
Georgetown University Institutional Repository
<http://www.library.georgetown.edu/digitalgeorgetown>

The author made this article openly available online. Please [tell us](#) how this access affects you. Your story matters.

Bond, S. D., Carlson, K. A., Keeney, R. L., (2008). Generating Objectives: Can Decision Makers Articulate What They Want? *Management Science*, 54(1), 56-70.
doi: <http://dx.doi.org/10.1287/mnsc.1070.0754>

Collection Permanent Link: hdl.handle.net/10822/1

© 2008 Institute for Operations Research and the Management Sciences (INFORMS).

This material is made available online with the permission of the author, and in accordance with publisher policies. No further reproduction or distribution of this copy is permitted by electronic transmission or any other means.

Generating Objectives: Can Decision Makers Articulate What They Want?

Samuel D. Bond

College of Management, Georgia Institute of Technology, Atlanta, Georgia 30318,
sam.bond@mgt.gatech.edu

Kurt A. Carlson, Ralph L. Keeney

Fuqua School of Business, Duke University, Durham, North Carolina 27708
{kurt.carlson@duke.edu, keeneyr@aol.com}

Objectives have long been considered a basis for sound decision making. This research examines the ability of decision makers to generate self-relevant objectives for consequential decisions. In three empirical studies, participants consistently omitted nearly half of the objectives that they later identified as personally relevant. More surprisingly, omitted objectives were perceived to be almost as important as those generated by participants on their own. These empirical results were replicated in a real-world case study of strategic decision making at a high-tech firm. Overall, our research suggests that decision makers are considerably deficient in utilizing personal knowledge and values to form objectives for the decisions they face.

Key words: decision analysis; multiple criteria; organizational studies; goals; decision making; theory

History: Accepted by David E. Bell, decision analysis; received August 2, 2006. This paper was with the authors 4¹/₂ months for 2 revisions.

Forgetting our objectives is the most frequent stupidity in which we indulge ourselves.

—Nietzsche (1879, p. 642)

1. Introduction

The ability to make decisions is what gives us control over our future. However, it is not merely control that we seek, but influence to make things better. Speaking broadly, the term “better” connotes achievement in terms of personal or organizational objectives. Therefore, across a variety of fundamental and applied decision-making disciplines (decision analysis, decision support, operations research, etc.), an important issue is the ability of individuals to generate personal or organizational objectives for a given decision. This paper addresses two fundamental questions that are relevant to this issue: (a) To what extent are decision makers able to list their own objectives, and (b) how important are objectives that decision makers overlook, but later acknowledge as relevant?

To our knowledge, the ability of individuals to generate personally relevant decision objectives has not been previously addressed. We examine this question with a series of studies designed to appraise the comprehensiveness of objectives generation in a variety of decision settings. The remainder of this paper is organized as follows. The next sections describe prior work relevant to the generation of objectives, summarize our research approach, and state our hypotheses. We then present three studies, each involving actual

decisions of substantial importance to the participants involved. The first two studies, which are retrospective but real, concern MBA and executive-MBA students choosing a business school for their education. The third study involves the important prospective decision of selecting a summer internship, using a sample of MBA students currently making that decision. Following these studies, we describe a real-world application in which one of the authors helped specify the objectives of a large high-tech company, based on views individually elicited from 12 executives. We conclude by discussing the implications of our findings for both individual and organizational decisions.

1.1. The Importance of Knowing One's Objectives

In any decision situation, what is “better” for an individual can be captured by that individual's set of objectives for the decision. Each objective represents a desire that the decision maker wants to achieve by making a decision.¹ As a simple example, three objectives of a summer vacation may be to see a new part of the country, to pursue outdoor activities, and to keep costs under control. For a manager overseeing the introduction of a new product, relevant objectives may include profit maximization, market-share growth, and enhancement of the company's reputation for quality offerings.

¹ Other terms such as criteria, goals, or value attributes are sometimes used to identify what is meant here by objectives.

Identification of objectives has long been considered a critical step in making good decisions (Keeney and Raiffa 1976). In his much-cited letter of advice to Joseph Priestly on making difficult choices, Benjamin Franklin (1772) outlines a method of balancing the pros and cons in terms of objectives. Peters and Waterman (1982, p. 279) refer to their "one all-purpose bit of advice for management" as "figure out your value system," by which they mean the articulation of a clear and complete set of objectives. Management gurus such as Peter Drucker (1954) have long contended that a business is best served when specified objectives are relied upon to coordinate company interests with courses of action. The idea that knowledge of objectives is fundamental to firm success continues to be a popular theme in managerial wisdom (e.g., Barrett 2006).

Whether one wishes to help a decision maker with clear qualitative thinking or with a quantitative analysis, a key element in the process is a comprehensive set of objectives. If a decision maker overlooks a key objective, then information that could be used to distinguish the options on this dimension may be ignored, increasing the likelihood of an inappropriate decision. The utility of considering one's objectives is emphasized by Payne et al. (1999, p. 251), who note that poor choices frequently occur because individuals "...may try to solve the wrong problem using options or objectives that would not be most critical to them upon reflection." Indeed, Baron (1997) argues that the failure to force individuals to examine their fundamental values is a critical flaw in common prescriptive techniques. In organizational decision making, Nutt (1998) demonstrates that the incorporation of a wider set of criteria results in better decisions, and that firms often include too few criteria when evaluating strategic alternatives. As an advocate of *value-focused thinking*, Keeney (1992) has provided broad theoretical support for the argument that contemplation of one's values is an important component of effective decision making, and various investigations have substantiated this idea in organizational settings; e.g., Leon (1999) demonstrates that value-focused thinking results in a more inclusive representation of the decision than does processing focused on alternatives.

1.2. Are Individuals Able to Generate Their Objectives?

Although little prior research has addressed the generation of decision objectives, scholars have spent decades examining the ability of individuals to deal with complexity in judgment and decision tasks (Kahneman et al. 1982). In general, the evidence reveals that decision makers are remarkably incomprehensive. Numerous empirical demonstrations have shown that individuals generally fail to

consider relevant information during the evaluation of alternatives, focusing only on that which is most prominent or salient (e.g., Tversky et al. 1988, Payne et al. 1993). When faced with more than a few alternatives, decision makers tend to restrict their processing to a narrow subset of possibilities (Lussier and Olshavsky 1979, Beach 1990, Gilbride and Allenby 2004). Furthermore, individuals are markedly deficient in considering the range of possible outcomes that may occur (Fox and Clemen 2005), and they generally do not consider alternative outcomes to an event that has transpired (Fischhoff 1975). Related research has documented the striking inability of individuals to generate comprehensive sets of explanations in causal reasoning tasks (e.g., the "pruning bias" of Fischhoff et al. 1978; cf. Russo and Kolzow 1994). Although these inadequacies constitute distinct phenomena, each involves the adoption of an overly narrow mental representation for the judgment or decision task at hand. Commonly, such failures are portrayed as the result of an adaptive mechanism by which individuals make efficient decisions in the face of processing limitations (Simon 1955). In addition, some researchers have suggested that certain choice settings tend to heighten the salience of certain decision objectives: being able to justify one's choice, minimizing effort, etc. (e.g., Bettman et al. 1998, Payne et al. 1999).

Despite widespread acknowledgment of the inadequacies just described, it has generally been assumed that as long as individuals possess adequate knowledge in the domain, they should be able to construct a list of relevant objectives for a particular decision. A different possibility is that the generation of decision objectives is subject to the same cognitive constraints that underlie the failures cited above. If decision makers adopt an overly narrow mind-set during the generation process, they may be incapable of conjuring all the objectives that matter to them. Indeed, this possibility was voiced long ago by Benjamin Franklin (1772, p. 236): "When these difficult cases occur, they are difficult, chiefly because while we have them under consideration, all the reasons pro and con are not present to the mind at the same time; but sometimes some set present themselves, and at other times another, the first being out of sight." Along similar lines, we suggest that even when individuals possess decision-relevant knowledge and personal values, they may be unable to generate a comprehensive list of objectives. As a consequence, decision making may typically be guided by whatever incomplete set of objectives is currently accessible.

Decision objectives are sometimes divided into two categories, referred to here as *fundamental objectives* and *means objectives*. Fundamental objectives represent desires that are an end in themselves for the situation

at hand, and accomplishment of a fundamental objective provides direct utility to the individual. Conversely, means objectives provide utility by facilitating the achievement of other objectives. In general, the set of fundamental objectives provides the basis to evaluate various alternatives, whereas means objectives help to stimulate the generation of alternatives. In educational settings, for example, some people pursue learning opportunities solely to enhance their likelihood of receiving a pay raise. For these individuals, the objective "increase my knowledge" serves as a means to the fundamental objective "increase my salary." Other people derive satisfaction from their store of knowledge in a particular domain; for these individuals, "increase my knowledge" is a fundamental objective for decisions involving educational opportunities. Hence, what is a means objective to one individual may be a fundamental objective to another facing the same decision. Therefore, we do not distinguish between fundamental and means objectives in the studies below, focusing instead on whether individuals are inadequate in generating either type.

2. Overview of the Present Research

Applying the arguments above to the realm of important personal decisions, it is plausible (but by no means certain) that individuals will be unable to bring to mind a sizable portion of personally relevant objectives. We examined this question with three studies; although the design of each study incorporated distinct features, all three applied a similar framework. First, participants read a brief introduction describing an important decision of personal significance. After doing so, they were asked to generate and list all objectives that were personally relevant to this decision. Next, participants were shown a *master list* containing a variety of potential objectives that might pertain to the decision, and they checked all the objectives from this master list that they deemed to be personally relevant. Therefore, for any particular respondent, the master list of objectives could be partitioned into three subsets: (1) objectives not considered relevant by that individual (i.e., those left unchecked), (2) relevant objectives that were generated without the aid of our master list (we label these *self-generated objectives*), and (3) relevant objectives that were omitted during the initial generation process but checked on the master list (*recognized objectives*). The objectives in the last two subsets served as the focus of our empirical exploration. That is, by examining the quantity and importance of objectives in the self-generated and recognized subsets, we are able to assess the extent to which individuals fail to generate objectives that are relevant and important.

For all three studies, we selected consequential decision contexts that were personally meaningful to

the participant population, ensuring that participants in each study were knowledgeable regarding the decision domain. The first two studies utilized current or incoming MBA students, and the decision involved the selection of an MBA program. Because the participants had recently undertaken the focal decision, decisions in these studies were retrospective. Therefore, the design provided a conservative test of our hypotheses, and any failure to generate important objectives would be especially striking. In contrast, the third study investigated our ideas in the context of a decision that was ongoing at the time: the selection of a summer internship by first-year MBAs. As such, Study 3 involved an impending decision of both immediate and long-term consequence.

Even for consequential, long-term decisions, it would hardly be surprising if individuals cannot bring to mind a few objectives of trivial importance. However, on the basis of arguments presented above, we expected the inadequacy to be considerably more severe. Two main hypotheses were examined, the first of which was that individuals will be incomplete in generating decision objectives. Evidence for this hypothesis would take the form of a substantial number of *recognized objectives* across the studies. Second, we hypothesized that objectives that were initially overlooked would be nearly as important as (or more important than) those generated independently. Evidence for this hypothesis is obtained by comparing the importance ratings of *recognized objectives* to that of *self-generated objectives* for each participant. Finally, we considered it likely that even when they are allowed multiple opportunities to ruminate on their objectives, decision makers will nonetheless overlook objectives of critical importance. This idea is examined in Study 3.

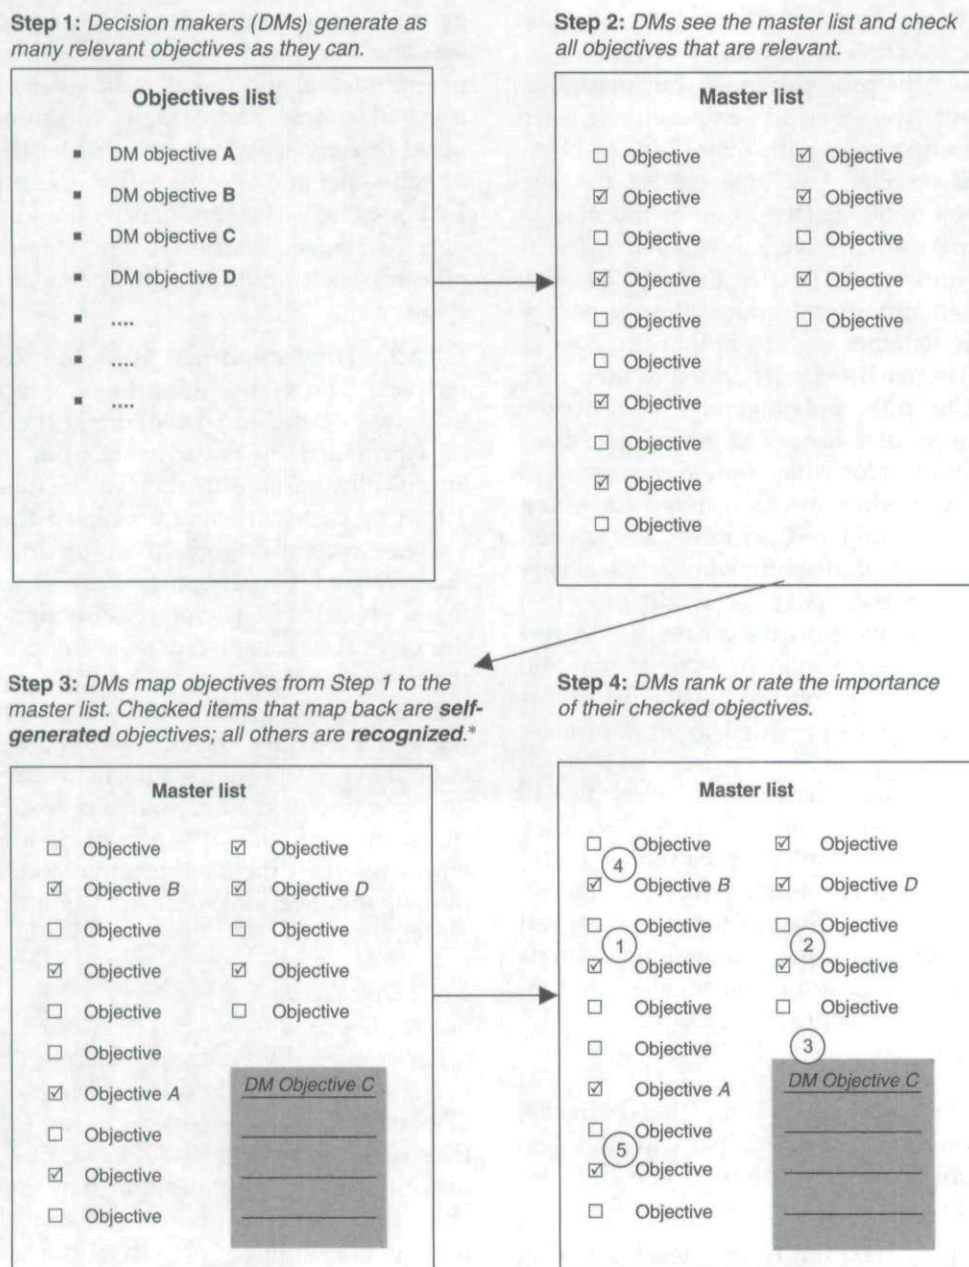
3. Study 1: Do Individuals Generate Incomplete Lists of Objectives?

The first study was an initial examination of our hypothesis that individuals will fail to generate important objectives during contemplation of an important decision. Because the designs of all three studies were similar, descriptions of Studies 2 and 3 will focus on modifications to the basic framework presented here.

3.1. Methods

3.1.1. Participants and Design. Participants were 84 full-time MBA students from eight different East Coast universities who were taking part in an annual charity event. The study was administered in a group setting along with other unrelated studies and a \$5 donation was made for each participant. The same packet of materials was given to each participant, and

Figure 1 Outline of Study Procedure



*Note that the mapping need not be one-to-one: Multiple items on the self-generated list might map to the same item on the master list, and one item on the self-generated list might map to many on the master list.

the entire study took approximately 15 minutes to complete.

3.1.2. Materials and Procedure. Figure 1 provides an overview of the procedure. The study was administered via paper and pencil. Participants were asked to envision that they were planning to pursue an MBA but had not yet decided on a school. Each participant was then instructed to list the objectives that (s)he thought were “relevant when making this decision.” The page provided 19 blank lines on which participants could list the objectives that mattered to

them. Beside each line were the letters A to S (used later to reference the objectives).

The next page presented a master list of objectives that might be relevant to the MBA decision. This list was created by the experimenters through brainstorming, internal discussion, and communication with prior MBA students. In total, the master list contained 28 items covering a wide variety of criteria (personal and professional, short- and long-term, etc.). Sample objectives included “personal growth potential,” “availability of electives,” and “geographic location.” Beside each objective was an empty checkbox

and a line. Participants were instructed to put a check in the box next to any objective that they thought was relevant to their selection of an MBA program.

The next step in the procedure was the "mapping" phase. Participants were told to compare their listed objectives on the first page with their checked objectives on the master list. Any time one of the previously listed objectives matched one of the master-list objectives, participants were instructed to write its letter (A-S) on the line next to the objective that it matched. This mapping procedure allowed participants to identify whether objectives they checked as relevant on the master list corresponded to their own listed objectives from the previous page; i.e., checked objectives on the master list could be mapped back to any listed objectives for which they were perceived substitutes. The procedure also accounted for differences in the specificity of listed objectives and master-list objectives; e.g., a participant who listed a very general objective on the first page might map that item to multiple objectives on the master list. A few participants listed one or more objectives that did not map to any objectives on the master list; therefore, additional space was provided to write in these unmapped objectives.

Next, participants were asked to identify the 10 most important objectives from the master list. They ranked these objectives by placing a number (1-10) above the corresponding objective on the master list (including any unmatched objectives that had been written in). Finally, a few follow-up questions requested information regarding gender, MBA concentration, and other demographic variables.

3.2. Results

Twenty-two packets were not completed correctly, yielding a usable sample of 62 participants. The follow-up questions revealed little of value and will not be discussed here.

3.2.1. Was Self-Generation Incomplete? On average, participants listed 7.4 objectives for the MBA decision (the maximum number of listed objectives was 14). On average, 13.6 objectives on the master list were checked as relevant. For each participant, we divided these relevant objectives into those that were mapped back to that participant's list (*self-generated objectives*) and those that were only recognized as relevant after exposure to the master list (*recognized objectives*).²

Results clearly indicated that participants failed to generate a comprehensive list of objectives. On average, the number of recognized objectives was 7.6, a magnitude reliably greater than zero ($t(61) = 12.87$, $p < 0.01$). More surprisingly, the number of recognized objectives was actually greater than the number of self-generated objectives ($M = 6.0$, paired $t(61) = 2.03$, $p < 0.05$). The abundance of objectives that were only recognized at the second stage indicates a significant shortcoming of participants at the generation stage.

3.2.2. How Important Were the Recognized Objectives? The relative importance of recognized objectives was examined by looking at the "top 10" rankings provided by participants from the master list. Specifically, an investigation of the objectives ranked 1 to 5 by each participant revealed that, on average, 1.4 were recognized objectives. Broadening the set to those ranked 1 to 10, an average of 4.0 were recognized objectives.³ Therefore, although a majority of the objectives considered most important by participants were self-generated, a considerable portion of their top 10 objectives would not have been identified without use of the master list. Figure 2 presents the cumulative proportion of participants who failed to generate one of their top-ranked objectives. As illustrated by the figure, 10% of participants did not generate even their most important objective, and 71% of participants overlooked at least one of their top five objectives.

3.3. Discussion

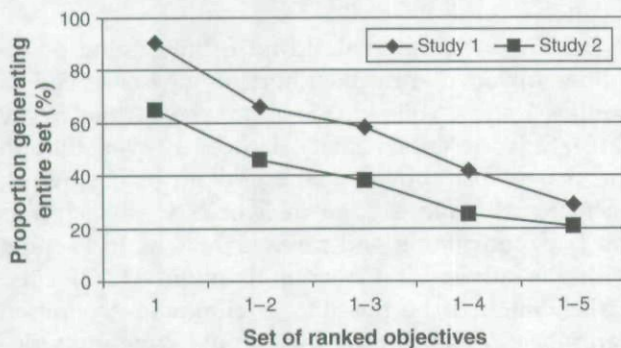
Study 1 presents a mixed picture regarding the capacity of individuals to generate objectives for a personally significant decision. On the one hand, participants produced a wide variety of decision objectives without any external aid, and they generated the majority of the objectives that they considered to be most important. On the other hand, the substantial number of recognized objectives indicates that numerous important objectives "slipped the minds" of participants during initial contemplation. Furthermore, participants acknowledged that many of these recognized objectives were in fact critical to the decision.

One conceivable explanation for the abundance of recognized objectives is that participants devoted inadequate effort to the generation task, generating too few objectives at the first stage and then recognizing many that they had missed. However, two arguments appear to contradict this account.

² As described above, the mapping procedure allowed participants to link one (or more) of their listed objectives to one (or more) objectives on the master list. Also, some listed objectives were not mapped to the master list. Thus, the number of self-generated objectives could be either more or less than the number of objectives originally listed.

³ Note that 4.0 is actually a conservative estimate. Fourteen of the 62 participants ranked fewer than 10 objectives (either because they did not check 10 items or they did not follow directions). Had everyone ranked 10 objectives, the average number of recognized objectives would have been higher.

Figure 2 Studies 1 and 2: Proportion of Participants Generating All Their Top-Ranked Objectives



Notes. The x-axis corresponds to the rankings provided by each participant: e.g., the data points above "1-4" indicate the proportion of participants generating their four top-ranked objectives without assistance. Study 2 incorporated the use of a one-week time delay, reducing the likelihood that participants would remember which objectives were self-generated.

First, if inadequate effort during the generation task was driving the difference in self-generated and recognized objectives, then participants who listed the fewest objectives at stage 1 should have recognized the most new objectives at stage 2, implying a negative correlation between the quantity of self-generated and recognized objectives. In fact, analysis revealed this correlation to be slightly positive ($r = 0.18$, ns). Second, the importance data reveal that almost all participants generated some objectives that they did not consider critical to the decision (i.e., objectives not included in their top 10), indicating that most participants took the generation task seriously enough to list both highly important and less-important objectives. Thus, we suggest that the failure to generate important objectives was not caused by inadequate effort or attention during the generation process, but rather by an incomplete cognitive representation of the decision.

Another possibility is that the number of self-generated objectives was understated because listed objectives that did not map to the master list were omitted from the analysis. We allowed this restriction to ensure that our analysis was performed on a uniform base of objectives for each participant. It is important to note that there were very few objectives of this nature ($M = 1.8$). In addition, we added any objectives that were "written in" by multiple participants in this study to the master list utilized in Study 2.

4. Study 2: Replication and Addition of a Time Delay

Study 2 was motivated by two concerns, the first of which was to replicate Study 1 using a different population. The second concern stemmed from a possible source of bias affecting importance ratings in the

previous study. Participants in Study 1 were asked to rank their most important objectives immediately after engaging in the mapping procedure. Out of a motivation to appear competent in their ability to generate objectives, therefore, participants may have overrated the importance of self-generated objectives. A similar possibility, based on self-inference, is that participants may simply have assumed that objectives generated independently were of particular importance. Alternatively, a demand effect may have occurred, whereby participants noticed the substantial number of recognized objectives and thus overrated their importance. We addressed these possible biases by introducing a time delay between the objective generation task and importance-ranking task. After the passage of time, participants should be less able to recall which of the objectives they had generated independently. Therefore, importance rankings should be relatively unaffected by the biases just described and more reflective of the objectives' actual underlying importance.

4.1. Methods

4.1.1. Participants and Design. Participants were 102 students in a weekend executive MBA program at a large university, and the study was administered as part of a class exercise. The design was similar to that of Study 1, with differences noted in the following section.

4.1.2. Materials and Procedure. In the first part of the study, participants were given a stimulus packet to assess their objectives for choosing an MBA program. The packet was similar to the one administered in Study 1, with two major exceptions. First, the number of objectives on the master list was expanded to 31 in order to incorporate popular "written-in" objectives from the first study ("diversity of students/faculty," "school ranking," and "alumni support"). Second, after mapping their listed objectives onto the master list, participants were told to circle only their top 10 objectives (but not to perform any ranking procedure).

Approximately one week later, each participant was sent an individualized e-mail. An introductory paragraph asked participants to "rank order the following objectives from most important to least important." Following these instructions, each participant saw a list of the 10 objectives (s)he had circled the prior week. Beside each objective was a line on which to assign its ranking. Participants simply replied to the e-mail to complete the study.

4.2. Results

At the first stage of the study, 17 participants failed to complete the master list, and 15 participants did not circle their top 10 objectives, leaving a usable sample of 70 respondents.

4.2.1. Was Self-Generation Incomplete? The average number of objectives initially listed was 5.9 (min = 2; max = 15). When presented with the master list, participants identified an average of 14.3 objectives as relevant. Of these, an average of 7.7 were recognized objectives, a magnitude clearly greater than zero ($t(69) = 12.25$, $p < 0.01$). More importantly, and consistent with results of the first study, the number of these recognized objectives was directionally greater than the number of self-generated objectives ($M = 6.6$, paired $t(69) = 1.44$, $p > 0.15$). Once again, participants had failed to list over half of the objectives that they later identified as personally relevant, supporting the contention that individuals encounter significant difficulty producing all the objectives that matter to them.

4.2.2. How Important Were the Recognized Objectives? As before, we examined the relative importance of recognized objectives in two ways. First, an average of 3.8 objectives within the top 10 were recognized objectives, corroborating the results of Study 1.⁴ Next, we examined the importance rankings assigned by participants at the second stage. As illustrated in Figure 2, 35% of participants did not generate even their most important objective, and 79% of participants included at least one recognized objective among those ranked 1–5. This inadequacy, observed after a one-week time delay, is even stronger than obtained in Study 1, so the prior result may be considered conservative (see the discussion of possible biases above). An alternative means of examining the importance of self-generated and recognized objectives is to compare their mean importance rankings. On average, recognized objectives in participants' top 10 sets were assigned a ranking of 5.8, a value only slightly lower than the 5.1 ranking assigned to self-generated objectives. Although this difference is reliable (paired $t(51) = 2.17$, $p < 0.04$),⁵ the two groups of objectives were strikingly similar in perceived importance.

4.3. Discussion

Study 2 replicated and extended the findings obtained in the prior study. Once again, results supported the notion that even for important, personally meaningful decisions, individuals do not generate a comprehensive list of objectives. Furthermore, the importance data indicated that objectives recognized only with external aid were, on average, nearly as important as those generated without assistance. Because these

rankings were obtained after a one-week delay, it seems unlikely that demand effects were driving our results.

The notion of external aid raises interesting possibilities for countering the shortcomings observed in Studies 1 and 2 (this idea is addressed further in the case study described later). However, even though the master list provided a useful means by which participants could expand their objectives, such aids are time consuming and rarely available. In keeping with the advice of Benjamin Franklin (1772) cited earlier, might it be possible to eliminate generation deficiencies more directly by simply asking individuals to consider their objectives repeatedly over time? Study 3 addressed this possibility in the context of an impending decision characterized by both short- and long-term consequences.

5. Study 3: Objectives for an Important, Prospective Choice

The first two studies illustrated that individuals tend to generate an inadequately narrow set of objectives, even when they have recently engaged in the focal decision. Although these studies provided a conservative test of our hypotheses, it would be fruitful to examine the same questions within an ongoing decision. Study 3 accomplished this goal by analyzing an important decision that was taking place at the time of data collection: the choice of a summer internship by first-year MBAs. Despite the obvious significance of this decision to their livelihood, results from Studies 1 and 2 raise the possibility that students will be unable to generate a comprehensive list of internship objectives. A second limitation of Studies 1 and 2 is that respondents listed their objectives during an experimental session. Although they were under no pressure to respond quickly, a time constraint was implied by the reality that the sooner they finished, the sooner they could leave. To ensure that time was not a limitation in Study 3, participants were given a week to generate their objectives, and they were specifically reminded of their task during the week.

5.1. Methods

5.1.1. Participants and Design. An e-mail recruitment was sent to approximately 250 first-year MBA students at a large southeastern university. The study was coordinated with the school's career center and was conducted during the spring term, at a time when most first-year students were interviewing for summer internships. To ensure that the decision was impending for all respondents, the recruitment specified that students who had already selected their

⁴ As before, this is a conservative estimate; 13 participants circled fewer than 10 objectives.

⁵ There are only 51 degrees of freedom because seven participants did not respond to the e-mail requesting importance rankings, and 11 individuals mapped all the objectives in their top 10.

internships were ineligible. Participants were 33 students who responded to this request. They were compensated for their participation with \$20 in cash and entry into a drawing for a \$100 gift certificate. In addition, all correspondence highlighted the potential value of participation as an aid to the decision process.

The design of Study 3 followed the same general format as Studies 1 and 2, with three notable exceptions. Most importantly, the focal decision in this study was the selection of a summer internship, and the master list included objectives relevant to this decision. Second, participants were explicitly requested to ruminate on their objectives over a one-week period before providing their initial list. Finally, in place of top 10 rankings, importance ratings were obtained for all of the objectives that were identified as relevant. Each of these features is described below.

5.1.2. Materials and Procedure. The study was administered through a combination of e-mail and paper and pencil over the course of two weeks. At the first stage, students responding to the recruitment e-mail were sent a reply with instructions for eliciting their objectives. After reading a general overview of the study, participants were informed that "We would like you to think periodically over the week about how you would answer the following question: What objectives are relevant to you when selecting from potential opportunities for your summer internship?" They were specifically instructed to think about the problem multiple times during the week and were encouraged to consult with others who might have relevant input. Approximately three days later, participants were sent an e-mail reminding them to ruminate on their objectives repeatedly throughout the week.

At the end of the week, participants were asked to send (via e-mail) the lists of objectives that they had generated. Subsequently, a master list of 29 objectives for the internship decision was created by the authors, utilizing a combination of participant responses, our individual analyses of the decision problem, and mutual discussion. Table 1 illustrates the resulting list; sample items include "enhances my knowledge in a particular industry," "allows me to meet interesting people," and "provides a structured program for learning and training."

Individualized packets were created and e-mailed to each participant approximately three days after the elicitation procedure; participants printed these packets and completed them by hand. The individualized packets contained three pages: (1) the master list of objectives, (2) the participant's self-generated list, and (3) a few follow-up questions. As in the prior studies, participants were asked to place a checkmark next to all objectives on the master list that they

Table 1 Master List of Summer Internship Objectives for Study 3

I would like to choose an internship that...

Improves my attractiveness for full-time job offers
Helps me make good networking contacts
Gives me pride from landing a prestigious internship
Is at a company that sponsors work visas for placement in U.S. offices
Helps me develop my leadership skills
Provides information to help select a job after graduation
Provides opportunities to interact with senior managers
Provides a structured program for learning and training
Uses skills I have learned in my first year of B-school
Allows me to meet interesting people
Is with a company whose culture I identify with
Is challenging
Helps me decide what courses and skills I need to develop next year
Is enjoyable to do
Could lead to a full-time offer from that firm
Helps me improve my communication skills
Is a job that I would like to do full-time after graduation
Compensates me well
Gives me a substantial project of which I can feel ownership
Allows me to experience a new geographical area
Helps me decide whether the internship field is good for me long term
Provides flexibility for personal interests during the summer
Enhances my resume
Is with an organization that I am passionate about
Lets me work with a diverse group of people
Enhances my knowledge in a particular industry
Offers the chance to learn new skills
Is at a well-recognized/respected company
Is in a specific location (e.g., near family or friends)

deemed important in selecting a summer internship, and they were then asked to map the objectives on their self-generated lists to those on the master list. After completing the packet, participants returned it to a drop-off location at their convenience.

Approximately five days later, the importance assessment was conducted via e-mail. Each participant received a personalized message with a list of all the objectives (s)he had identified as relevant. In addition, each list included a "bogus objective" and a "baseline objective." The bogus objective, "allows me to mentor other employees," was included to ensure that participants were devoting adequate attention. This objective was not a part of the master list, nor was it mentioned by any participant during the generation task; thus, we expected that diligent participants would assign it little importance. The baseline objective, "helps me improve my communication skills," was an objective from the master list checked by only five participants. As a reasonable objective that was not unanimously checked, the baseline provides a benchmark for interpreting the importance assigned to other objectives.

In contrast to Studies 1 and 2, the importance assessment of Study 3 utilized a rating procedure. Instructions in the e-mail asked participants to "rate each objective in terms of its importance to you."

Beside each objective on their personalized list, participants indicated the importance of that objective by typing a number from 1 (not important) to 9 (extremely important). After determining their ratings, participants replied to the e-mail to finish the study.

5.2. Results

Five individuals failed to complete the entire study, leaving a sample size of 28. Responses to the follow-up questions did not correlate with the key results and will not be discussed further.

5.2.1. Was Self-Generation Incomplete? As before, the number of *self-generated* and *recognized* objectives was compared in order to measure the comprehensiveness of participants in generating the objectives that mattered to them. Upon exposure to the master list, participants identified an average of 14.8 objectives as important. Of these checked objectives, participants mapped an average of 6.8 back to their self-generated lists, but an average of 7.9 were only recognized as relevant after being encountered on the master list. Therefore, participants failed to generate over half of the objectives that they later acknowledged to be relevant (paired $t(27) = 1.24$, $p > 0.20$). It is important to remember that unlike Studies 1 and 2, participants were actively engaged in the focal decision outside of the study, and they were explicitly encouraged to deliberate on their objectives over a one-week period. Nonetheless, results converged with those of the prior studies, revealing a remarkable inadequacy on the part of participants facing this real-world decision.

5.2.2. How Important Were the Recognized Objectives? The relative importance of self-generated and recognized objectives was compared by examining ratings provided at the final stage. As mentioned above, a "bogus objective" and a "baseline objective" were also included in the ratings assessment. The average importance rating of the bogus objective was only 2.9 on the nine-point scale, confirming that participants were adequately deliberating on their responses. The average rating of the baseline objective, which provides a standard of comparison for evaluating the importance of other objectives, was 4.7 (paired $t(27) = 5.82$, $p < 0.01$).

On average, participants assigned a rating of 6.8 to their set of recognized objectives, indicating that many of the objectives they had left out during the generation stage were in fact critical to their internship decision. In contrast, the average rating assigned to self-generated objectives was 7.4. Although the difference in importance of recognized and self-generated objectives was statistically reliable (paired $t(27) = 3.98$, $p < 0.01$), it is remarkably small in absolute terms. This difference appears especially minor

in light of the fact that the average importance of recognized objectives was nearly two points higher than that of the baseline objective (paired $t(27) = 6.81$, $p < 0.01$). Taken together, therefore, the results indicate that participants failed to consider numerous objectives that they later acknowledged to be vital for the evaluation of alternatives.

5.3. Discussion

Replicating the inadequacies observed in the first two studies, Study 3 extended their findings to a prospective decision setting. The results are particularly remarkable when one considers the real-world significance of the domain we examined. Participants were MBA students actually deciding between summer internships, a decision of undeniable short- and long-term consequence. Even so, and even with extended opportunities to engage in the generation process, the students produced a strikingly deficient list of objectives for their internship decision. Furthermore, the missed objectives were nearly as important as the ones that were evoked independently. It is alarming to consider the downstream effects of such inadequacy on actual choices; on the other hand, we are optimistic that the student participants may have benefited from their involvement.

In interpreting the results of all three experiments, it is useful to consider a potential bias brought about by the mapping procedure. After receiving the master list of objectives, each participant checked those that were relevant and then matched their original objectives to those they had checked. In the reported analyses, we accepted all of these mappings at face value. However, examination of the packets in each study suggests that participants were somewhat liberal in matching their independently generated objectives to those on the master list. In Study 3, for example, one participant mapped their listed objective "has a multicultural and multidisciplinary group of coworkers" to the master-list objective "provides opportunities to interact with senior managers." Another participant mapped the listed objective "provides an enjoyable environment" to "helps me make good networking contacts," and a third mapped "has a good chance of providing a full-time offer" to "helps me decide whether the internship field is good for me." In general, therefore, it is our view that the mapping procedure overstates the completeness of participants' self-generated lists. If we eliminated only the most questionable matches in the present studies, our results would be even stronger.

6. Case Study: Objectives in a Real-World Problem

Experiments about decision making naturally focus on specific aspects that are only a part of any real-world decision process. Although such a restriction is

necessary to carefully study the elements of the situation, it is the relevance of insights from the experiments to real-world decisions that ultimately matters. The issue is whether, at one extreme, the empirical results are only artifacts of the studies themselves, or, at the other extreme, they capture a fundamental weakness in real-world decision making and suggest improvements.

Fortunately, we possess the means of examining our ideas in a real-world, highly significant business decision. In 1993, Seagate Technology was the largest manufacturer of disc drives and components in the world, with annual revenues of \$3 billion. In August of that year, Al Shugart, chairman of the board and CEO of Seagate, announced his vision for the firm through the end of the century. This vision included the creation of a software company via the acquisition of companies offering products for information, network, and storage management. After Seagate had purchased approximately a dozen firms, it needed to integrate these firms into a single organization with a common mission and objectives. Steve Luczo, who had been hired from Bear Stearns to build Seagate Software, began working with one of this paper's authors (Keeney) in 1995. Their intention was to develop a comprehensive set of objectives that clearly stated Seagate Software's mission.

Initially, Keeney held individual discussions with 12 executives, including Shugart, Luczo, and several previous owners of the acquired firms who now held top management positions at the new company, in order to identify objectives of Seagate Software. Discussions lasted about one hour each and involved a value-focused probing procedure described in Keeney (1992).

At the beginning of the discussion, each individual was asked to discuss any objectives, hopes, aspirations, desires, or plans (s)he had for the new organization. As each individual spoke, Keeney created a list of any statements that indicated an objective. Subsequently, the discussion was guided into potential categories of objectives that may have been only lightly covered. For example, one question asked at this stage was, "What are Seagate Software's objectives for its customers?" Additional categories concerned objectives for other stakeholders, including Seagate Technology, shareholders, and society. Another important category was personal objectives of the respondents themselves, brought into focus with the question "What professional or personal objectives do you hope to achieve by being part of Seagate Software?" Additional details of this assessment process are outlined in Keeney (1999).

After these discussions, each individual was sent a written list of the objectives (s)he had articulated and a preliminary means-end objectives network relating

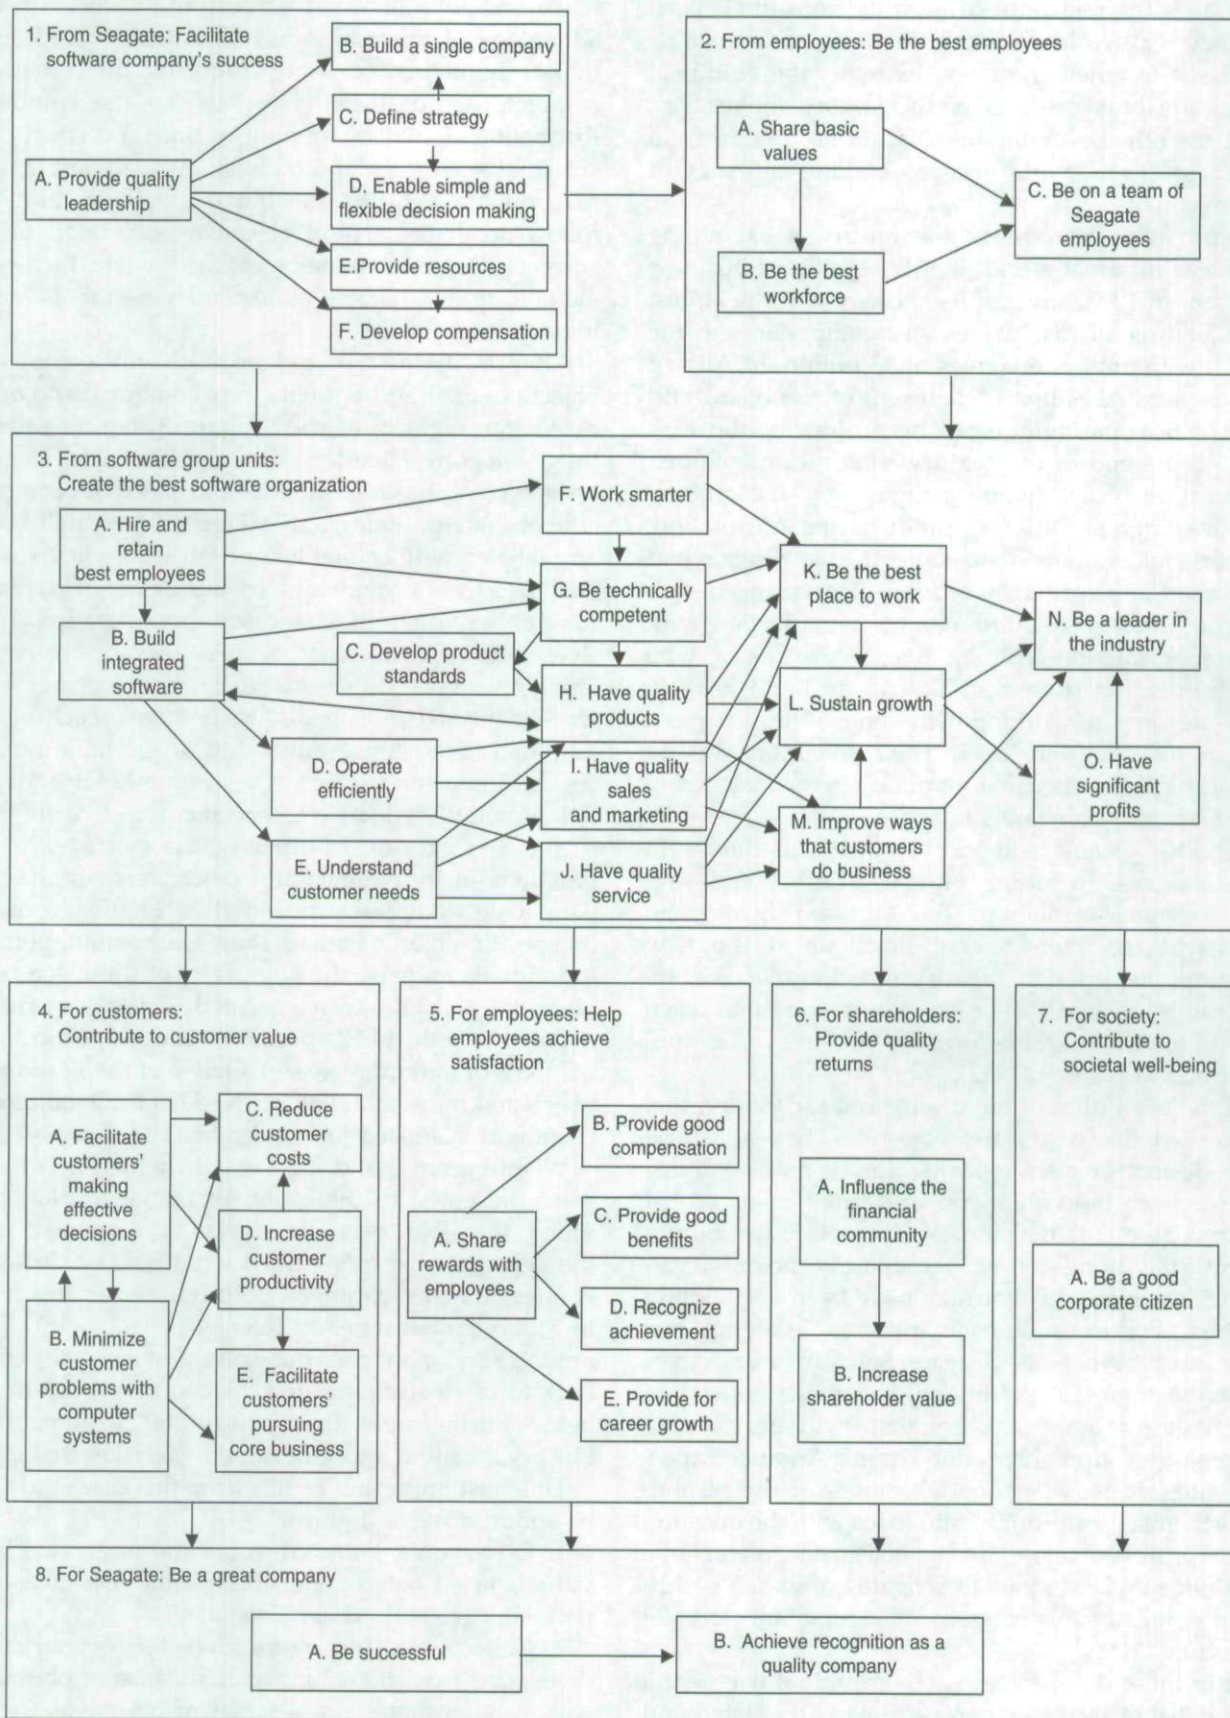
all of the objectives. Respondents were encouraged to make additions and modifications to this list. A major advantage of this real-world elicitation procedure is that respondents: (a) were motivated to provide a complete set of their objectives for the company, (b) spent a significant amount of time and effort constructing such a set, and (c) were allowed ample time to ruminate between the initial elicitation and the follow-up. If the sets of objectives generated under such conditions are nonetheless incomplete, then such factors are less tenable as explanations for the findings of Studies 1–3.

After the discussions and feedback were completed, objectives of all respondents were combined and organized into eight categories; Figure 3 depicts each of these categories as a large box. Within each category of objectives, the respondents had listed several specific objectives, illustrated in Figure 3 by small boxes and labeled with capital letters. For example, the specific objective "provide quality leadership" is referred to as objective 1A. In most cases, more-detailed objectives were given for each of these specific objectives. Table 2 summarizes the categories and specific objectives identified during discussions with each of the 12 individuals, and examination of the table reveals several interesting results. On average, each individual provided at least one specific objective for 5.16 of the 8 categories. In other words, out of 96 combinations of individuals and objective categories, 62 were assigned at least one objective. Figure 3 contains 39 specific objectives, and thus 468 combinations of individuals and specific objectives; of these 468 combinations, only 170 were specified. This represents an average of only 14.17 specific objectives per individual (36% of the complete set). Only 1 of the 12 individuals listed more than half of the 39 specific objectives (2 individuals listed just under half).

When Figure 3 and the logic on which it is based were presented to the group of managers that provided the objectives, there was a consensus that the aggregate list represented what Seagate Software wanted to achieve. Steve Luczo referred to this work in a subsequent letter to Keeney (1999, p. 15), "The organization and communication of our objectives back to us clearly indicated the common purpose we held." Furthermore, this "contributed to integrating our organization and guiding our decision making."

The most important results from this case study can be summarized as follows:

1. On average, knowledgeable and motivated individuals listed only 36% of the specific objectives that they later recognized as relevant.
2. Collectively, the aggregated list of objectives developed from the 12 individuals' lists of objectives was a much more complete set of objectives for the business situation.

Figure 3 Objectives of Seagate Software Illustrated in a Means-End Objectives Network

Notes. (Adapted from Kenney 1999.) An arrow between two objectives indicates that achieving the former influences achievement of the latter.

Table 2 Objectives of Seagate Mentioned by Individual Leaders in the Organization

Objective	Individual												Individuals with objective	Individuals in category
	1	2	3	4	5	6	7	8	9	10	11	12		
1				1									1	11
1A			1	1		1							3	
1B	1				1	1				1	1		5	
1C	2			1	1		4	1		1			6	
1D		2	1	2	1					1			5	
1E		3	1	2					1				4	
1F			1		2			1	1		2		5	
2														8
2A				1						1			2	
2B				1		2	1	2	2			1	6	
2C										1	1		2	
3														12
3A		1	2					1	1	1		2	6	
3B	3	1			2	1	2	1	3	1		1	9	
3C			1					1					2	
3D	2	3	3							1		1	5	
3E	1		2		1	2	4			1			6	
3F			1			1	1	1	1	2			6	
3G		1								1	1		3	
3H			2		1		1	1	2		2	1	7	
3I			1			1	2	1				1	5	
3J				1	2	1							3	
3K		1		2	3	2		4		2			6	
3L	1		1	2	2		1	1		2	1	1	9	
3M		1		2		1						1	4	
3N		3	5			1	2	4	1			1	7	
3O	1		1	1			1			2	1	1	7	
4	1								1				2	6
4A			1	1									2	
4B			1			1							2	
4C			1	1						1			3	
4D			2							1			2	
4E						1							1	
5		3	1			2			2			1	5	9
5A			1		1	1		1	1				5	
5B			1			1				1			3	
5C		1			2								2	
5D										1			1	
5E			1	1		1		1				1	5	
6	1	1											2	7
6A		1	1			1							3	
6B	1	1	1	1					2			1	6	
7										1			1	1
7A										3			1	
8				1					1				2	8
8A	1	1		2		1			2			1	6	
8B		1		1		1		1			1		5	
Objectives 39	9	14	23	17	12	19	10	15	11	19	8	13	170	

The first result is consistent with results of the laboratory studies discussed earlier. The second result is also consistent if one interprets the aggregated list above as analogous to the "master list" of the earlier studies.

It is useful to consider implications of this real-world demonstration. The basic circumstances were

weighted against finding results consistent with Studies 1 through 3. Each of the 12 individuals surveyed knew the purpose of the assessment task, and each was motivated and knowledgeable. They had more than sufficient time to respond, were prompted verbally with cues to help identify objectives, and were explicitly asked for modifications or additions to their

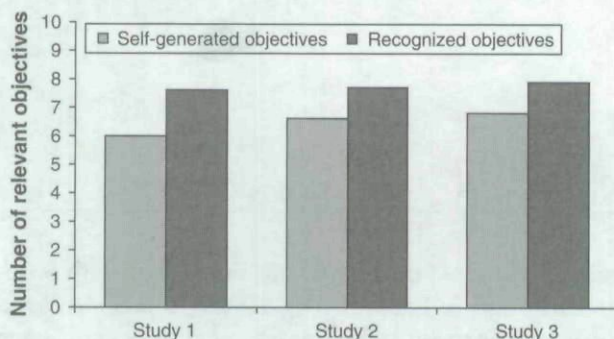
initial lists. They were also encouraged to discuss their original lists of objectives with others in order to identify missing items. Even in such a favorable setting, the observed inadequacies were consistent with the findings of our empirical studies. It seems clear that specifying a reasonably complete set of objectives is a formidable task for an unaided individual.

7. Discussion

Taken together, the data presented in this article indicate that individuals may commonly undertake important decisions without considering many personally relevant objectives. That is, even when decision makers think explicitly about what matters to them, it is likely that they will omit many of their most important objectives. Figure 4 summarizes the results of the empirical studies in terms of the quantity of self-generated and recognized objectives. In each case, the two sets were similar in magnitude, implying that participants were remarkably incomprehensive during the generation process. Even more notably, the importance data for each study reveal that individuals were not merely “leaving out” objectives that were in fact trivial to them. Instead, many of the recognized objectives were acknowledged to be critical to the decision.

One alternative explanation for the present findings is that participants accurately reported the decision objectives that mattered to them at the generation stage, but then checked a number of additional objectives from the master list in order to appear thorough. This possibility seems improbable given the anonymity of participants in all three studies. Furthermore, a demand explanation is problematic in accounting for the importance data, because it seems unlikely that objectives added “for appearance sake” would have been rated highly important.

Figure 4 Quantity of Self-Generated vs. Recognized Objectives in Studies 1–3



Notes. Bars represent the average number of checked objectives on the master list that were either mapped to a participant's own list (*self-generated objectives*) or left unmapped (*recognized objectives*). Studies 1 and 2 involved the choice of an MBA program; Study 3 involved the choice of a summer internship.

We believe that our findings are consistent with a wide body of evidence that decision makers respond to complexity by simplifying their environment. In particular, our data support the notion that during the contemplation of a decision, individuals become mired in an overly narrow mental representation that prevents them from considering a comprehensive set of objectives. Instead, they generate only the objectives that are cued by this incomplete representation, and a variety of uncued (but substantially important) objectives are not evoked. One implication of this pervasive failure is that decision making will be guided by whichever incomplete set of objectives is made salient by the simplified frame that is adopted. Another implication is that the generation process can be improved by helping decision makers to adopt a broader construal and approach the task from multiple perspectives. Study 3 presented evidence that merely considering the decision at multiple points in time is insufficient to redress the problem. Instead, it seems likely that external assistance is required in order to stimulate the articulation of objectives that would otherwise be overlooked. Further research will be useful to identify the specific cognitive mechanisms underlying our results and explore alternative means of ameliorating this shortcoming.

7.1. Limitations

The decisions in Studies 1 and 2 were retrospective. As mentioned previously, we believe that this feature enabled a more conservative test of our hypotheses, because current MBA students presumably possess some expertise in the decision domain of school selection. However, it is possible that having recently engaged in the focal decision may have caused participants to respond differently to the objectives generation task, in that objectives reported retrospectively may not match those that would have been reported prior to the decision. We offer four responses. Most importantly, Study 3 focused on an ongoing decision that is not subject to this criticism, and the results were remarkably consistent with those of the prior studies. Second, if participants in Studies 1 and 2 were biased by having already made their MBA decision (e.g., by wanting to defend the decision they made in real life), then the data may present an inaccurate depiction of their underlying values, but it is unclear how this bias would explain the observed differences in self-generated and recognized objectives. Third, the Seagate case study addressed objectives for a complete range of future decisions; even with motivation, knowledge, and assistance in generating objectives, the results in this real-world application were consistent with those of our lab-based studies. Fourth, many consequential decisions are made repeatedly (budget

allocations, personnel hiring, etc.), and thus it is noteworthy that the failures we document are not eliminated by prior experience with the decision.

This research offers no direct evidence that the inadequacies observed in our studies (or the Seagate application) would degrade the quality of eventual decisions. Our focus was on the objectives generation process itself, and we believe these studies provide valuable evidence regarding the inadequacy of unaided generation. Nonetheless, given the widespread belief that consideration of one's objectives is beneficial to effective decision making (Keeney and Raiffa 1976), an obvious implication of our findings is that decisions will be poorer as a consequence of the failure to adequately generate objectives. Research investigating this possibility would be a natural extension of our work. For example, investigators might present a hypothetical decision scenario and examine the extent to which assistance during the objectives generation stage reduces the incidence of normative violations.

7.2. Implications for Decision Models

A frequent criticism of decision modeling is that analysts tend to utilize objective functions that are incomplete representations of a decision maker's values. The present research can be viewed as supporting this criticism, but in a way that has not previously been addressed. That is, our findings imply that choice models based on dimensions an agent has reported to be important will frequently suffer from omitted variable bias, because the agent will have overlooked several factors of genuine importance to the decision. In the examples presented in this paper, decision makers appear to have overlooked both fundamental objectives (hampering the evaluation of alternatives) and means objectives (hampering the creation of alternatives). Therefore, our results suggest that unaided listing is an insufficient method for eliciting a complete set of objectives, and analysts must be creative in finding procedures to facilitate the task. Exploiting the distinctions between means and fundamental objectives may be fruitful in this endeavor.

On a more positive note, our studies also suggest what these procedures might look like. Participants in the empirical studies benefited greatly from a "master list" that provided an abundance of potential decision objectives. The master lists were initially created by compiling potential objectives elicited from a number of individuals. As such, the value of the master list in these studies can be viewed in light of the established benefits of aggregating group responses in complex decision tasks (Kerr and Tindale 2004): e.g., the use of multiple experts in forecasting (Clemen 1989) or managerial problem solving (Rulke and Galaskiewicz 2000). The Seagate case study provides

a compelling demonstration that strategic firm objectives can be more fully captured by combining the responses of numerous individuals within the company. More broadly, our results support the notion that helping individuals to consider their underlying objectives should be a focus of prescriptive analysis (e.g., Baron 1997).

We close by noting that there are numerous real-world opportunities for individual decision makers to uncover their objectives. In consumer decision making, for example, persons may benefit from independent guides like *Consumer Reports*, which describe the features of the products they review. By utilizing these guides, diligent consumers may recognize important personal objectives that they had not generated on their own. Second, consultation with others making the same decision (e.g., other individuals applying for an MBA) may be helpful in recognizing one's objectives. Similar benefits may be obtained by asking friends or acquaintances with intimate knowledge of the decision maker for objectives they believe to be most appropriate. Third, the opportunity sometimes exists for a trial period in which one may "audition" alternatives at limited cost (e.g., temporary employees, money-back guarantees, rent-to-own housing). In addition to conveying information about attributes of the auditioned alternative, these trial periods provide decision makers with the opportunity to confront their objectives experientially by reflecting on factors that contribute to satisfaction with the choice. To the extent that any of these methods improve our ability to "articulate what we want," they represent valuable tools for overcoming inadequacies in the generation process.

Acknowledgments

The authors thank James Bettman, Bob Clemen, Robin Gregory, Jay Russo, Jim Smith, and Bob Winkler for valuable feedback on earlier versions of this paper.

References

- Baron, J. 1997. Confusion of relative and absolute risk in valuation. *J. Risk Uncertainty* 14 301–309.
- Barrett, R. 2006. *Building a Values-Driven Organization*. Elsevier, Burlington, MA.
- Beach, L. R. 1990. *Image Theory: Decision Making in Personal and Organizational Contexts*. Wiley, New York.
- Bettman, J. R., M. F. Luce, J. W. Payne. 1998. Constructive consumer choice processes. *J. Consumer Res.* 25(3) 187–217.
- Clemen, R. T. 1989. Combining forecasts: A review and annotated bibliography (with discussion). *Internat. J. Forecasting* 5 559–583.
- Drucker, P. F. 1954. *The Practice of Management*. Harper, New York.
- Fischhoff, B. 1975. Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *J. Experiment. Psych.: Human Perception Performance* 1(3) 288–299.
- Fischhoff, B., P. Slovic, S. Lichtenstein. 1978. Fault trees: Sensitivity of assessed failure probabilities to problem representation. *J. Experiment. Psych.: Human Perception Performance* 4(2) 330–344.

- Fox, C. R., R. T. Clemen. 2005. Subjective probability assessment in decision analysis: Partition dependence and bias toward the ignorance prior. *Management Sci.* 51(9) 1417-1432.
- Franklin, B. 1772. *Letter to Joseph Priestly*. (Reprinted in W. Isaacson, ed. 2005. *A Benjamin Franklin Reader*. Simon & Schuster, New York.)
- Gilbride, T. J., G. M. Allenby. 2004. A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Sci.* 23(3) 391-406.
- Kahneman, D., P. Slovic, A. Tversky. 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, New York.
- Keeney, R. L. 1992. *Value-Focused Thinking: A Path to Creative Decisionmaking*. Harvard University Press, Cambridge, MA.
- Keeney, R. L. 1999. Developing of a foundation for strategy at Seagate software. *Interfaces* 29(6) 4-15.
- Keeney, R. L., H. Raiffa. 1976. *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*. Wiley, New York. (Reprinted by Cambridge University Press, 1993.)
- Kerr, N. L., R. S. Tindale. 2004. Group performance and decision making. *Ann. Rev. Psych.* 55 623-655.
- Leon, O. G. 1999. Value-focused thinking versus alternative-focused thinking: Effects on generation of objectives. *Organ. Behav. Human Decision Processes* 80(3) 213-227.
- Lussier, D. A., R. W. Olshavsky. 1979. Task complexity in brand choice. *J. Consumer Res.* 6(2) 154-165.
- Nietzsche, F. 1879. *The Wanderer and His Shadow*. (Reprinted in G. Colli, M. Montinari, eds. 1980. *Friedrich Nietzsche, Sämtliche Werke: Kritische Studienausgabe*. de Gruyter, Berlin.)
- Nutt, P. C. 1998. Evaluating complex strategic choices. *Management Sci.* 44(8) 1148-1166.
- Payne, J. W., J. R. Bettman, E. J. Johnson. 1993. *The Adaptive Decision Maker*. Cambridge University Press, Cambridge, MA.
- Payne, J. W., J. R. Bettman, D. A. Schkade. 1999. Measuring constructed preference: Towards a building code. *J. Risk Uncertainty* 19(3) 243-270.
- Peters, T., R. Waterman. 1982. *In Search of Excellence*. Harper & Row, New York.
- Rulke, D. L., J. Galaskiewicz. 2000. Distribution of knowledge, group network structure and group performance. *Management Sci.* 46(5) 612-625.
- Russo, J. E., K. Kolzow. 1994. Where is the fault in fault trees? *J. Experiment. Psych.: Human Perception Performance* 20(1) 17-32.
- Simon, H. A. 1955. A behavioral model of rational choice. *Quart. J. Econom.* 69 99-115.
- Tversky, A., P. Slovic, S. Sattah. 1988. Contingent weighting in judgment and choice. *Psych. Rev.* 95(3) 371-384.

Copyright 2008, by INFORMS, all rights reserved. Copyright of *Management Science* is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.