, however, that there is some chance the product will not be a financial success (i.e., will result in a negative NPV). In fact, there appears to be some small chance that the project could have a NPV over the life of the project of -$1,000,000 or less. Still, most of the occurrences are positive and if the launch resulted in the most optimistic forecast, an NPV over $4 million could be obtained.

Exhibit 2

Actual Results

You were assigned to manage the product. It took you a year to roll out the new valve. Over that time you were not able to capture as much market share as you had planned. In fact, after the first year the new valve only had 20.3% market share. The primary reason you got less than the expected 35% market share appears to be a result of the production process. While the valve tested extremely well when produced in small batches, production in mass quantities appears to create a significant amount of heat, which degrades the valve coating. Unfortunately this condition does not appear to be correctable, but even with this problem the valve is an improvement over any competitive offering, as indicated by the valve's penetration rate.

In fact, the salesforce was very pleased with the new offering and the R&D staff were elated to see their new discovery introduced into the marketplace. During the second year things did not change substantially. As predicted, the market share decreased somewhat to 18.2%, due to price cutting from a few of the small players. Great Lakes, however, did not lower its price and in no way seemed to retaliate. No new information was obtained with respect to their possible R&D activity or probability of introducing a new valve.

On a positive note, the total industry seemed to be growing faster than anticipated. Over the last two years the growth averaged 5% per year and does not seem to be slowing down.

Finally, the accounting department has provided you with aggregate cashflow estimates for the project which put the current NPV (i.e. the NPV of cashflows through year 2 only) at -$1,250,000, indicating that the project had already recouped $1,250,000 of the initial $2,500,000 investment.

The analyst that did the initial work on evaluating the success of the launch updated his analysis of expected NPV to reflect the actual results obtained in the last two years. This updated analysis reflects the actual results at the end of two years. Also, at your request he conducted an analysis to determine what would occur if Quality was to drop out of the market and use the freed up money to invest in other projects. In doing so he verified that the equipment could be sold currently for book value and that there were a number of other projects that could be funded if this venture were discontinued.

The results of his analysis are attached below and reflect the best prediction of the future NPV cashflows of the project (i.e. $257,183) compared with cashing out the investment (i.e. $669,466).

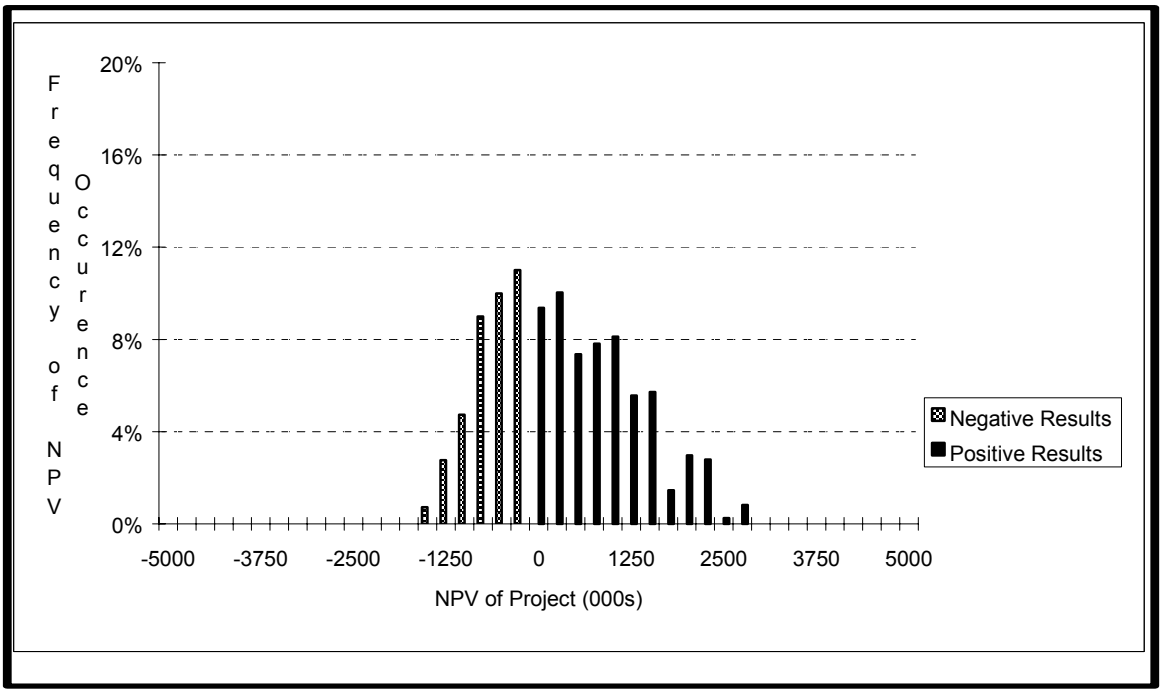

Option 1 - Stay in the market

Expected (Mean) Result	$257,183
Maximum Result	$3,050,443
Minimum Result	-$1,643,159

Option 2 - Cash out now

Expected result $669,466