DESIGNING DECISION SUPPORT SYSTEMS TO HELP AVOID BIASES & MAKE ROBUST DECISIONS, WITH EXAMPLES FROM ARMY PLANNING

B. Chandrasekaran
Department of Computer & Information Science
The Ohio State University
Columbus, OH 43210

ABSTRACT

This paper is concerned with two sets of issues related to optimality in planning. The first is a proposal that design of decision support systems (DSS's) for planning should aim to support the planner in generating a plan that is *robust*, i.e., has satisfactory performance even when reality differs from assumptions. Such a plan would sacrifice optimality when reality is as assumed for reasonable performance over a larger range of situations. We discuss how this proposed refocus follows from the in-principle incompleteness, and common errorfulness, of domain models required to assess the performance of plans. The second issue related to optimality is the degree to which human judgment in planning is subject to a number of biases, all detracting from optimality. The Framing Bias arises from the Bounded Rationality of human cognition. The Transitivity Bias is a result of treating small and large differences in the criteria values as of equal importance. Our analysis leads to design recommendations for DSS's that provide a measure of protection against these biases. The ideas are motivated and illustrated with Army planning examples.

1. INTRODUCTION

Decision support systems (DSS's) for planning generally support activities related to selecting an optimal plan, i.e., a plan that maximizes or minimizes, as appropriate, the performance criteria of interest. We have reported (Josephson, et al 1998; Hillis, et al, 2003) on the development and use of an architecture for generating plans, simulating their performance on multiple, potentially contending criteria, selecting Pareto-optimal subsets, and, finally, applying trade-offs visually to narrow the final decision. Such simulations of course require domain models. In Army planning, such models might be of Blue and Red forces and the environment. With the increasing availability computer-based simulations, military Courses of Action (COA) planning can go beyond the doctrinally required three distinct COAs to accomplish the mission. In principle, planners can generate, simulate, compare and select from a large number of COAs, thus vastly increasing the chances for finding good COAs.

Our focus in this paper is two sets of issues related to optimality in relation to DSS's for planning. The first is a proposal that DSS's for planning should aim to support the planner in generating a plan that is *robust*, i.e., has satisfactory performance even when reality differs from assumptions. Such a plan would sacrifice optimality when reality is as assumed in favor of reasonable performance over a larger range of situations. The second issue is how DSS's might be designed to overcome some of the biases in human decision-making that result from the structure of human cognition and lead to humans making sub-optimal choices.

2. FROM OPTIMALITY TO ROBUSTNESS

Any evaluation of choice alternatives can require use of models. Performance equations, simulations, etc., all make use of models of the relevant part of the world. This reliance on models is a property of all prediction, whether by computer programs or by thinking.

There are two fundamental sources of uncertainty in prediction. The first is the fact the world is stochastic in nature, so at best we can only get a probability distribution of the values for the various criteria of interest. The second problem is more serious: models are in principle incomplete, and even in the dimensions that are included, they may be inaccurate to various degrees. These problems imply that when a COA is deemed to be superior to another COA based on criteria values predicted by model-based simulations, this decision could be wrong, and sometimes grievously so; the rejected COA might in fact be better.

The inevitable gap between model and reality is not an issue for computer simulations alone; as Secretary Rumsfeld's now-famous remark about "unknown unknowns" suggests, simulations that human minds perform are also subject to the same problem. However, humans often, though not always, have a sense of which of their models may be questionable, and. have evolved a set of heuristics by which they sometimes critique their models and simulations to test their robustness. The awareness of this gap between models and reality, and is implications for the quality of decisions, is not widespread in research on simulation-based decision

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE **01 DEC 2008**	2. REPORT TYPE **N/A**	3. DATES COVERED **-**

4. TITLE AND SUBTITLE **Designing Decision Support Systems To Help Avoid Biases & Make Robust Decisions, With Examples From Army Planning**	5a. CONTRACT NUMBER 5b. GRANT NUMBER 5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER 5e. TASK NUMBER 5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Department of Computer & Information Science The Ohio State University Columbus, OH 43210**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S) 11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release, distribution unlimited

13. SUPPLEMENTARY NOTES
See also ADM002187. Proceedings of the Army Science Conference (26th) Held in Orlando, Florida on 1-4 December 2008, The original document contains color images.

14. ABSTRACT

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**UU**	**8**	

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std Z39-18

support. While the builder of a model might have a sense of the strengths and weaknesses of his models that may temper his reliance on the results of planning, today's computer simulations use models composed out of pieces built by different people, without the planner having access to the unarticulated intuition of the model builders.

2.1. Uncertainty and Gaps in the Simulation Models

Aren't gaps and errors in the models are all simply kinds of uncertainty, to be handled by some kind of overarching probabilistic framework? There are many methodological and practical issues in treating all model gaps and errors in a probabilistic framework. The difference between uncertainty and gaps is that in the former, we know what we don't know and have some idea of the structure of missing knowledge. We model the uncertainty by a probability distribution over the structure. In the case of gaps, we either aren't aware of important aspects of reality, or have mistakenly decided that they are not relevant. How does one assign distributions over structures whose existence the modeler is unaware of and has no mathematical model for?

In the case of probabilistically modeled uncertainty, the standard approach is to generate estimates of expected values for the various outcomes of interest by running Monte Carlo simulations. However, as Bankes (Bankes, 2002, Lempert, et al, 2002) points out, difficult decisions about the future cannot be made on the basis of expected values for at least two reasons. The first is that if the outcome measures are correlated, then individual expected values will give a false picture of the future, but this can be taken care of by a more sophisticated stance towards computing the joint expected values. More seriously, however, expected value-based assessments fail to indicate both dangers and opportunities that may lie nearby, and possibilities for driving the future states to avoid dangers and exploit opportunities. In addition to developing such a sense of the decision landscape, what the decision maker (DM) needs is an understanding of how sensitive the outcomes are to the assumptions about the uncertainty, and correspondingly, how to reduce the sensitivity of the desired outcomes to the uncertainties

Both of these issues – protecting against uncertainties in the models used and protecting against the gaps and errors in models – require a shift in point of view from optimality based on expected values to robustness. In order to realize the full potential of the vastly increased search spaces made possible by computer simulation, it is essential that the decision support system empower the planner to explore the plans for robustness of the selected plans.

2.2. Robustness of COA's

There are two different but related robustness concepts:

i. Alternatives A and B have approximately similar expected outcomes, but the one whose ratio of upside/downside is higher is the more robust one[1].

ii. Alternatives A and B have approximately similar expected outcomes, but the one whose outcomes are less sensitive to simulation assumptions and uncertainties is the more robust one.

The robustness we have in mind to cope with model errors is type (ii) above. While everyone would agree, when pressed, that indeed simulations and reality differ, I don't think people in the field are sensitized to the extent to which simulation-based decision comparisons can be problematic, and hence the decision finally chosen may be deeply flawed as a result of this mismatch between reality and simulation. While there is no absolute guarantee against this problem – as mentioned earlier, human thinking itself is based on world models and so limits of simulation apply to human thinking in general – nevertheless, decision support systems ought to empower the DM to evaluate how sensitive a decision alternative is with respect to the specific assumptions in the simulation. We regard development of such systems as the next major challenge for decision support.

This issue is drawing attention in decision theory in general (Bankes, 2002, Lempert, et al, 2002). For example, in (Lempert, et al, 2002) the authors remark, "Reliance by decision makers on formal analytic methodologies can increase susceptibility to surprise as such methods commonly use available information to develop single-point forecasts or probability distributions of future events. In doing so, traditional analyses divert attention from information potentially important to understanding and planning for effects of surprise." They use the term "deep uncertainty" to refer to the situation where the system model is uncertain. They propose an approach that they call robust adaptive planning (RAP), which involves creation of large ensembles of plausible future scenarios. They call a strategy robust if it "performs reasonably well, compared to the alternatives, over a wide range of plausible scenarios." The idea is that, using a user-friendly interface, the DM looks for alternatives that are robust in this sense, rather than optimal in the traditional sense.

There is no universal solution, or even attitude, to the problem of protecting oneself against surprises. According to historian Eric Bergerud, "Bismarck in particular never thought that events could be predicted

[1] Due to Paul K. Davis (personal communication).

with precision. When a policy was pursued a range of outcomes could be expected. The trick was to develop policy where the minimum outcome (today we might call it a worst case scenario) was acceptable. If a triumph ensued, great. If it was something in between, don't die of surprise." This assumes that the goal is to minimize the worst case scenario. Not all missions can be approached this way. There are times when one might wish maximize the best, taking chances of bad outcomes. The decision support system has to support exploring the space of alternatives under a variety of attitudes, conservative or risk-taking, as the commander deems appropriate.

2.3. Exploring Robustness of Plans

A model consists of a set of assumptions of the relevant part of reality. As mentioned, robustness of a plan (or decision, more generally) is a measure of the sensitivity of a plan to the various assumptions *and* to any erroneous or missing relevant knowledge. In principle, a plan cannot be uniformly robust with respect to *all* its assumptions, let alone with respect to knowledge that it doesn't contain. If *nothing* about the world is as assumed, well, the planner is out of luck. So in practice, robustness exploration should be oriented towards making use of islands of relative certainty to critique and protect against less certain parts of models. Further, techniques for testing for robustness with respect to assumptions explicitly incorporated in the model are different from those to test for robustness against missing knowledge.

First, the planner might explore the sensitivity of relevant outcomes to specific assumptions whose reliability he has concerns about. For example, it would be good to know if the likelihood of taking the objective changes much or little for a given COA whether or not it rains, even though the actual uncertainty about the rain might be quite large. Sensitivity analyses of this type are common. Such a sensitivity analysis may be numerical parametric, as in: "How much will the performance change if parameter *p* in the model is off by 10%?" It can be qualitative & structural as in: "In my model, I am assuming that the mullahs would rather give up their support of the Sadr faction than risk extending US presence closer to their border. How does the likelihood of a deal change if I give up this assumption?" Such analyses call for varying the specific model assumptions, simulating over contingencies of each model, and getting a statistical estimate of the range of variability of the performance. In (Chandrasekaran & Goldman, 2006) we report on conducting this kind of exploration in planning, using a visual interface that makes relations between variables quite salient.

Given that an outcome is sensitive to an assumption, the planner might look to see if there are additional elements that impact the sensitivity. In the rain example, the DM might be able to determine that while the probability of achieving the objective varies considerably depending on whether and how much it rains, if subgoal *S* is achieved, then rain doesn't affect the final outcome very much, and that increasing a certain COA parameter to a specific value increases the chances of achieving *S* independent of rain. What the planner has done in this case is an example of identifying the "dangers and opportunities" mentioned earlier and how they can be used to make the plan more robust. As another example, let us consider a large uncertainty in enemy position. The DM might explore how to make a COA robust with respect to this uncertainty. He might see if he can design a COA such that the initial stages of the COA are largely oriented towards getting more information about the location, and based on the information, later stages of the COA can be assigned appropriate parameters. Again, the goal is less to simply obtain estimated values for the outcomes than to explore the COA space to take advantage of opportunities.

Second, the planner might try to identify assumptions whose validity is key to achieving the desired performance. This is a bit harder than the goal we just discussed, since we are not examining the sensitivity of a specific assumption, but trying to *discover* the assumptions, among all assumptions in the model, that may be key to performance. Example: "In the COA for disrupting enemy communication network, what are the variables highly correlated with the outcome variable, and which model assumptions do they especially depend on?" In this case, once we identify the assumptions, we can devote additional empirical resources to verify or test the assumptions for completeness and accuracy. A visual interface such as in (Josephson, et al, 1998) may be useful for this as shown in (Chandrasekaran & Goldman, 2006).

A variation on identifying key assumptions is to perform reverse simulations from desired results so as to identify ranges for the values of key variables that have a bearing on the outcomes of interest. The planner might then be able to identify model assumptions that relate to these variables. For example, an analysis might indicate that capture of an asset cannot be guaranteed if a certain class of events happens, e.g., ambush at a certain location. This in turn could enable the planner to identify the model assumption, "No enemy presence in area A," as crucial to whether or not an ambush is possible. Notice that this is not a case where the planner was especially concerned about the reliability of this assumption in the model. It is that this assumption has been identified as worth double-checking (in comparison with other assumptions) because it has been identified as key to the performance.

Highest in the level of difficulty is exploring the existence of relevant model gaps – missing knowledge, including unknown unknowns. Ultimately, there is no real guarantee that we would be able to identify all the relevant gaps in the models and fix them. But the situation is not hopeless. If we start with the assumption that the model we have is generally good, but potentially missing important components, we can adopt the strategy of identifying regions of certainty and completeness which might then be used to suggest areas of the model of whose completeness we are less confident. Several strategies might be considered.

First, certain features or behaviors might suggest the need for additional domain knowledge of specific kinds. If taking a specific enemy asset turns out to be crucial in achieving plan goals, additional mental and modeling resources might be devoted to buttressing the model with additional knowledge about enemy capacities to defend the asset, or weaknesses in Blue capabilities to take the asset. Second, a reverse simulation – from the goal backwards – and additional reasoning might help. In the situation where the final approach to the objective is identified, the planner might realize that his model has only crude terrain information, quite insufficient to confidently rule out the prospect of an ambush. This kind of exploration is quite focused; our search for potentially relevant missing knowledge is driven by evidence that it is especially related to the outcome. Third, odd or counterintuitive behaviors in the simulation might trigger search not only for incorrect, but missing, knowledge.

An important consideration is that we are not abstractly interested in critiquing or improving the simulation model, but with respect to evaluating one or a small set of COAs, and in response to specific observed behaviors of the simulation. This is likely to focus the need for additional investment in model-building to specific questions, reducing the amount of work involved.

2.4. Increasing the Robustness of plans

Some of the ways in which a plan robustness can be increased have been mentioned earlier in the context of robustness exploration, but it is still useful to bring them together.

If robustness analysis indicates sensitivity to certain assumptions or certain simulation states, the planner has several options. In the case of the assumptions, time and resources permitting, verification efforts can be selectively focused on the specific assumptions. An example q a key simulation state is the discovery that one nay to reach the objective is along a specific avenue of approach. Because of its importance, verification ant modeling efforts can be focused on knowledge related to this avenue of approach. The analysis has moved the epistemic state of the planner from one of anything and everything could be faulty, which makes it impossible to practically verify any component of his models, to one of selectivity in critiquing models.

The planner might also modify the plan in various ways. First, he can try to reduce the sensitivity, e.g., by adding components that restore performance in case reality is different from the assumption. Second, he can organize the plan so that some of the initial steps have the twin goal of achieving some of the plan goals as well provide information about reality, and later steps in the plan take advantage of the information gained. For example, a network disruption plan might have components that attempt to destroy some enemy assets, but the results also give information about enemy defenses. Branches in the plan might be taken depending upon this information: "We are assuming enemy strength is weak here, but if it turns out wrong, we can reinforce from here." Finally, the plan might be buttressed with elements that will help achieve goals even in the absence of relevant knowledge about some aspects of reality. One way to protect a plan against missing weather models is to have components in the plan that ensure success over a variety of weathers.

Based on insights about events on edge of phase transitions, plans may be modified to cause or avoid key events identified to be critically related to outcomes. Example: From an analysis of a simulation, we find that when variable v takes values above say v_0, an outcome variable of interest radically changes in value in an unsatisfactory direction. This then can be the basis for explicitly examining the domain to see which of the phenomena in the domain might have the effect of moving v beyond the critical value and identifying ways in which v can be kept within the desirable range.

Notice that many of these strategies are in fact heuristics used in human planning on an everyday basis. We usually have a sense of trust or lack of trust in our models of the world that we use in our mental simulation, and adopt strategies of early least commitment, information gathering, consistency and reasonable values checking, and so on. However, the challenge for DSS's is to provide computational support to perform such activities. We have outlined above some available approaches, but much remains to be done before a framework is available that integrates multiple types of explorations with plan modification.

3. PROTECTING DECISION MAKER FROM COGNITIVE BIASES

We now shift our focus to how DSS's may be designed to protect a DM against certain biases that seem inherent to human cognition and that often can result in the DM making suboptimal choices. We start by placing the decision-making problem in a larger naturalistic context.

3. 1. Decision- Making in the Wild

It is not surprising that investigations of biases in human cognition use formally stated reasoning problems for which there are clear rational answers. However, everyday decision making – what one might call decision-making in the wild (DMITW) – involves situations that are pre-formal: the DM has first to convert the situation into a problem statement. A man wins $20,000 in a lottery. What should he do: Should he pay off some of his loans? Buy a new car? Repair his house? Perhaps something that hasn't occurred to him yet, but even more essential? Once the alternatives are all laid out, and for each alternative, either a utility (*u*) is specified, or criteria of interest identified and evaluated, the formal problem can be solved as utility maximization or a Pareto-optimal solution for a multi-criterial decision problem.

While mistakes made in solving the formal version are interesting, rarely is the "rationality" in how the DM formulates the problem is discussed. Some readers might remember the old TV ad in which a man makes the choice of something to drink, starts drinking it, sees someone drinking V8, and hits his forehead in mock self-blame, "I could have had a V8." The implication was that he had known about V8. His brain failed him; it missed generating the choice alternative V8. If he had not known about V8, it would not be a cognitive failure. A similar cognitive failure would occur in the case of the lottery winner if he failed to generate an important choice alternative, or mis-assessed the value of an alternative by forgetting to consider a way it could be useful. A cognition that fails to bring to the fore relevant information that is available is not firing on all cylinders; it is being less than optimal. As a result, the decision made is less than optimal.

Conversely, a decision that might appear not quite rational when viewed narrowly within the formal statement of a problem might actually make sense if viewed in the larger context of the work done by cognition. Let us keep in mind the following anecdote about Philosopher Sidney Morgenbesser for later discussion. Once, after finishing dinner, Morgenbesser decided to order dessert. The waitress told him he had two choices: apple pie and blueberry pie. He ordered the apple pie. After a few minutes the waitress returned and said that they also had cherry pie, at which point Morgenbesser said, "In that case I'll have the blueberry pie." This was supposedly a joke, but might the above behavior actually be rational?

These issues turn the focus on the cognitive architecture. Once we understand the limitations of cognition and their origin, we might seek ways in which DSS's might help the DM avoid, or mitigate, these limitations. So-called "biases" are a window into the workings of the architecture. Let's consider two, Framing and Transitivity biases.

3.2. Framing Bias

Consider the following experiment (De Martino, et al, 2006). Subjects in a college campus received a message indicating that they would initially receive £50. Next, they had to choose between a "sure" option and a "gamble" option. The "sure" option was presented in two different ways, or "frames" to different groups of subjects: either "keep £20" (Gain frame) or "lose £30" (Loss frame). The "gamble" option was identical in both frames and presented as a pie chart depicting the probability of winning. In the experiments, the subjects who were presented the "sure" option as a Gain frame chose it over "gambling" at a statistically significant higher rate than those who were presented the "sure" option as a Loss frame, though of course both frames had identical payoffs, and the calculations involved were trivial. How a choice was framed –described – made more difference that we might expect.

3.3 Architecture of Cognition and Problem Solving

Cognition can be analyzed architecturally in many layers. We focus on the problem solving (PS) level of analysis, such as embodied in Soar (Laird, et al, 1987). In this account, the agent has *goals* in his working memory (WM). The presence of goals causes relevant knowledge to be retrieved from long-term memory (LTM) and brought to WM. This knowledge helps set up a *problem space*, which may result in *subgoals*. Some subgoals may fail, in which case other subgoals are pursued. This recursive process goes on, until all the subgoals are achieved, and the main goal is as well. The net effect of all this is that problem solving involves search in problem spaces. The search is open-ended, i.e., the size of the search space cannot be pre-specified. It is also *opportunistic*, being the result of the descriptions of the specific goals, specific world states, and specific pieces of knowledge. Successful searches are cached in LTM for reuse. Heuristics and pre-compiled solutions can reduce search, which is why most every day problems are solved rapidly. One specific kind of heuristic is rapid selection between alternatives based on their types. Additionally, the architecture may have built-in preferences for certain kinds of alternatives over others.

In the example of a person winning a sum of money in the lottery, both the generation of alternatives and their assessments will typically involve such searches. A subgoal of "keep spouse happy" might bring to WM from LTM recent episodes of disharmony, which over a sequence of additional retrievals might eventually suggest remodeling the kitchen as a possible choice to spend the money on. Evaluating the relative merit of one alternative over another, e.g., remodeling kitchen *vs* buying a new car, would typically involve a mental simulation of the consequences of each choice, which again requires repeated access to LTM. Limitations of time, knowledge and memory access all in various ways contribute to the output of these processes not being as complete or accurate as they could be, leading to the concept of the Bounded Rationality of the human as a cognitive agent (Simon, 1972). Which choice is rational may be obvious once the alternatives are generated and multi-criterially evaluated, but as mentioned, that doesn't speak to the rationality of the generation and evaluation of the alternatives themselves.

Let us look a bit more closely at two steps in the process, retrieval from LTM, and choosing between alternatives in the state space. LTM is accessed to solve subgoals – relevant content is brought to WM to setup problem spaces. The associative memory retrieval process works by using the contents of WM as indexes to identify relevant LTM elements. The indexes are provided by the terms in the description in WM of the states, goals and subgoals.

In selecting between alternatives, the architecture has to assign preference values to the alternatives. LTM may directly provide them either as previously compiled values or as heuristics, or the architecture may set up computing preferences as a subgoal. In either case, it is simply another instance of solving for a subgoal, in a manner described in the previous paragraph. However, under certain conditions, the architecture may use hard-wired heuristic preferences; this does not require accessing LTM. It seems a plausible hypothesis that our evolutionary ancestors derived some advantage from a hard-wired choice selection mechanism that, under requirements of rapid decision making, preferred alternatives that evoked gain over those that evoked loss, and alternatives that evoked safety over those that evoked danger.

A feature that is common to the operation of both of these mechanisms, *viz*., retrieval from LTM and selection among alternatives by hard-wired preferences, is that they depend on the descriptive language used, i.e., the terms in which the relevant items in WM are couched. The same state or subgoal described using different terms, but in a logically equivalent way, *may* retrieve different elements of LTM. The hard-wired preference mechanism, similarly, responds to the terms, such as loss or gain, in the description, not to the totality of the meanings, of the alternatives.

3.4. Architectural Explanation: The Framing Bias

We can now look at the Framing Bias example in the light of the architecture and its properties. "Loss of £20," and "Gain of £30" may be equivalent in the experimental setup, but the architectural hard-wired preference mechanism has only the terms "loss" and "gain" to go by. The subjects had limited time, and the hardwired mechanism left its mark in the experimental results.

3.5. Architectural Explanation: Rationality of the Morgenbesser choice

Consider this hypothetical account of Morgenbesser's behavior. When he was told of apple and blueberry pies as choices, the choices and his dining goals evoked two criteria: taste and calorie content. Using additional knowledge and access to his own preferences, suppose he found apple pie to be the Pareto-optimal choice, which of course he then ordered. When the waitress came back with cherry pie as a choice, suppose that this new term evoked an LTM element about pesticide use in cherry farming. Morgenbesser now decides to use pesticide use as a third criterion, and in this new formulation of the decision problem – three criteria -- blueberry pie becomes Pareto-optimal over apple and cherry pies. In this hypothetical account, Morgenbesser was not joking, but making a rational decision.

As this hypothetical shows, the same feature of the architecture that played a role in the generation of the Framing Bias also can potentially increase the overall rationality of DMITW. This is also a caution that issues of rationality in DM are more complex than laboratory studies of formal reasoning problems might indicate.

3.6. Transitivity and Other Biases

Another bias that is relevant is what has been called Arrow's Paradox in theories of voting, and shows up as a kind of transitivity failure in individual decision-making. A group of voters may prefer candidate A to B, B to C, but C to A, when the choices are presented in pairs. The corresponding example in individual decision-making is when the DM is given pairs of decision alternative's, (A,B), (B,C) and (C, A), and prefers A to B, B to C and C to A.

Why does this kind of ostensibly irrational behavior arise? Looking at the voting example, think of each voter as a criterion in a 3-choice, *n*-criterion (where n is the number of voters) MCDM problem. Suppose each voter is asked to distribute 1 unit of preference over the three candidates, each allocation representing a degree of liking of each candidate. For example, if V_1, V_2 and V_3 are voters, suppose V_1 assigns 0.5, 0.3 and 0.2, V_2 assigns 0.3, 0.36 and 0.34, and V_3 assigns 0.3, 0.2 and 0.5, to A, B and C respectively. Straight voting will exhibit Arrow's Paradox, but choosing the candidate with the highest total allocation by the voters does not. Viewing each voter as a criterion in an MCDM problem, the Pareto-optimal decision rule will not exhibit a transitivity error. However, more than one candidate may end up in the Pareto-optimal set. In the specific example above, all three chairs will be Pareto-optimal. The DM will need to call upon additional trade-off preferences to select between them.

3. 7. Decision Support Systems And Framing Effects

Retracing what we said, real word problems don't come fully formalized, and DM's often have to engage in open-ended problem solving to figure out the goals, the alternatives, the criteria to evaluate the alternates, and the criteria values themselves. This activity requires repeated access to information in LTM, and may also make use of architecturally hard-wired preference mechanisms. Rationality in this process consists of making sure that the DM makes full use of all the relevant knowledge, unlike the person who wished he had chosen V8. If our analysis is correct, framing – how goals, states, criteria are described – has an impact on the decision process because retrieval of knowledge from LTM and the architectural preference mechanisms depend on the terms in the description in WM. Thus, framing may reduce the rationality by failure to bring to WM all the relevant knowledge the agent has.

Consider a COA planning problem, which the planner wishes to view as a multi-criterial decision-making (MCDM) problem. To illustrate the issue of framing, let us assume that the problem has been formalized, i.e., the COA alternatives, criteria and their values have been specified. The DSS can help by generating the Pareto-optimal subset. Narrowing the choice calls for injection of additional information by the DM. A common form of additional information is trade-off preferences. For example, he may notice that while alternative A_1 scores higher on criterion C_1, alternative A_2 scores higher on criterion C_2, but A_2's C_1 value is only 5% smaller, while its C_2 value is 20% higher. He may decide that on balance, A_2 is really the better choice. This decision is not simply one of balancing 5% against 20%, since the criteria may be quite incommensurable. Rather, he applies additional domain knowledge that convinces him that the utility of A_2 will be higher than that of A_1. Decision theory assumes the existence of a differential utility to explain the ability of a DM to make a trade–off, but the processes a human DM engages in to access or estimate this differential utility have not been much discussed.

I suggested earlier in the discussion on the example problem of how to spend a modest lottery win that both generating relevant criteria and coming up with relative scores on these criteria for various alternatives often requires a kind of mental simulation of the consequence of a choice. The effectiveness of this simulation, like that of all problem-solving, depends on relevant knowledge being retrieved to set up the problem space and evaluate the states.

Suppose one of the criteria in a COA planning task is "fuel consumption during the operation," and suppose that a DM is trying to apply trade-off preferences between two Pareto-optimal COA's. If COA1 is better than COA2 by 10% on the fuel consumption criterion, and COA2 is better than COA1 on the criterion, "time to complete mission," the DM, by thinking through various scenarios, might conclude that COA2 is better for his purposes. Suppose now that the same fuel criterion is framed differently, "fuel remaining for future mission." The latter can be obtained quite readily from the former, so it is logically equivalent. However, the process of mental simulation might be rather different in this case. The DM might now be reminded of a possible mission immediately following the current one, something that had not explicitly been part of the specifications for the current mission planning task. Now, his simulation of the consequences of the two alternative COA's yields different results. The relative unimportance of a 10% difference in fuel consumption is replaced by the significant potential advantage of a 10% additional fuel for the next mission. Even though future missions were not explicitly part of the formal statement of the current COA problem, the re-framing of the criterion has elicited new knowledge, which in turn has made available to the commander an additional comparison criterion, readiness for next day's mission.

The lesson is not that one or the other framing must be chosen – perhaps the designer should consider including both framings, if then is reason to believe that each will help the DM invoke different relevant knowledge. In general, the DSS should encourage the DM to bring to bear as much of his potentially relevant knowledge as possible.

4. CONCLUDING REMARKS

This paper has focused on two issues related to optimality in planning; and their implications for the design of DSS's. The first was that the real goal of planning should be to come up with plans that are robust over a variety of situations rather than optimal with respect to the assumptions embedded in the models used. The second focused on how the structure of human cognition can explain the Framing Effect that has been observed in human decision-making and that often leads to sub-optimal decisions. Additionally, we noted that much of the discussion of optimality has been in the context of formalized decision situations. However, every day human decision making, what we called "decision-making in the wild," is devoted to formulating – in some ways formalizing – the decision task, i.e. to getting to a place where the traditional notion of rationality or optimality can even be applied. While it is useful for DSS's to help a DM to arrive at an optimal solution for a well-formulated problem, the deeper challenge is to help the DM formulate the problem by bringing to bear his knowledge, understand the decision space, and explore the robustness of his solutions by changing the assumptions.

DSS's should exploit the strengths of the human cognitive architecture, and mitigate its weaknesses. Its strengths are a large, potentially open-ended stove of knowledge, including knowledge about accessing external knowledge; and an ability to bring relevant knowledge, if very slowly, to bear on the task at hand. Its weaknesses are that it is slow, its evolutionarily hard-wired architectural constraints can cause biases, as can the fact that access to knowledge depends on cues in descriptions & may miss relevant knowledge; and has other memory limitations, including a low capacity working memory.

ACKNOWLEDGMENTS

This research was supported by the author's participation in the Advanced Decision Architectures Collaborative Technology Alliance sponsored by the U.S. Army Research Laboratory under Cooperative Agreement DAAD19-01-2-0009. Years of discussions with John Josephson have helped in the evolution of these ideas. Stimulating collaboration with Michael Barnes and Patricia McDermott on decision support systems is acknowledged with appreciation.

REFERENCES

Bankes, S.C., Tools and techniques for developing policies for complex and uncertain systems. PNAS, 2002. 99 (suppl. 3): p. 7263-7266.

Chandrasekaran, B. and Mark Goldman, "Exploring Robustness of Plans for Simulation-Based Course of Action Planning: A Framework and an Example," Proceedings of the 2007 IEEE Symposium on Computational Intelligence in Multicriteria Decision Making (MCDM 2007), April 1-5, Honolulu, HI, Institution of Electrical and Electronic Engineers, ISBN: 1-4244-0698-6, pp. 185-192.

De Martino, Benedetto; Dharshan Kumaran, Ben Seymour, Raymond J. Dolan, Frames, Biases, and Rational Decision-Making in the Human Brain, Science, Vol. 313. no. 5787, pp. 684 - 687, 2006.

Hillis, David; Michael Barnes, Liana Suantak, Jerry Schlabach, B. Chandrasekaran, John R. Josephson, Mark Carroll, "Collaborative Visualization Tools For Courses Of Action (Coa) In Anti-Terrorist Operations: Combining Coevolution And Pareto Optimization," Proceedings of the Army Research Laboratories Collaborative Technology Alliance Symposium, April 29 - May 1, 2003.

Josephson, John R., B. Chandrasekaran, Mark Carroll, Naresh Iyer, Bryon Wasacz, Giorgio Rizzoni, Qingyuan Li, David A. Erb, "An Architecture for Exploring Large Design Spaces," Proc. of the National Conf on AI (AAAI-98), AAAI Press/The MIT Press, pp. 143-150

Kaste, Richard; Janet O'May, Eric Heilman, B. Chandrasekaran and John Josephson, "From Simulation To Insights: Experiments In The Use Of a Multi-Criterial Viewer To Develop Understanding Of The COA Space," Proceedings of the Army Research Laboratories Collaborative Technology Alliance Symposium, April 29 - May 1, 2003.

Laird, J.E., Newell, A., & P.S. Rosenbloom. (1987). Soar: An architecture for general intelligence. Artificial Intelligence, 33, 1-64.

Lempert, R., S. Popper, and S. Bankes, Confronting Surprise. Social Science Computer Review, 2002. 20(4): p. 420-440.

Simon, Herbert A. 1982. *Models of Bounded Rationality, Vols. 1 and 2*. MIT Press.