

User Perceptions of Decision Support Effectiveness: Two Production Planning Experiments*

Fred D. Davis

Carlson School of Management, University of Minnesota, Minneapolis, MN 55455

Jeffrey E. Kottemann

School of Business Administration, University of Michigan, Ann Arbor, MI 48109

ABSTRACT

Decision support systems continue to be very popular in business, despite mixed research evidence as to their effectiveness. We hypothesize that what-if analysis, a prominent feature of most decision support systems, creates an "illusion of control" causing users to overestimate its effectiveness. Two experiments involving a production planning task are reported which examine decision makers' perceptions of the effectiveness of what-if analysis relative to the alternatives of unaided decision making, and quantitative decision rules. Experiment 1 found that almost all subjects believed what-if analysis was superior to unaided decision making, although using what-if analysis had no significant effect on performance. Experiment 2 found that decision makers were indifferent between what-if analysis and a quantitative decision rule which, if used, would have led to significant cost savings. Thus, what-if analysis did create an illusion of control: decision makers perceived performance differences where none existed, and did not detect large differences when they were present. In both experiments, decision makers exhibited difficulty realizing that their positive beliefs about what-if analysis were exaggerated. Such misjudgments could lead people to continue using what-if analysis even when it is not beneficial and to avoid potentially superior decision support technologies.

Subject Areas: Decision Support Systems and Production/Operations Management.

INTRODUCTION

What-if analysis has long been among the most widely used capabilities of computer-based decision support systems [4] [31] [47]. What-if analysis is a method for manipulating a quantitative model of a business situation in which decision makers specify alternative values of decision variables and environmental assumptions, and the computer solves the model and displays predicted results. Numerous computer programs, including IFPS and Lotus 1-2-3, are available for creating business models and interrogating them via what-if analysis (see [48]).

Despite this popularity, research to date on the effectiveness of what-if analysis, and decision support systems in general, has produced mixed results (see [7] and

*The order of authorship is alphabetical; authors contributed equally to the research. The authors wish to acknowledge the helpful comments of Ed Miller, Janet Sniezek, Dan Stone, Paul Zarnoff, as well as the anonymous reviewers and associate editor.

[46]). Researchers have often been unable to find any performance differences between assisted and unassisted decision making (e.g., [3] [21] [23]). Significant DSS effects, when found, have sometimes been positive [6] [8] [22] [40] [46], and sometimes negative [10] [35] [36].

Why is what-if analysis so popular when research has been unable to demonstrate consistent advantages? Do decision makers overestimate the effectiveness of what-if analysis? As discussed in [1], and demonstrated in the context of information reporting systems [15], real and perceived performance impacts of decision support may differ. This research examines whether what-if analysis gives people the illusion that it helps their decision performance more than it really does. Such an illusion would have troubling implications in light of research linking user acceptance of computer tools to users' beliefs about how much the tools improve performance [11] [12]. If users overestimate their decision performance using what-if analysis, they may unintentionally degrade their own performance by using ineffective tools or avoiding alternative tools that may well be more effective than what-if analysis.

What-if Analysis and the Illusion of Control

We hypothesize that what-if analysis leads to an illusion of control, defined by Langer [37] as a person's expectation of success on a task that is inappropriately higher than objective circumstances warrant. In general, an illusion of control tends to occur when factors ordinarily associated with increased performance in skill situations, including "exerting effort while actively engaged in the task" [37, p. 313], are introduced into situations at least partially determined by chance. Other research supports the effort-confidence link [42] [50]. By allowing decision makers to manipulate decision variables and assumptions in a model, and observe their effects on predicted outcomes, what-if analysis increases the degree of active involvement in the decision task relative to unaided decision making. As Langer points out, "In skill situations, the more one actively participates in the event, the more control one has over the outcome. Hence it is predicted that the greater one's active involvement in a chance event, the greater the illusion of control" [37, p. 318]. Langer and Roth [39] also found that active involvement in a purely chance task inflated subjects' sense of control over outcomes. Although these experiments were done in situations where outcomes are completely determined by chance, such as lotteries and coin flips, Langer argues that the illusion of control should be even greater in situations that are partially controllable due to the increased plausibility of attributing performance to skill factors [37].

Forecasting and planning activities fit the kind of situation in which Langer would predict the occurrence of the illusion of control [27]. That is, decision makers generally have only partial control over the environment, but engage in cognitive activity with the purpose of increasing control. The present research is concerned with the broad category of routine decision making under uncertainty, specifically examining a simulated multi-period production scheduling task. Outcomes in such decision-making environments are partially determined by chance. Therefore, we expect the additional cognitive effort associated with what-if analysis to generate an illusion of control. In Experiment 1, we compare what-if analysis to unaided decision making, hypothesizing that people will systematically overestimate the advantage of what-if analysis.

What-if Analysis Versus Quantitative Decision Rules

If what-if analysis creates an illusion of control, as we suggest, then it may cause people to overestimate the effectiveness of what-if relative to non-interactive, quantitative decision rules. Quantitative decision rules are an alternate form of decision support, quite different in spirit from what-if analysis. For several decades, it has been shown that a wide variety of such rules are generally superior to unaided decision making (for an extensive review, see [32]). Examples include linear statistical models [13] [14] [41], bootstrap models [9], linear decision rules [29], unit weighting schemes [13] [18], and simple mechanical decision criteria [5] [44], which have proven effective in such diverse contexts as medical diagnosis, graduate admissions, auditing, marketing, law, finance, and production planning.

A major reason for the effectiveness of decision rules is that they reduce the inconsistency present in unaided decision making, while bringing relevant information to bear on the decision [9] [13] [32]. Despite their proven effectiveness, however, decision makers have been reluctant to use decision rules. One explanation for this reluctance, given by Dawes [13], is cognitive conceit: the illusion that the environment is more predictable than it really is and that greater cognitive effort will lead to better predictions than those afforded by linear models known to be imperfect. This explanation parallels the illusion of control concept described by Langer [37] by linking active involvement (effort) to perceptions of improved performance.

How will people judge the effectiveness of decision rules in comparison to what-if analysis? Since decision rules generate recommendations via nonjudgmental, mathematical, statistical, or mechanical means [32], their use involves less active involvement than does what-if analysis, putting the decision maker in a more passive role. In Experiment 2, we hypothesize that, due to the illusion of control created by what-if analysis, people will overestimate the effectiveness of what-if relative to a quantitative decision rule which, if used, would generate significant performance advantages.

One known determinant of user acceptance of decision rules is the way they are described to the decision maker. Arkes, Dawes, and Christensen [5] found that people became more likely to use a decision rule if they were warned beforehand that trying to beat the rule would lead to worse performance. Similarly, Powell [44] found that increasing the reported accuracy of a decision rule (50 percent versus 70 percent versus 90 percent) had a significant effect on peoples' willingness to use it. In Experiment 2, we present the decision rule in two different ways: neutral versus 90 percent accuracy.

Learning from Experience

Will people learn from experience? If the illusion of control phenomenon is temporary, disappearing as decision makers gain experience in the task environment, the resulting problems of dysfunctional adoption (or rejection) of decision tools would correct themselves over time. However, there are theoretical reasons to expect that the illusion of control will not diminish with experience. As Einhorn and Hogarth [19, p. 395] state, "in real-world situations, judgments are made for the purpose of choosing between actions. This means that outcome information, which is available only after actions are taken, is frequently the only source of feedback with

which to compare judgment” (see also [17] [20]). In the present context, people judge the effectiveness of what-if analysis as a basis for deciding whether or not to use it in preference to alternatives such as unaided decision making or quantitative decision rules. Operating under the illusion that what-if is superior to its alternatives, decision makers are apt to use what-if analysis to the exclusion of its alternatives, thus failing to generate the feedback needed to disconfirm their erroneous beliefs.

We investigate both the typical situation, in which decision makers have only the evidence generated by their own actions as a basis for confirming or disconfirming their beliefs, and the less typical situation, in which outcome feedback regarding actions not taken is provided explicitly. In both Experiment 1 and 2 we expect that, when people are not provided outcome feedback, their inflated beliefs about what-if will persist over time. In both experiments, we also examine whether introducing explicit outcome feedback that would ordinarily not be available to decision makers will mitigate the illusion of control surrounding what-if analysis.

EXPERIMENT 1: WHAT-IF ANALYSIS VERSUS UNAIDED DECISION MAKING

Experiment 1 was designed to approximate the real life situation in which a decision maker accustomed to making unaided decisions is offered the use of what-if analysis on a trial basis, and then forms an opinion about the effectiveness of what-if. Subjects made unaided decisions within a 24-period production planning simulation. After witnessing their decision performance unaided, subjects were given an interactive, what-if model for four periods of trial use. After trial use, subjects were asked their beliefs about their performance with, versus without, the use of what-if analysis. Next, subjects used what-if for a full 24-period simulation, and then were asked again how effective they think it is relative to unaided decision making. The first round of questioning occurred before outcome feedback (results of actually using what-if over 24-periods), the second round after outcome feedback. To rule out order effects, a second group received the treatments in reverse order: first they used the what-if model for the first 24-periods; tried four practice periods unaided; reported their beliefs about the effectiveness of what-if (before outcome feedback); completed 24 periods unaided, and again reported their beliefs (after outcome feedback).

- H1: Based on prior research using a similar task, we expect no performance advantages due to the use of what-if analysis compared to unaided decision making [35] [36].
- H2: However, due to the illusion of control, we expect subjects to believe that what-if analysis is superior to unaided decision making.
- H3: We expect that outcome feedback (i.e., seeing the results of using what-if analysis versus unaided decision making over 24-periods) will significantly reduce subjects' beliefs in the superiority of what-if analysis.

Method

Task. The experimental task is based on the Holt, Modigliani, Muth, and Simon (HMMS) model of the production planning problem (see [28] [29] [35] [36]). In

this model, the decision maker faces future demand and must decide upon a production level and a workforce level over a series of periods. The objective is to minimize cumulative total costs. The cost function for a given new period is the sum of three quadratic cost components [28]:

$$\text{Workforce level change cost} = 64.3 (\text{current workforce} - \text{new workforce})^2,$$

$$\text{Worker overtime/idle time cost} = .8 (\text{new workforce} \cdot 5.67 - \text{new production})^2,$$

$$\text{Cost for nonoptimal inventory} = .02 (\text{current inventory} + \text{new production} \\ - \text{new demand} - 320)^2,$$

Total cost is the sum of the above three component costs.

These cost components are highly interrelated in that one component can be avoided only at the expense of another. For example, if the workforce is set too low relative to production, higher worker overtime cost is incurred. In turn, if production is set too low relative to demand, inventory outage cost is driven higher. The demand trend used in the experiment begins at 2500 units and increases at a rate of 50 units per period. This "unadjusted" demand is randomized ± 200 (uniform distribution) units to generate the actual demand. Each subject got a different random demand trend drawn from this distribution.

Appendix A shows the screen layout of the unaided version of the task simulator, which was programmed using the Lotus 1-2-3 macro language. The lower lines on the screen display the worker productivity index (5.67) and the optimal level at which to maintain inventory (320); both remain constant throughout the task. The upper right corner displays the current workforce and inventory levels. The upper left displays the total cost thus far incurred. The upper middle of the screen displays the sales orders for the next three periods. The middle of the screen is used to accept subjects' production and workforce level inputs. For the what-if treatment, a spreadsheet what-if function was embedded into the task simulator interface (Appendix B). With this function, subjects can simulate any number of scenarios per period to assess cost ramifications. They may enter workforce and production levels, and have the model calculate resulting costs. Subjects receive the following information for each period: the current workforce level, the current inventory level, and the next three periods of demand. Subjects enter "G" for "GO" when ready to process the decision they have entered for the current period. Users are then asked to confirm the "G" before the decision is processed. Upon confirmation, the results are computed and displayed. As shown in Appendix C, the display includes the new inventory and workforce levels, workforce change cost, worker over/idle time cost, inventory over/outage cost, total cost for the period, and the updated average total cost thus far incurred.

Subjects and Procedure. Forty-five managers attending an evening MBA course at a large midwestern university participated in Experiment 1, which employed a counterbalanced, within-subjects design. Subjects were randomly assigned to one of two treatment groups. One group ($n=23$) performed the 24-period production planning task using what-if analysis first and then unaided. The other group ($n=22$) performed the task first unaided and then using what-if analysis. Two days prior to the experiment, one of the authors lectured the entire group of subjects on the

production planning task. The general nature of the quadratic cost functions was explained and illustrated, although subjects were not given the specific coefficients. Subjects were told that their objective was to minimize long-run total cost. They were told that the purpose of the exercise was to give them the opportunity to learn about production planning using a realistic task simulator. They were also told that a what-if analysis function would be available, but that it is unknown whether what-if analysis is useful in the production scheduling task.

When the subjects arrived at the computing lab they were given an instruction sheet, directed to a personal computer, and given one of two versions of the experiment diskette. Subjects were unaware that there were two different versions of the diskette. Subjects using the unaided-first diskette: (1) performed four practice periods unaided, (2) performed 24 actual periods unaided, (3) performed 4 practice periods using what-if analysis, (4) were asked three belief questions before beginning 24 periods of actual what-if analysis, (5) performed 24 actual periods using what-if and observing outcome feedback on their performance, and (6) were asked the same three belief questions again. Subjects in the what-if-first treatment undertook the same basic sequence except that they first performed the task with what-if analysis and then without what-if analysis. To make subjects aware of the chance element of the task, they were informed that the randomized demand streams for the what-if and unaided segments of the task were independently generated from a common demand distribution. As an incentive, 20 \$10 cash prizes were awarded: \$10 was awarded to each of the top ten performers using what-if analysis and to each of the top ten unaided performers a week after the experiment. Since each subject performed the 24-period task both with and without what-if analysis, it was possible for subjects to win up to \$20. Subjects were allowed up to 2.5 hours to complete the task. Most finished in less than two hours (Mean = 1.97, Standard Deviation = .26).

Belief measurement. Users' beliefs regarding the performance benefits due to what-if analysis were measured with three items administered by the task simulator:

1. Compared to unaided decision making, my cost performance is (options arranged vertically: (A) Better using the interactive decision model. (B) The same with or without the interactive decision model. (C) Worse using the interactive decision model).
2. Compared to unaided decision making, the use of the interactive decision model (Significantly enhances/Does not affect/Significantly degrades) my cost performance. The descriptors in parentheses were used to anchor the left, center, and right of a horizontal 1-to-7 scale.
3. I would be most confident in the production planning task (options arranged vertically: (A) If I DID NOT use the interactive decision model. (B) Whether or not I used the interactive decision model. (C) If I DID use the interactive decision model).

Each question screen initially highlighted the center (neutral) response and instructed subjects to "Highlight your response using the arrow keys. Press the Enter key to select your response." These items were adapted from a more general measure of perceived usefulness developed and validated by [11] and [12], who showed that

such beliefs are theoretically and empirically linked to people's willingness to use computer-based tools in the workplace. The wording of the items were modified to be more specific to the decision simulation task and what-if model being studied in the present experiment. Multiple questions were used in order to enhance reliability and validity. A mixture of 3-point and 7-point scales and horizontal and vertical response formats was employed in order to reduce common method variance.

These three questions were asked twice during the task: before and after receiving outcome feedback regarding how much the use of what-if analysis helped their performance in this particular situation. At both questioning times the three questions were significantly correlated with each other and loaded on a single dimension according to principal components analysis (Table 1). A three-item belief scale was formed by: coding item one from one (A) to three (C); linearly transforming item two from a 7-point to a 3-point metric $((X-1)/3)+1$; coding item three from three (A) to one (C), and then averaging the three items. The resulting belief measure can range from one (subject negative toward what-if analysis) to three (subject positive toward what-if analysis), and had Cronbach's alpha reliability of .84 when asked the first time (prior to full outcome feedback) and .88 the second time. Table 2 gives the means and standard deviations for the belief measurements.

Results

What-if analysis had no significant effect on actual performance compared to unaided decision making, and yet most subjects reported believing what-if to be beneficial. As shown in Table 3, mean cost performance was \$2.7 million using what-if analysis versus \$2.9 million unaided (paired- $t(44)=.019$, n.s.). The standard deviations of these means were \$6.7 million and \$4.8 million, respectively. What-if analysis improved cost performance for only 26 out of 45 subjects (58 percent).

Consistent with the illusion of control hypothesis, subjects overestimated how much what-if analysis helped their performance. Table 4 shows the number of subjects who felt it had a positive effect, no effect, or a negative effect on their performance, for each treatment, and for each of the two times the questionnaire was administered. Thirty-nine out of 45 subjects (87 percent) thought their decision performance was better using what-if analysis, five indicated no difference, and only one thought his or her decision performance was worse using what-if analysis (per belief question 1). Recall that the 3-item belief measure described above was scaled so that values could range from one (what-if positive) to three (what-if negative), with a neutral midpoint of two. Mean beliefs on this measure (Table 5) were significantly less than 2.0 both before ($t(44)=-14.178$, $p<.001$) and after ($t(44)=-8.362$, $p<.001$) receiving outcome feedback.

Did outcome feedback have any effect on the illusion of control? Most subjects remained overconfident in what-if analysis even in the face of outcome feedback about their aided versus unaided performance: only two of the people who did worse using what-if analysis changed their evaluation of its benefit (Table 4). Among the 19 subjects who actually performed worse using what-if analysis, there was a significant shift away from believing what-if to be beneficial, from a mean of 1.23 to 1.55 (Table 5). This represents a significant difference from the 1.19 rating given by the what-if beneficial group ($t(43)= 2.383$, $p<.05$). Nevertheless,

Table 1: Correlation and principal component analyses of belief items: Experiment 1.

	Correlations ^a		Principal Component ^b
	Q1	Q2	Factor 1
Before Full Outcome Feedback			
Question 1			.88
Question 2	.65		.87
Question 3	.63	.61	.86
Eigenvalue			2.26
Percentage variance explained			75.3
After Full Outcome Feedback			
Question 1			.89
Question 2	.76		.93
Question 3	.65	.73	.88
Eigenvalue			2.42
Percentage variance explained			80.8

^aAll correlations were significant at $p < .001$.

^bEigenvalue for second factor was below 1.0 both before (.40) and after (.36) full outcome feedback.

Table 2: Means and standard deviations of user perceptions of what-if effectiveness: Experiment 1.

	Mean	Standard Deviation	Cronbach's Alpha
Before Full Outcome Feedback			
Question 1	1.18	.49	
Question 2	1.87	1.14	
Question 3	1.16	.42	
Three-Item scale	1.21	.38	.84
After Outcome Feedback			
Question 1	1.36	.71	
Question 2	2.20	1.42	
Question 3	1.27	.58	
Three-Item scale	1.34	.53	.88

most of these subjects continued to believe in the efficacy of what-if. Given that performance outcomes are partly determined by chance, and subjects were made aware of this, it is reasonable for these subjects to attribute their particular results to random factors, and to believe that what-if might be beneficial over the long term. Thus, while outcome feedback did have a significant effect on beliefs, the

Table 3: Effect of what-if analysis on actual cost performance: Experiment 1.

Mean Cost:	What-if analysis	\$2,675,932
	Unaided	\$2,873,932

Paired-*t* with 44 degrees of freedom=.109, non-significant.
 What-if analysis helped 26 of 45 subjects (58%).

Table 4: Frequency of belief responses on belief question 1: Experiment 1.

Group	Stated Belief	Before Outcome Feedback	After Outcome Feedback
Total sample (<i>n</i> =45)	1. What-if positive	39	36
	2. No difference	5	6
	3. What-if negative	1	3
Subsamples:			
What-if actually better (<i>n</i> =26)	1. What-if positive	23	22
	2. No difference	2	3
	3. What-if negative	1	1
What-if actually worse (<i>n</i> =19)	1. What-if positive	16	14
	2. No difference	3	3
	3. What-if negative	0	2

Note: Based on belief item #1.

Table 5: Mean beliefs about effectiveness of what-if analysis: Experiment 1.

	Total Sample (<i>n</i> =45)	Subsample		Difference
		What-if Better (<i>n</i> =26)	What-if Worse (<i>n</i> =19)	
Before outcome feedback	1.21	1.19	1.23	Non-significant
After outcome feedback	1.34	1.19	1.55	<i>p</i> < .05

Note: Based on 3-item scale:
 1="What-if Positive"
 2="No difference"
 3="What-if Negative"

resulting beliefs were still significantly in favor of what-if analysis. Note that even this fairly ambiguous outcome feedback would not ordinarily be available to real world decision makers who, acting on their beliefs in the superiority of what-if analysis, may avoid using unaided decision making enough to assess its relative performance.

Consistent with the idea that what-if analysis creates an illusion of control through more active involvement in the decision task, people took longer using what-if than unaided (Wilcoxin signed-rank test, 34/44, $p < .01$). The order of within-subject treatment (what-if first versus unaided first) had no effect on what-if performance ($t(43) = .871$, n.s.), unaided performance ($t(43) = .588$, n.s.), or performance beliefs ($t(43) = .974$, n.s.). Unaided performance was therefore not significantly different before versus after the what-if session (mitigating concerns about whether the what-if-first group experienced learning effects that carried over to the unaided session) or vice versa.

In summary, Experiment 1 found that, although what-if analysis had no performance advantage over unaided decision making, subjects attributed a performance advantage to what-if. Full outcome feedback had a limited effect on these beliefs.

EXPERIMENT 2: WHAT-IF ANALYSIS VERSUS A QUANTITATIVE DECISION RULE

In Experiment 2, decision makers using what-if analysis were given the recommendations of a quantitative decision rule, described below, which they were free to follow or not. Subjects completed 24 periods of the same production simulation task using the same what-if model as used in Experiment 1, with no unaided group.

H4: Because of the illusion of control, we expect decision makers to overestimate the value of what-if analysis and to avoid following the recommendations of the quantitative decision rule.

There were two orthogonal between-subjects manipulations. Half of the subjects received feedback regarding what their cumulative costs would have been had they used the recommendations of the quantitative decision rule exclusively, whereas the other half received no such feedback (and thus were more representative of real-life decision makers who ordinarily do not receive feedback regarding actions not taken). The decision rule was described in two different ways: for half of the subjects, the decision rule was discussed in a neutral manner (see below), whereas the other half were told the decision rule outperformed decision makers not using it about 90 percent of the time.

H5: We expect full outcome feedback [2] [19] to significantly reduce the illusion of control, thus increasing confidence in the decision rule.

H6: We expect rule description [5] [44] to significantly reduce the illusion of control, thus increasing confidence in the decision rule.

Method

Task. Experiment 2 used the same production scheduling task as Experiment 1. The input screen was modified to incorporate the advice of a quantitative decision rule (Appendix D). The quantitative decision rule used to generate the recommendations shown in Appendix D employs a straightforward moving average of demand to

determine production level, and then adjusts workforce to be the number of workers needed to produce that level by dividing by the worker productivity index of 5.67:

$$\text{Production} = .5 \text{ demand}_t + .3 \text{ demand}_{t+1} + .2 \text{ demand}_{t+2} \\ - \text{current inventory} + \text{safety stock of 320,}$$

$$\text{Workforce} = \text{Production} / \text{worker productivity constant of 5.67.}$$

The rule was expected to be effective because it tracks demand trends more systematically than decision makers using what-if analysis. Demand growth was increased to 100 units per period (as opposed to 50 in Experiment 1), with the same starting point of 2500 units and randomization of ± 200 units, in order to take greater advantage of the rule's tracking capability. Again, each subject got a different random demand trend drawn from this distribution. Pretesting the rule using sample demand streams confirmed its effectiveness. Thus, we have constructed a situation in which a fairly simple decision rule is expected to outperform decision makers using what-if analysis, analogous to situations pitting the unaided decision maker against simple but effective decision rules, discussed earlier, in order to investigate whether the illusion of control caused by what-if analysis, together with the ambiguous feedback characteristics of real life decision making, will impede subjects from learning that the rule outperforms unaided what-if analysis.

Subjects and Procedure. Fifty-two daytime MBA students attending an elective class on management decision models at a large midwestern university participated in Experiment 2. On a five-point question asking personal computer experience, with anchor points "Very Little" (1), "Average" (3) and "Extensive" (5), mean experience of the sample was 3.2, and 98 percent of subjects selected two, three, or four. Experiment 2 employed a 2X2 factorial between-subjects design. The two manipulations were outcome feedback (year-to-date performance of the decision rule versus no feedback) and rule description (neutral versus 90 percent accuracy). Feedback was manipulated by including or not including a display of the rule's performance on the results screen of the task simulator (included on example screen, Appendix D). The results screen was the same as Appendix C except that the feedback group received a report of year-to-date rule performance similar to that shown in Appendix D. Rule description was manipulated in the instructions, as discussed below.

Two days before the experiment, one of the authors lectured the entire sample of subjects on the production planning task. To reduce hypothesis guessing and demand characteristics, subjects were told that the experimenters were interested in how much importance decision makers placed on each of the three cost components in arriving at production and workforce decisions. The general nature of the quadratic cost functions was explained and illustrated, although subjects were not given the specific coefficients. Subjects were told they would be provided the advice of a "decision heuristic" which they were free to follow or not. The formulas of the quantitative decision rule were presented to subjects in the instructional session prior to the experiment as well as in the instructional handout given at the beginning of the experiment. The rule was referred to as a "heuristic" in the context of the experiment to relate it to the concept of heuristics presented in class, namely, a sometimes helpful but not necessarily optimal rule for decision making.

When subjects arrived at the computing laboratory, they were given an instruction sheet, led to a personal computer, and given one of four versions of a diskette selected at random. Subjects did not know that there were multiple treatment groups. Eleven subjects were assigned to the neutral-description/no-feedback group; 12 to the neutral-description/feedback group; 14 to the 90 percent-accurate/no-feedback group, and 15 to the 90 percent-accurate/feedback group. On the first screen of the task simulation, instructions for subjects in the neutral-description group said "The heuristic produces good, but rarely optimal, decisions" whereas instructions for subjects in the 90 percent-accuracy group said "The heuristic has historically outperformed 90 percent of your peers." As an incentive, \$10 cash prizes were awarded to each of the top 20 performers a week after the experiment. All subjects completed the task within the 1.5 hours allowed.

Belief measurement. After performing the 24-period simulation, subjects were asked questions automatically by the task simulator. For all questions, instructions on the screen asked subjects to "Highlight your response using the arrow keys. Press the Enter key to select your response." The first three were filler items designed to reinforce the cover story, asking subjects to rate from one to seven how important inventory costs, workforce change costs, and overtime/idle time costs were in their decision making. The fourth question asked subjects to "Please indicate your degree of past experience using personal computers" on a 1-5 scale with anchor points: "Very Little" (1), "Average" (3), and "Extensive" (5).

The next four items measured subjects' beliefs about the performance advantages of using the quantitative decision rule.

1. My decision performance was, or would have been, better if I followed the heuristic's recommendations. (Options arranged vertically: (A) Better following the heuristic. (B) The same with or without following the heuristic. (C) Worse following the heuristic.)
2. Compared to not using the heuristic at all, strictly following the heuristic's recommendations would have (on a horizontal 1-7 scale: Significantly Improved/Not Affected/Significantly Degraded) my decision performance.
3. I would be most confident in my performance in the production planning task by (options arranged vertically: (A) Following the heuristic's recommendations. (B) Partially following the heuristic's recommendations. (C) Not following the heuristics recommendations).
4. Do you think that following the heuristic's recommendations improved, or would have improved, your overall decision performance on this task? (Definitely Yes/Neutral/Definitely No) on a horizontal 1-7 scale.

These items were adapted from a more general measure of perceived usefulness developed and validated by [11] and [12], who showed that such beliefs are theoretically and empirically linked to people's willingness to use computer-based tools in the workplace. The wordings of the items were modified to be made more specific to the decision simulation task and what-if model being studied in the present experiment. Multiple questions were used in order to enhance reliability and validity. A mixture of 3-point and 7-point scales and horizontal and vertical response formats was employed in order to reduce common method variance.

Table 6: Correlation and principal component analyses of belief items: Experiment 1.

	Correlations ^a			Principal Component ^b
	Q1	Q2	Q3	Factor 1
Question 1				.87
Question 2	.73			.83
Question 3	.58	.40		.78
Question 4	.58	.62	.63	.85
Eigenvalue				2.78
Percentage variance explained				69.4

^aAll correlations were significant at $p < .01$.

^bEigenvalue for second factor was below 1.0 (.63).

The four belief items were significantly correlated and loaded on a single dimension according to principal components analysis (Table 6). A 4-item belief index was formed by first linearly transforming the two 7-point questions to a 3-point metric $((X-1)/3)+1$, and then averaging the four belief items. The resulting 4-item belief scale can range from one (decision rule positive) to three (decision rule negative), and had Cronbach's alpha reliability of .85. Lastly, one question asked subjects how much they used the advice from the quantitative decision rule. "How often did you follow the heuristic's recommendation?" on a 1-7 scale with anchor points "Not at all" (1), "About half the time" (4), and "All the time" (7). Table 7 gives the means and standard deviations for the belief questions and self-reported rule use.

Results

As expected, the quantitative decision rule, if used, would have led to large performance advantages. Also as expected, users overestimated the effectiveness of what-if analysis relative to the decision rule. Overall, the average cost of running the production operation for 24 periods was \$6,599,471 (Table 8). If subjects had adhered to the advice offered by the decision rule, their costs would have averaged \$472,525, for a mean savings of \$6,126,945, a substantial improvement (paired- $t(51)=3.810$, $p < .001$). Despite this vast performance difference, subjects not shown the rule's performance (true of the real-world avoider of such a decision rule) failed to detect any performance difference between what-if analysis and the decision rule. Beliefs were assessed using the 4-item measure described above which was scaled from one (rule positive) to three (rule negative) with a neutral midpoint of two. Mean beliefs (Table 9) of subjects not receiving outcome feedback on the rule's performance was 2.1, not significantly different from the neutral midpoint ($t(24)=.714$, n.s.). Subjects receiving feedback on the performance of the decision rule had mean beliefs of 1.5, significantly favoring the decision rule ($t(26)=-9.196$, $p < .001$).

As expected, feedback on the rule's performance had a significant effect on beliefs about the relative effectiveness of what-if analysis versus the quantitative

Table 7: Means and standard deviations of beliefs and self-reported rule use: Experiment 2.

		Mean	Standard Deviation	Cronbach's Alpha
Beliefs	Question 1	1.67	.83	
	Question 2	2.90	1.86	
	Question 3	2.02	.58	
	Question 4	3.12	1.67	
	Scale	1.76	.54	.85
Self-Reported rule use		2.24	1.60	

decision rule. Counter to expectations, rule description (neutral versus 90 percent) had no significant effect on either beliefs or performance. The effect of rule feedback and description were assessed by two-way analysis of variance. Outcome feedback comparing actual performance to that which would have occurred given strict adherence to the decision rule significantly increased subjects' beliefs in the superiority of the decision rule ($F(1,48)=8.045, p<.01$), but had no effect on actual performance ($F(1,48)=.359, n.s.$). Differences in rule description affected neither beliefs ($F(1,48)=.020, n.s.$) nor performance ($F(1,48)=.064, n.s.$). There were no interaction effects between feedback and description with respect to either beliefs ($F(1,48)=.230, n.s.$) or performance ($F(1,48)=.126, n.s.$). Self-reported personal computer experience had no effect as a covariate on either beliefs ($F(1,47)=.988, n.s.$) or performance ($F(1,47)=1.095, n.s.$), and experience did not differ significantly among the four treatment groups ($F(2,49)=.812, n.s.$).

Although outcome feedback caused people to significantly revise their beliefs about the relative effectiveness of what-if analysis versus the quantitative decision rule, this was not accompanied by a shift in rule usage, or resulting performance: there was no significant effect of feedback on either performance or self-reported rule usage. The self-reported measure of rule usage was unaffected by feedback ($F(1,48)=.103, n.s.$), description ($F(1,48)=.559, n.s.$), and their interaction ($F(1,48)=.444, n.s.$). Thus, subjects generally chose not to follow the recommendations of the decision rule, irrespective of feedback or description.

It is surprising that the \$4.6 million average cost performance for subjects receiving feedback was not significantly lower than the \$8.8 million average for non-feedback subjects. As in Experiment 1, the standard deviations were quite large (\$9.4 million for feedback subjects and \$13.4 for non-feedback subjects). A non-parametric test of this difference was performed in case departures from normality undermined the power of the F -test, but was also non-significant (Mann-Whitney $U=291, Z=-.85, n.s.$). To investigate the reasons for these large standard deviations, we examined individual cost trends over the 24 time periods of the simulation for each of the three cost components. Although the variety of patterns was great, two discernible patterns emerged. Some people were apparently "trackers" who quickly discovered that moderate adjustments in workforce or production in the short term might help avoid larger, costlier changes in the future as the business responds to growing demand. Note that the decision rule is also based on a tracking

Table 8: Effect of feedback and description on cost performance: Experiment 2.

Rule Description ^b	Actual Performance		
	Rule Feedback ^a		Combined
	No	Yes	
Neutral	8,128,181 (<i>n</i> =11)	5,193,369 (<i>n</i> =12)	6,596,975
90% Accurate	9,322,207 (<i>n</i> =14)	4,062,078 (<i>n</i> =15)	6,601,450
Combined	8,796,835	4,564,874	6,599,471
Decision rule performance ^c			472,525

^aRule feedback had no significant effect on actual performance.

^bRule description had no significant effect on actual performance.

^cDecision rule outperformed decision makers using what-if analysis ($p < .001$).

strategy. Other people, the “myopics,” apparently tried to optimize in the short run by avoiding increases in production or workforce, possibly because they believed the upward trend may reverse itself. In any case, these people kept short-term costs low initially and later incurred spiraling costs in order to satisfy demand. The trackers tended to keep costs from going up non-linearly, while the myopics often experienced non-linear cost run-ups. This is an example where a seemingly conservative decision strategy, for example, avoiding the hiring of additional staff in the short run, has traumatic consequences for the business. There seemed to be a mixture of trackers and myopics in every treatment group, hence the high standard deviations in cost performance.

We then investigated whether feedback may have caused some of the subjects to switch from a myopic to a tracking strategy (i.e., to behave more like the rule even if not adhering to it fully). Trackers tended to incur less than \$1 million in total costs. Forty-four percent of the feedback subjects incurred less than \$1 million compared with 32 percent of the non-feedback subjects ($\chi^2(1)=.8$, n.s.). We also looked at degree of rule adherence by calculating the absolute difference between the rule’s recommendations and the subject’s choice of workforce and production levels, finding no differences across treatment group. This is consistent with self-reported rule use, which also found no differences across groups. Thus, the high variance seems to be due to individual variations in strategy choice (e.g., tracker versus myopic) more so than strategy switching resulting from feedback or rule description. People receiving feedback may have been resistant to switching over to using the rule midstream in light of the large one-time penalty associated with that. That is, the feedback they received showed performance of the decision rule if used from the first period, whereas there was no way to get a “clean slate” once several decision periods have transpired and the efficacy of the rule becomes apparent. This may explain why rule use did not increase with full outcome feedback while perceived efficacy of the rule did. In other words, by the end of the game, people may have realized that it would have been much better to use the

Table 9: Effect of feedback and description on mean performance beliefs: Experiment 2.

Rule Description ^b	Rule Feedback ^a		Combined
	No	Yes	
Neutral	2.1	1.6	1.8
90% Accurate	2.1	1.3	1.7
Combined	2.1	1.5	1.8

Note: Based on 4-item scale:

1="Positive Toward Decision Rule"

2="Neutral Toward Decision Rule"

3="Negative Toward Decision Rule"

^aRule feedback had a significant effect on performance beliefs ($p < .01$).

^bRule description had no effect on performance beliefs.

rule from the outset, even though it may not have been as beneficial to switch in mid-game. This may be an artifact of using a finite simulation in that, over a longer time horizon, subjects may have been impelled to incur the switching costs in order to take advantage of the highly effective decision rule.

In summary, decision makers generally did not initially recognize the substantial benefits of the quantitative decision rule compared to what-if analysis, chose not to follow its recommendations, and therefore incurred costs which were much higher than those obtained by the decision rule. This occurred despite the fact that the subjects were students in a management science type class, and so might be expected to favor OR/heuristic models to other approaches. When decision makers were shown how well they would have done by following the decision rule's advice, they shifted their beliefs in favor of the decision rule (as measured at the end of the 24-period simulation). This shift in beliefs was not accompanied by a change in either rule usage or performance, as neither were significantly affected by feedback. Changing the way the rule was described had no effect on either beliefs or performance.

DISCUSSION

Does what-if analysis create an illusion of control? The present results suggest yes (Table 10). Experiment 1 found that people strongly believed that what-if analysis was superior to unaided decision making, even though no actual performance difference occurred. Experiment 2 found that people generally failed to detect a substantial performance advantage of a quantitative decision rule over what-if analysis. In both cases, people exaggerated the effectiveness of what-if analysis relative to its alternatives. By promoting more active involvement in the decision-making task, what-if analysis appears to inflate decision makers' sense of how well they perform in the production planning task. The illusion of control associated with what-if analysis is problematic because it may cause decision makers to use what-if when it doesn't really help, and to prefer what-if over superior decision support technologies.

Decision makers had difficulty learning from experience that their beliefs about what-if were inflated. In Experiment 1, subjects overrated what-if analysis

Table 10: Perceived versus actual effectiveness of decision support: Summary of findings.

Experiment 1			
Performance	What-If Analysis	Unaided	Significance
Perceived	Better	Worse	$p < .001$
Actual	No difference	No difference	Non-significance

Experiment 2			
Performance	What-If Analysis	Decision Rule	Significance
Perceived	No difference	No difference	Non-significance
Actual	Worse	Better	$p < .001$

even when given the opportunity to witness their own decision performance both with and without the use of what-if analysis. In Experiment 2, decision makers did not perceive any significant performance difference between what-if analysis and the quantitative decision rule unless they were shown explicitly how well the rule would have performed if used, information often not available in practice.

What can be done to minimize the illusion of control due to what-if analysis? Exaggerated control beliefs are sometimes adaptive and beneficial [38], leading people to seek to exercise greater control in partially controllable circumstances. However, if the inflated sense of control associated with what-if analysis causes decision makers to avoid superior methods which are readily available, as we found, then the consequences are dysfunctional, and methods for reducing the perceptual bias should be sought. We addressed two interventions for debiasing: providing explicit outcome feedback showing relative performance of alternative tools, and changing how decision tools are described to decision makers. Outcome feedback was only partially effective: in Experiment 1 it did not overcome the illusion of control, while in Experiment 2 it did. Probably the large size of the difference between what-if and rule performance in Experiment 2 explained why outcome feedback was effective there only. In contrast to Powell's [44] finding, changing the way in which the decision rule was described (in our case from neutral to 90 percent accurate) did not significantly influence beliefs or behavior with respect to the decision rule in Experiment 2. Thus, more research is needed to identify interventions that would be effective in reducing the illusion of control.

One interesting finding of the present research is that greater cognitive effort, in the form of what-if analysis, did not pay off in terms of improved decision making. In Experiment 1, users of what-if analysis took significantly longer to make their decisions, with no significant performance gain. In Experiment 2 people used what-if instead of a much more effective decision rule which would have been nearly effortless to employ. Research in other contexts has similarly demonstrated limitations in the effort-performance link, and shown that cognitively simple decision rules often perform quite well [24] [26] [30] [33] [34] [49]. Computer-augmented or not, cognitive effort has only limited capability for generating performance

improvements in task environments partially governed by uncontrollable factors, true of many real life situations.

Will our results generalize to other situations? One possible limitation relates to our use of MBA students. In an earlier study using this production planning task, no significant differences were found between MBA's and managers, although neither group had work experience as production planners [45]. Also, in Experiment 1, all participants were evening students employed full time. A second limitation concerns the extent to which the production planning task is representative of real-world decision making. While recognizing that differences do exist, the basic nature of the task—recurring decisions under uncertainty—is characteristic of many domains: maintaining an investment portfolio over time, periodic planning, budgeting, medical diagnosis, and others [25] [33] [34]. A third limitation concerns the extent to which the findings will generalize to tasks beyond the class of recurring decisions under uncertainty. Suggesting that the illusion of control may indeed generalize, Pentland's [43] field study of 1110 Internal Revenue Service accountants found that they believed using personal computers was superior to manual methods although such use did not improve performance. These considerations increase confidence in our findings, although additional research is needed to verify the existence of the illusion of control, to gauge its prevalence, and to better understand its nature, contingencies, determinants, and remedies.

The present findings have practical implications for designers and users of decision support systems. Many design techniques emphasize guiding the design process according to user perceptions and preferences. While such a practice is important from the standpoint of increasing the chances for user acceptance, since users are more apt to accept tools they believe are effective [11] [12], possible discrepancies between perceived and actual performance imply that resulting tools, though accepted, may in some cases actually be ineffective or even dysfunctional (see, [1] [15] [35]). For example, although it is commonly believed that "greater flexibility and choice in software aids will promote improved user performance," system-guided model manipulation significantly improved performance compared to user-guided manipulation [16, p. 461]. Designers should be cautious when designing according to user preferences, and should seek independent confirmation of the efficacy of user-preferred designs. Our findings suggest paying particular attention to decision support techniques that increase active involvement in the decision task, which are especially susceptible to the illusion of control. Thus, this research raises several important questions for future research. We need to gain a better understanding of how pervasive the illusion of control is in DSS contexts, what its ramifications are, and how to manage it. [Received: June 9, 1992. Accepted: November 17, 1993.]

REFERENCES

- [1] Ahituv, N. Assessing the value of information: Problems and approaches. *Proceedings of ICIS-89*, Boston, December 1989, 315-325.
- [2] Alavi, M., & Henderson, J. C. An evolutionary strategy for implementing a decision support system. *Management Science*, 1981, 27, 1-9.
- [3] Aldag, R. J., & Power, D. J. An empirical assessment of computer-assisted decision analysis. *Decision Sciences*, 1986, 17, 572-588.
- [4] Alter, S. L. *Decision support systems: Current practice and continuing challenge*. Reading, MA: Addison-Wesley, 1980.

- [5] Arkes, H. R., Dawes, R. M., & Christensen, C. Factors influencing the use of a decision rule in a probabilistic task. *Organizational Behavior and Human Decision Processes*, 1986, 37, 93-110.
- [6] Benbasat, I., & Dexter, A. S. Individual differences in the use of decision support aids. *Journal of Accounting Research*, 1982, 20, 1-11.
- [7] Benbasat, I., & Nault, B. R. An evaluation of empirical research in managerial support systems. *Decision Support Systems*, 1990, 6, 203-226.
- [8] Benbasat, I., & Schroeder, R. G. An experimental investigation of some MIS design variables. *MIS Quarterly*, 1977, 1, 37-50.
- [9] Bowman, E. J. Consistency and optimality in managerial decision making. *Management Science*, 1963, 9, 310-322.
- [10] Chakravarti, D., Mitchell, A. A. & Staelin, R. Judgment-based marketing decision models: An experimental investigation of the decision calculus approach. *Management Science*, 1979, 25, 251-262.
- [11] Davis, F. D. Perceived usefulness, perceived ease of use, and user acceptance of information technology. *MIS Quarterly*, 1989, 13, 319-340.
- [12] Davis, F. D., Bagozzi, R. P., & Warshaw, P. R. User acceptance of computer technology: A comparison of two theoretical models. *Management Science*, 1989, 35, 982-1003.
- [13] Dawes, R. M. The robust beauty of improper linear models in decision making. *American Psychologist*, 1979, 34, 571-582.
- [14] Dawes, R. M., & Corrigan, B. Linear models in decision making. *Psychological Bulletin*, 1974, 81, 95-106.
- [15] Dickson, G. W., Senn, J. A., & Chervany, N. L. Research in management information systems: The Minnesota experiments. *Management Science*, 1977, 23, 913-923.
- [16] Dos Santos, B. L., & Bariff, M. L. A study of user interface aids for model-oriented decision support systems. *Management Science*, 1988, 34, 461-468.
- [17] Einhorn, H. J. Learning from experience and suboptimal rules in decision making. In T.S. Wallsten (Ed.), *Cognitive processes in choice and decision behavior*. Hillsdale, NJ: Erlbaum, 1980, 66-80.
- [18] Einhorn, H. J., & Hogarth, R. M. Unit weighting schemes for decision making. *Organizational Behavior and Human Performance*, 1975, 13, 171-192.
- [19] Einhorn, H. J., & Hogarth, R. M. Confidence in judgment: Persistence of the illusion of validity. *Psychological Review*, 1978, 85, 395-416.
- [20] Einhorn, H. J., & Hogarth, R. M. Behavioral decision theory: Processes of judgment and choice. *Annual Review of Psychology*, 1981, 32, 53-88.
- [21] Fripp, J. How effective are models? *International Journal of Management Science (Omega)*, 1985, 13, 19-28.
- [22] Fudge, W. K., & Lodish, L. M. Evaluation of the effectiveness of a salesman's planning system by field experimentation. *Interfaces*, 1977, 8, 97-106.
- [23] Goslar, M., Green, G., & Hughes, T. Decision support systems: An empirical assessment for decision making. *Decision Sciences*, 1986, 17, 79-91.
- [24] Hammond, K. R., Hamm, R. M., Grassia, J., & Pearson, T. Direct comparison of the efficacy of intuitive and analytical cognition in expert judgment. *IEEE Transactions on Systems, Man & Cybernetics*, 1987, SMC-17, 753-770.
- [25] Hogarth, R. M. Beyond discrete biases: Functional and dysfunctional aspects of judgmental heuristics. *Psychological Bulletin*, 1981, 90, 197-217.
- [26] Hogarth, R. M., & Makridakis, S. The value of decision making in a complex environment: An experimental approach. *Management Science*, 1981a, 27, 93-107.
- [27] Hogarth, R. M., & Makridakis, S. Forecasting and planning: An evaluation. *Management Science*, 1981b, 27, 115-138.
- [28] Holt, C. C., Modigliani, F., & Muth, J. F. Derivation of a linear decision rule for production and employment. *Management Science*, 1956, 2, 159-177.
- [29] Holt, C. C., Modigliani, F., Muth, J. F., & Simon, H. A. *Planning production, inventories, and work force*. Englewood Cliffs, NJ: Prentice-Hall, 1960.
- [30] Johnson, E. J., & Payne, J. W. Effort and accuracy in choice. *Management Science*, 1985, 31, 395-414.
- [31] Keen, P. G. W., & Scott Morton, M. S. *Decision support systems: An organizational perspective*. Reading, MA: Addison-Wesley, 1978.
- [32] Kleinmuntz, B. Why we still use our heads instead of formulas: Toward an integrative approach. *Psychological Bulletin*, 1990, 107, 296-310.
- [33] Kleinmuntz, D. N. Cognitive heuristics and feedback in a dynamic decision environment. *Management Science*, 1985, 31, 680-702.

- [34] Kleinmuntz, D. N., & Thomas, J. B. The value of action and inference in dynamic decision making. *Organizational Behavior and Human Decision Processes*, 1987, 39, 341-364.
- [35] Kottemann, J. E., & Remus, W. E. Evidence and principles of functional and dysfunctional DSS. *International Journal of Management Science (Omega)*, 1987, 15, 135-144.
- [36] Kottemann, J. E., & Remus, W. E. The performance effects of what-if analysis in a production planning task. *Proceedings of conference on environments for supporting decision processes*, Budapest, 1990.
- [37] Langer, E. J. The illusion of control. *Journal of Personality and Social Psychology*, 1975, 32, 311-328.
- [38] Langer, E. J. *The psychology of control*. Beverly Hills, CA: Sage, 1983.
- [39] Langer, E. J., & Roth, J. Heads I win, tails it's chance: The illusion of control as a function of the sequence of outcomes in a purely chance task. *Journal of Personality and Social Psychology*, 1975, 32, 951-955.
- [40] McIntyre, S. An experimental study of the impact of judgment-based marketing models. *Management Science*, 1982, 28, 17-23.
- [41] Meehl, P. E. *Clinical versus statistical prediction: A theoretical analysis and a review of the evidence*. Minneapolis: University of Minnesota Press, 1954.
- [42] Paese, P. W., & Snizek, J. A. Influences on the appropriateness of confidence in judgment: Practice, effort, information and decision-making. *Organizational Behavior and Human Decision Processes*, 1991, 48, 100-130.
- [43] Pentland, B. T. Use and productivity in personal computing: An empirical test. *Proceedings of ICIS-89*, Boston, December 1989, 211-222.
- [44] Powell, J. L. An attempt at increasing decision rule use in a judgment task. *Organizational Behavior and Human Decision Processes*, 1991, 48, 89-99.
- [45] Remus, W. E. Graduate students as surrogates for managers in experiments on business decision making. *Journal of Business Research*, 1987, 14, 19-25.
- [46] Sharda, R., Barr, S. H., & McDonnell, J. C. Decision support system effectiveness: A review and an empirical test. *Management Science*, 1988, 34, 139-159.
- [47] Sprague, R. H., & Carlson, E. D. *Building effective decision support systems*. Englewood Cliffs, NJ: Prentice-Hall, 1982.
- [48] Turban, E. *Decision support and expert systems*. New York: Macmillan, 1988.
- [49] Wilson, T. D., & Schooler, J. W. Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology*, 1991, 60, 181-192.
- [50] Yates, J. F., & Kulick, R. M. Effort control and judgments. *Organizational Behavior and Human Performance*, 1977, 20, 54-65.

Appendix A: Unaided input screen: Experiment 1.

Total Cost To Date	43677	Current Inventory	173
		Current Workforce	490
Period	4	5	6
Sales	2523	2699	2710
Units to Produce	2976		<---Input
Workforce Level	490		<---Cells
Worker Productivity is 5.67 units per worker			
The Optimal Level of Inventory is 320			
Type G to Go with current decision			

Appendix B: What-if analysis input screen: Experiment 1.

Total Cost To Date	396226	Current Inventory	14
		Current Workforce	525
Period	9	10	11
Sales	3419	3572	3581
Units to Produce	3419	3572	3575 <---Input
Workforce Level	550	555	560 <---Cells
Ending Inventory	14	14	8 <---
Workforce Level Change Cost	40188	1608	1608 <---Simulated
Worker Over/Idle Time Cost	72270	144602	127872 <---Results
Non-Optimal Inventory Cost	1873	1873	1947 <---
Total Cost	114301	148083	131427 <---

Worker Productivity is 5.67 units per worker
 The Optimal Level of Inventory is 320

Type G to Go with current decision

Appendix C: Results screen: Experiment 1.

Total Cost To Date	222785	Current Inventory	253
		Current Workforce	451
Period	3		
Sales	2834		
Units to Produce	2834		
Workforce Level	451		
Ending Inventory	253		<---
Workforce Level Change Cost	0		<---Actual
Worker Over/Idle Time Cost	61308		<---Results
Non-Optimal Inventory Cost	90		<---
Total Cost	61398		<---

Worker Productivity is 5.67 units per worker
 The Optimal Level of Inventory is 320

Type C to Continue with next decision period

Appendix D: Input screen: Experiment 2.

Your Total Cost:	37368	Current Inventory	256
What Heuristic:	26940	Current Workforce	483
would have done			
Heuristic's	Period	4	5
Recommendation	Sales	2883	2994
2869	Produce	2869	2994
506	Employ	506	515
			3000 <---Input
			530 <---Cells
242	Ending Inventory	242	242
34015	Hire/Layoff \$	34015	5208
0	Over/Idletime \$	0	4375
122	Inventory \$	122	122
34137	Total Cost \$	34137	9705
			111 <---
			14468 <--Simulated
			21 <--Results
			874 <---
			15363 <---
Worker Productivity is 5.67 units per worker			
The Optimal Level of Inventory is 320			
Type G to Go with current decision			

Note: Display of "What Heuristic would have done" was provided only to subjects in the outcome feedback condition.

Fred Davis is Associate Professor of Information and Decision Sciences at the Carlson School of Management, University of Minnesota. He received his Ph.D. in management science from the Massachusetts Institute of Technology's Sloan School of Management. He conducts research on user acceptance of information technology and computer-supported decision making which has appeared in journals such as *Decision Sciences*, *Journal of Applied Social Psychology*, *Management Science*, *MIS Quarterly*, and *Organizational Behavior and Human Decision Processes*.

Jeff Kottemann is Associate Professor of Computer and Information Systems at the University of Michigan Business School. He received his Ph.D. in management information systems from The University of Arizona. His research interests include assessing the impact of computerized decision aids on performance and decision makers' perceptions of their performance, and devising tools and techniques for development of large-scale information systems and modeling systems. His work has been published in such journals as *Communications of the ACM*, *Decision Sciences*, *Journal of MIS*, *MIS Quarterly*, and *Organizational Behavior and Human Decision Processes*.