My ideas divide naturally into two parts and produce three recommendations. The first part focusses on the structure of decision making under uncertainty. I strongly embrace a constructivist viewpoint, and this leads to some new ideas about norms for "good" decisions. Two concrete recommendations are suggested by these ideas.

1. (a) We should focus on the art of structuring decision analyses, rather than on formal models for decision making. In particular, we need to teach broadly effective heuristics for framing decision problems.

1. (b) We need to develop and demonstrate capacity for risk assessment for complex plans, in which physical, biological and social uncertainties all enter.

The second part, which is much less well thought through, considers decision making in games against rational malevolent opponents (various "axes of evil" come to mind). I depart from standard game theory by viewing goals as labile, rather than fixed. This leads to the third recommendation:

2. One should deal with rational opponents by seeking to identify and activate shared goals, thus perhaps turning enemies into something else—collaborators, at least, if not friends.

1. Individual/Cooperative Planning under Uncertainty

I came to decision theory by way of analysis of basic measurement structures. In my early work (e.g., chapters 5 and 8 of Foundations of Measurement) I assumed that norms for subjective probability and individual choice could be formulated along the lines of Savage's axioms. (Duncan Luce and I developed a sort of hybrid between von Neumann & Morgenstern and Savage.) The fact that actual human decisions are not well described by such axioms just shows up human imperfections. To this way of thinking, that mankind survives at all is a bit puzzling.

My viewpoint now is sharply different, but the interest in fundamental underlying structure remains unchanged.

"Preferences" are constructed in the choice setting. Goals are activated, and sometimes novel goals are adopted, within the immediate context of the decision problem; attributes of an outcome are weighted contingent on the overall choice setting; and rules of behavior that seem to fit the problem situation, or past similar decision problems may be sought in memory. The closest psychological approximation to utility lies in "gut feelings" or emotional reactions; however, where there are many goals and much uncertainty, complex cognitive processes may be involved in assigning a structure to the decision problem, and the assigned structure may determine the emotional reactions. The processes used to solve decision problems thus do not lead to a transitive ordering across all potential options. Nor is there any reason to assert there should be such an ordering that underlies choice. We cannot postulate a transitive ordering for social choice (Arrow's Theorem); and the combination problem for an individual, who must try to integrate conflicting goals and

disparate ways of structuring each decision problem is just as complex as for a society trying to integrate conflicting preferences of individuals. In short, transitive ordering is not a norm for decision making, thus, maximization (of SEU, "value", or any other scalar quantity) is also not a norm.

The abandonment of maximization as a norm for decision making does not, however, mean that there is no guidance for making better and worse decisions, and no role for decision analysis. On the contrary: most convincing decision analyses do not require actual calculation of SEU, rather, the important analysis consists of identifying goals, options, tradeoffs, and uncertain events in such a way that one can first contemplate a breakdown in terms of utilities and subjective probabilities. Often, the right qualitative formulation of the problem makes the right solution obvious; and where this fails, finer-grained calculation is unlikely to help. The proper norms for decision making are the heuristics that help one to set up the problem analytically. These can be taught on an abstract level—in fact, decision analysts are people who have more or less mastered this teaching. The question is, how easily and how widely can these heuristics of good decision making be taught?

Another implication of the change in norms concerns the practice of risk assessment. Much work may go into the construction of a family of models that permits fairly valid empirical assessment of a risk. Once that is complete, then what? The usual story is that the resulting probability estimate is ignored. One may as well not bother.

If we think of humans as imperfect Bayesian decision makers, then it makes some sense to replace their wildly inappropriate probability guesses with the results of careful modelling. However, if decision are based on a variety of complex structuring processes, identifying a rule that fits, or a previous decision problem that partially matches the present one, how do we expect people to make use of particular risk assessments? It seems to me that risk assessment should be guided by prior understanding of how the decision maker structures a particular decision problem. Only if one knows what risks the decision maker is trying to assess, can one know what models are needed.

One of the common types of chances that people focus on is the chance that a particular plan will succeed. (Exactly what constitutes "success" needs to be elicited.) That chance, and the cost estimate for the plan, determine how it fares in competition with alternatives. Success of a plan, however, usually depends on many things working right, or at least well enough. Some of these may be easy to model, others hard. We need to develop the capacity to do risk assessment for complex plans, integrating physical, biological and social systems into our modelling.

2. Conflict with Rational, Malevolent Adversaries

There seems to be a widespread feeling that experts on risk assessment should have something to say about the world that we woke up to on September 11.

However, risk of terrorist attack is fundamentally different from risk of earthquake. The earthquake is uncertain, but seismologists have come far in their ability to estimate the probabilities of various possible magnitudes, locations, and other aspects; over a human generation, the probabilities for a particular region may remain constant or may evolve in a lawful fashion. One can try to plan for earthquakes—California works hard at this—and one may hope to assess the costs and the likelihood of success of one's plans, as recommended in Part 1. But planning carefully and well against a terrorist attack simply makes it much less likely that the attacking strategy for which one has prepared will actually be used. The rational and malevolent opponent will try to devise a novel strategy for which we are not prepared. Planning carefully and well against all possible forms of terrorist attack—including ones not yet invented—seems to me not a well-defined or feasible task.

Those who have some responsibility for planning—government officials, experts on terrorism and on risk assessment, and others—feel pressure to do something, but also to give the appearance of doing something. Based on the very limited observations I am in position to make, I judge that the effort devoted to the appearance of action—to reassurance, if you will—has been larger than that devoted to actually reducing danger. If protective measures merely serve to stimulate terrorist creativity, then perhaps actually reduction of danger cannot be expected.

Playing cutthroat games for low stakes can be fun, depending on one's taste. Playing for life or death is to be avoided, if possible. The most obvious strategy for avoiding a highly competitive game is to convert it into a cooperative one. This is a lesson of the cartel, learned early and well by American businessmen.

The lability of goals—the fact that new ones may be suggested by a choice context and may be adopted in order to resolve an intrapersonal conflict—suggests that conflict between parties may also sometimes be resolvable through identifying goals that could be shared and adopting them. "Atoms for peace" or the (aborted) plans for shared development of water resources in Jordan and Israel are salient examples in which this strategy has been suggested and pursued in international conflict.

The advice that decision theorists should give to the world post September 11 is that we are playing a game whose equilibria are very unfavorable; we need to change the game rather than search for an unfavorable minimax. We will do far better with fewer enemies; and making some marginal enemies into collaborators, by developing shared goals, is perhaps the only way to change the game we are playing.