4

When Lives Are in Your Hands: Dilemmas of the Societal Decision Maker

SARAH LICHTENSTEIN, ROBIN GREGORY,

PAUL SLOVIC, AND WILLEM A. WAGENAAR

A societal decision maker (SDM) is a person who makes risky decisions on behalf of others. Most of the time, such decisions should be based on the wishes and beliefs of the affected people. This chapter explores a few cases in which it could be argued that the SDM, in making the decision, should in good conscience disregard the desires or beliefs of the affected people. Several simplifying assumptions are made: the SDM uses decision analysis in making the decision; the affected people speak with one voice on the matter under dispute; the SDM and the people disagree on an issue vital to the decision problem; the SDM cannot delay the decision or otherwise avoid the disagreement; the SDM is motivated only to make the right decision; and the SDM can effectuate an unpopular decision. In this context, the following dilemmas are discussed. What if the people object to the use of decision analysis? What if the people reject the axioms of decision analysis? What risk attitude should the SDM adopt? What concerns should be included in the analysis? What if people are misinformed? What if individual and societal perspectives differ? Do people really want what they say they want? For some of these questions, we argue that the SDM should make decisions against the wishes of the people; for others, we are not sure how to resolve the dispute.

CONSIDER THE PERSON whose job it is to make risky decisions on behalf of others; that is, decisions with outcomes affecting other people, perhaps many other people, to a far greater degree than the outcomes affect the decision maker. We call such a person a societal decision maker, or SDM. Our society has many such people, making decisions about energy options, drug-testing standards, genetic research, automobile emissions, and the like. Such problems are challenging social issues because the stakes are high and the value issues complicated.

In seeking assistance for these complex problems, it is natural that SDMs would turn to research in decision making for assistance. As people who earn our living by producing and disseminating research results, we would hardly want things to be otherwise.

There is a rich collection of advice available. Over the past several decades, research on decision making has developed a considerable body of knowledge about how risky decisions are made and how they can be improved (e.g., Kahneman, Slovic, and Tversky 1982; Slovic, Lichtenstein, and Fischhoff 1988). Although studies from many disciplines have made important contributions to the topic, in this chapter we focus on contributions from the psychological study of decision making and risk perception. Research in behavioral decision theory examines how individuals and groups actually do make decisions, in contrast to normative prescriptions about how such decisions ought to be made (Einhorn and Hogarth 1981; Slovic, Fischhoff, and Lichtenstein 1977). Studies of public perceptions of risk examine people's expressed opinions about hazardous activities or technologies and attempt to determine how information about uncertain outcomes should be communicated among decision makers, lay people, and technical experts (Slovic 1987).

We have long argued (Fischhoff et al. 1981), and still believe, that policymakers need to take account of public values and perceptions in societal decisions about risk. Failure to do so entails several dangers.

First, the people may know something that the experts are missing. For example, one finding from risk-perception studies (Slovic, Fischhoff, and Lichtenstein 1979) is that experts tend to assess the risks of a technological option in terms of its expected fatalities and injuries whereas lay people typically use a broader evaluation scheme that includes the voluntariness of exposure to a risk, the degree to which it is understood scientifically, and a number of other psychological factors. It is not hard to argue that such richer, more comprehensive views express important criteria that should be included in the decision process.

Second, the disregard of public opinion may result in a decision that cannot be implemented because of outspoken public opposition. Even if it is possible to force the decision on an unwilling public, the outcomes may be quite different from, and worse than, anticipated. Did politicians in 1917 predict that the adoption of the Prohibition amendment to the U.S. Constitution would lead to a law enforcement crisis and the rise of a wealthy, well-armed, and organized underworld?

Finally, the SDM courts trouble in saying, "I know better than you," and, "I'm doing it for your own good." We resented such claims made by our parents when we were young; as adults, we resent them even more. Our society is structured, by and large, democratically; we would not willingly live under authoritarian rule, even the benevolent regime of a technological elite.

But when SDMs look to the research literature for guidance in representing the public's wishes in their decisions, another danger emerges. Experimental results, naively applied to complex social issues, may result in poor social decisions. In order to know which of the experimental findings are relevant, and when, and under what conditions, SDMs need to ask probing questions of the findings and their interpretation. For example, it is important to know why

the public and experts view risks differently. Surely the policy implications are different if the public is misinformed about the facts than if the experts are defining the problem too narrowly.

Thus, our starting point is that incorporating public input will usually lead to better social decisions regarding risk, but we do not believe that this happy outcome is guaranteed. We will try, in this chapter, to suggest when and why SDMs might, in good conscience, go against public opinion in order to make a better social decision.

This is a chapter about "oughts" and "shoulds," about ethical dilemmas. We find ourselves with more opinions than skills as ethicists. We present our views forcefully here in the hope that others, whether agreeing or disagreeing with us, will be moved to contribute to ongoing discussions about the proper role of public input; we believe that these issues contain strong implications for the management of risks in our society. We warmly dedicate this chapter to the memory of Hillel Einhorn, who we believe would have approved of our attempts while arguing with us heartily.

SETTING THE STAGE

Our prototypical decision maker is a regulator who forms one link in a chain of command but whose opinions play a decisive role in the regulatory process. Such people do not stand alone in the world of social dilemmas, forced to approach each decision with only their natural instincts and ethics. Instead, we assume that the SDM will use what has been learned about decision making, as an aid to structuring the decision problem and as a guide to evaluating social needs and the probable consequences of alternative actions.

Although a range of prescriptive approaches is available to the decision maker, we advocate the use of decision analysis. Our reasons are straightforward. First, decision analysis is explicit about its assumptions. It makes clear what is being done and what could be done differently. Second, at its core is subjective expected utility theory, with all its prescriptive power. Third, the multiplicity of objectives and consequences that characterize many decision problems can be incorporated into the decision-analytic framework. Unlike cost-benefit analysis, in which economic attributes so often dominate, decision analysis can fuse "hard, objective" knowledge with "soft, subjective" knowledge and values, thereby encouraging the richest, most responsive characterization of problems.

Our view of the efficacy of decision analysis is perhaps broader than that of authors who see the technique as a tool to help express and organize one's already well-formed beliefs and preferences. In contrast, we suspect that in many situations beliefs and preferences are vague, ill formed, or even nonexistent (Fischhoff, Slovic, and Lichtenstein 1980). In working through a decision analysis, then, a user may be not reporting but actually creating beliefs and preferences. This process produces a better understanding of the problem

and may lead to new insights. Ideally, when the decision analysis is finished, the right decision seems obvious. The entire decision analysis then provides an elaborate justification for the decision.

Decision analysis is, we believe, the best method for making complex decisions. But it is not perfect; it cannot be expected to capture every possible aspect of the problem and resolve all issues. Thus, one could complete an analysis and then decide to go against its prescriptions. We do not know the conditions under which this might occur. If it does occur, one at least can feel comforted that the decision analysis has illuminated the issues and provided a fuller understanding of the situation.

We recognize that, particularly in the context of the broadly based social problems with which we are concerned, it is not always true that one single individual has sole decision-making power. But this does not materially affect the nature of the dilemmas we address. We therefore assume such sole authority. If, instead, decisions are made in a group (e.g., by a national parliament or by the three-member U.S. Nuclear Regulatory Commission), the reader can suppose that the others' opinions are evenly split and that one person holds the deciding vote. Of course, elected representatives such as members of the U.S. Congress face a more complicated situation in which current decisions (votes) need to be played off against the desire to ensure future coalitions, protect the good of the political party, and so forth. Nevertheless, a basic characteristic of such decision environments is that individuals are expected to act on behalf of, and for the good of, others.

As a further simplification, we here assume that all members of the affected public agree with one another with regard to the critical aspects of the problem under consideration. This is a useful dodge because it means we can use terms like *the public* as if there were only a single public (and we will say *the experts* as if this group, too, were undivided). However, as anyone with experience in public participation or communication knows, the public rarely speaks with one voice; a key to successful interaction is to identify and understand some of the major distinctions that mark the central actors or stakeholder groups (see Edwards and von Winterfeldt 1987). This is an important topic but one that will not be discussed in this chapter.

Furthermore, we assume that the SDM and the affected group disagree on an issue vital to the decision problem. Thus, for whatever reason, what the affected group says it wants is not the same as what the SDM wants or believes or thinks is right. Moreover, the disagreement is consequential. Showing, via sensitivity analysis, that the disagreement would not materially affect the decision would let our SDM off the hook; we thus exclude this possibility.

Let us continue to press our SDM by taking away other possibilities that might forestall the decision or enable the SDM to avoid squarely facing the problem at hand. We do this not because these possibilities are unrealistic but because they provide convenient excuses to ignore the dilemmas that we want to address. We list four such "outs" that we will not permit.

1. The SDM cannot plead limited resources or limited knowledge and just wait until more or better information comes in. At any rate, there is no reason to believe that more resources or knowledge would necessarily bring about agreement between the SDM and the affected group.

2. The SDM is not motivated by other personal or cultural factors such as, "If I decide that way I'll lose my job," or, "Everyone else in my position has always done it this way." Instead, we assume that the entire motivation for the SDM is to make the right decision.

3. The SDM is able to go against the wishes of the affected public without producing so much objection that the goal of the decision is itself threatened or ruined. Clearly, if this were not true, the SDM should take the objections into account. More interesting, to us, is the question of what SDMs should do when it is possible to implement an unpopular decision.

4. The SDM cannot automatically justify a course of action by simply going along with the desires of the affected group. We consider this a fink out of the lowest order. Without the urgency of a central conflict between the group and the SDM, there is probably no need for the decision maker at all and certainly no need for this chapter.

At this point, we have a competent SDM who is responsible for making an important decision using decision analysis; the SDM faces a major dilemma involving a discrepancy between the unanimous wishes of the affected group and the beliefs of the decision maker (or, as often happens, the advice of the experts). We turn now to consider several such dilemmas.

DILEMMAS OF THE SDM

What If People Object to the Use of Decision Analysis?

People may oppose decision analysis because they believe that formal structuring and the codification of values as numbers can never capture the intuitive essence of a complex problem. We reject this view and urge our SDM to reject it. Although a decision analysis can never be complete, the process of constructing it forces the decision maker to think about the problem in an orderly way, which is a virtue for any difficult decision.

People may also object to decision analysis because it calls for explicit trade-offs between attributes such as lives and money (MacGregor and Slovic 1986). They may object to such trade-offs because they believe that the government has deep pockets ("Spend whatever money it takes to save the lives") or because they find such trade-offs morally repugnant. For this dilemma, we have firm advice for the SDM. Go against the public's wishes. We believe that social decisions are made under real resource constraints, that such constraints tacitly imply trade-offs, and that it is better to make the trade-offs explicit. Decision analysis did not invent these trade-offs; if both the approach and its advocates were to vanish tomorrow, the trade-offs would still remain.

This is quite a strong position. However, we declare an even stronger one.

Because awkward trade-offs (between lives and dollars, pollution and jobs, or the safety of the old and the young) are both so prevalent and so difficult to face, it is the responsibility of the SDM to make them explicit and point them out to the public. Being attentive to public opinion does not require fooling people into thinking that tough decisions need not be made.

What If People Reject the Axioms?

There is plenty of evidence in the research literature showing that utility theory is not a good descriptive theory. Some of its descriptive failures may be traced, directly (e.g., Tversky 1969) or indirectly (e.g., Slovic and Lichtenstein 1983), to violations of one or more axioms. One response to these findings, currently much in vogue, is to invent new versions of utility theory that omit the offending axiom in an effort to make utility theory more descriptively accurate (see Machina 1987). Although we admire much of this work for its creativity, we are not convinced that our SDMs should adopt this new approach.

One reason for sticking with the axioms is that, when they are presented to subjects in transparent form, they are less likely to be violated than when they are not made explicit (Tversky and Kahneman 1986). One view of the experiments showing violations of rationality is that they are designed to study our customary intuitions and decision habits to gain insight into how people think but that the underlying structure of the tasks is sometimes not obvious, even tricky. Tversky's (1969) subjects, for example, vigorously proclaimed that they neither violated nor wanted to violate transitivity, although they did make intransitive choices. It is possible that, if the public were to apply as much careful thought to, say, Allais's paradox as did Savage (1972, 101–3), they, like he, might end in accepting the sure-thing principle despite strong initial intuitions against it.

It might be wrong to be too optimistic, however, about the public's eventual acceptance of the axioms. People may ascribe, instead, to Samuelson's (1950) suggestion that they should "satisfy their preferences and let the axioms satisfy themselves." For example, Lichtenstein interviewed some preference reversal subjects extensively (Lichtenstein and Slovic 1971). One subject, when he came to understand that the pattern of his responses could be used as a money pump against him, readily agreed to change those responses to make them consistent. But he then plaintively noted that he had been instructed, at the outset, to report his true preferences; he still felt that his original responses were faithful to that instruction. And the majority of Slovic and Tversky's (1974) subjects were unconvinced by Savage's analysis of the Allais problem.

One or more axioms of utility theory are violated whenever choices are made on the basis of the avoidance of ambiguity (Ellsberg 1961). Ambiguity and vagueness about probabilities, even when expressed as second-order probabilities, are formally irrelevant in decision analysis but seem to play a large role in personal beliefs. There is also increasing evidence that the avoid-

ance of ambiguity is affecting public policy. For example, insurers are unwilling to cover chemical and waste-processing firms because it is so difficult to specify the distribution of anticipated claims (Kunreuther 1987). Einhorn and Hogarth (1987) have suggested that ambiguity aversion may explain why some technologies are feared more than their first-order probabilities of failure or accident warrant.

Dislike of ambiguity may be a pure preference, like the preference for a "fair" coin over a coin biased in some unknown way. However, it is possible that what appears to be an aversion to ambiguity really represents other concerns. For example, the public may legitimately recognize that the experts have been wrong before—while being highly confident in their wrong beliefs (Henrion and Fischhoff 1986). Thus, the public might quite rightly trust probabilities based on abundant data more than probabilities assessed by experts lacking such evidence. It would be appropriate for the SDM to include such concerns in the analysis, perhaps by finding some attribute, such as dread or worry, that could serve as a proxy variable for ambiguity. Researchers could help SDMs in this arena by learning more about the structure and correlates of ambiguity aversion.

The jury is still out on the public's eventual acceptance of the axioms of decision theory. But we have a more personal reason for suggesting that SDMs stick to the axioms. We find that, even when we are comfortable with our own violations, we want to hold our SDMs to a higher standard of rational thought. Consider, for example, the theory proposed by Loomes and Sugden (1982), in which transitivity is given up to accommodate feelings of regret. In our own lives, we may make decisions in a way that minimizes regret even if so doing violates transitivity; we might even believe that this is a good thing for us to do. But we do not feel an equivalent need to protect our SDMs from regret in the decisions that they make on our behalf. We, not they, must suffer the consequences, so to hell with their tender feelings. As another example, the SDM might be tempted to prefer options with low ambiguity because a decision made on that basis can be more easily justified and defended to the public (Curley, Yates, and Abrams 1986). We urge our SDM to resist this temptation and make the best decision without regard for later difficulties in justifying it.

In summary, we are inclined to believe that the SDM should stick to the axioms, but we acknowledge that our case for doing so is not strong. Perhaps our best defense of the axioms is a pragmatic one. The SDM who accepts the axioms can get on with analyzing the problem at hand, comforted by the central theorem of decision theory, which states that, when the axioms are accepted, utilities can be measured.

What Risk Attitude Should the SDM Adopt?

Consider a decision situation with one or more attributes having clear, explicit numeraires, such as lives or money. There is plenty of evidence that

people's utility curves for such attributes are often risk averse in the domain of gains and risk seeking for losses. Sometimes risk attitudes may apparently reverse with a different response-elicitation method (Hershey and Schoemaker 1985) or a different wording (Tversky and Kahneman 1981). But we suspect that risk attitudes cannot be attributed solely to response biases; our preference for a sure outcome of, say, $50 over a 50-50 gamble paying either $200 or −$100 is likely echoed in the preferences of many.

The SDM might want to choose a utility function that has the following characteristics: bad outcomes are minimized, catastrophic outcomes are avoided, and the risk of harm is spread equitably among the affected people. Keeney (1980) has explicated the form of a utility function that has each of these characteristics.

Keeney started with n individuals, some or all of whom are at risk of death. If the risk of death to any one individual is independent of the risk to any other, the whole situation can be represented by a vector of probabilities (p_1, p_2, . . . , p_n) showing the probability, p_i, of death for each individual (more complex lotteries are used to show dependent risks).

These characterizations of risky situations can then be translated into a probability distribution of total number of fatalities, x. If all the deaths are equally bad, the central question is to find the appropriate utility function over x, the number of fatalities. Keeney explores three possibilities.

1. A *risk-neutral* utility function is the only function that will minimize expected fatalities.

2. A *risk-prone* utility function is the only function that is consistent with a preference for risk equity. Keeney defines risk equity as a preference for

$$(p_1, \ldots, p_i, \ldots, p_j, \ldots, p_n)$$

over

$$(p_1, \ldots, p_i + e, \ldots, p_j - e, \ldots, p_n)$$

when the difference between p_i and p_j is less than the difference between $p_i + e$ and $p_j - e$. It follows that the most equitable vector of probabilities is the one in which all the p's are equal. That is, a preference for risk equity is a preference for equally spread risk.

3. A *risk-averse* utility function is the only function that is consistent with a preference for catastrophe avoidance, defined as a preference for probability π of x fatalities over a probability π' of x' fatalities for any $x < x'$ such that $\pi x = \pi' x'$.

One can, thus, achieve only one of the apparently laudable but inconsistent goals of fatality minimization, risk equity, and catastrophe avoidance. Which one should the SDM adopt? We advocate an extreme position, that the SDM

should always use a risk-neutral utility function for lives—and for money—for the following reasons.

The SDM surely realizes that the currently considered risk is only one of many. From that perspective, the concept of risk equity gets fuzzy. Is it the SDM's goal, under risk equity, to ensure that all citizens of the country have an equal risk of dying when all regulated activities are taken into account? Short of that absurdity, it is unclear how to choose the number of people and the number of different risks over which risk equity should be sought. If risk equity is difficult to define, risk proneness loses its appeal.

As for catastrophe avoidance, we are ethically uncomfortable with the position that it is better to avoid 10 deaths from a single accident than to avoid one death in each of 10 separate accidents, other things being equal. It seems more compelling to us that it is the moral obligation of the SDM to save as many lives as possible; that implies risk neutrality.

Where money is concerned, a broader perspective will reveal that the current decision is just one of a multitude of demands for expenditures. Any risk attitude other than risk neutrality increases the expected costs and will, in the long run, buy us less. Moreover, risk aversion is sometimes motivated by fears of a large loss. But a possible loss that would threaten the budget of any one individual can be more easily absorbed by the larger budget of a federal agency.

A simplified but not entirely unrealistic conception of many safety decisions is that they involve trade-offs between lives and money. For example, a regulatory agency may have a fixed budget with which to research and regulate a wide variety of risks to life. Alternatively, the agency may realize that there is a limit to the number of different safety regulations that can be imposed on a particular risky industry—too many regulations would bankrupt or collapse the industry. Here, too, the wish to maximize the number of lives saved per dollar spent is consistent only with risk neutrality.

A catastrophic accident, say 100 deaths and 1,000 injuries in an area populated by 100,000 persons, would not only engender great suffering and grief but also place severe strains on community resources such as hospitals and morgues. But these real costs should not be captured by using risk-averse utility functions for death and injury. Instead, they should be included directly as additional attributes in the decision analysis. As to the suffering and grief, we acknowledge that they might be greater than they would be were such deaths and injuries spread out over time and place. But the SDM knows that individual losses are constantly occuring in numbers greater than this—without accompanying newspaper headlines and live television coverage. We believe that the SDM should accept the responsibility of preserving life in small chunks as well as big chunks; risk neutrality fosters this goal.

A concept related to risk equity is called benefit/risk equity. This is the idea, explored experimentally by Keller and Sarin (1988), that it is fair for

those who receive more of the benefits to accept more of the risks. As far as we know, this idea is not inconsistent with risk neutrality over number of fatalities, as discussed above.

A further distinction is between the risk to known and unknown lives. We do not know how to advise SDMs regarding the relative importance of these kinds of deaths. We would not want to live in a culture in which the SDM decides not to make a rescue attempt to save a (known) miner now trapped in a cave-in because the rescue money would save more (expected but unknown) lives if spent tomorrow on mine safety. Here may be a case (not well delineated, we admit) in which the SDM should depart from the prescriptions of decision analysis. We fear a possible result of slavish obedience to the analysis: a host of coldhearted regulators disregarding feelings and emotions.

Which Concerns Should Be Included in the Analysis?

We believe that all the attributes that are important to the affected people should be included in the analysis (assuming that the decision is sensitive to those attributes). However, we see some exceptions to this rule.

Attributes that are illegal or for which there is a clear societal consensus of moral objection should be excluded. For example, residents of a community may be opposed to a federally funded building project because they fear that racially mixed construction and operations crews would settle in their area. Such prejudice has been rejected by our society at large and thus should not be given formal standing in the evaluation process.

Other concerns, while valid, may fall outside the SDM's mandate. For example, an SDM charged with making decisions about safety regulations for an existing nuclear reactor should not consider whether the country is becoming too reliant on nuclear power. In this case, the proposed attribute lies outside the bounds of the legislation that grants authority to the SDM; it is a broad social issue that should be played out in some other arena (probably Congress).

A special case of an attribute that may have to be disregarded by the SDM is "not in my backyard," the syndrome whereby each community simultaneously acknowledges the need for a risky facility but refuses to serve as its host. One approach in such cases is for the SDM to emphasize the costs of noncooperation while offering rewards for cooperation (Kunreuther and Kleindorfer 1986). If no community comes through with an acceptance, then the SDM may need to override the narrowly focused interests of any single community in order to serve the broader interests of the region as a whole.

What If People Are Misinformed?

People may disagree with the SDM because they are misinformed about some aspect of a decision. The obvious solution, to inform them, is attractive if it works, but sometimes that may not be feasible.

Communicating the facts may be impossible if the facts are highly complex and technical. For example, scientists are beginning to understand the com-

plex chain of events linking exposure to low-frequency electromagnetic fields (e.g., from electric blankets or can openers) and potentially dangerous changes in human cell structures (Morgan et al. 1987). Without extensive training in areas such as chemistry and cell biology, however, people may be incapable of understanding the risks involved. Thus, resentment may result if the scientists end up saying, "Just take our word for it, folks, this is a risk that should be reduced."

In addition, people may not trust the SDM or the agency represented by the SDM. Thus, whatever they are told is viewed as a probable lie and heavily discounted. In such cases, assurances of trustworthiness are likely to fall on deaf ears. We can suggest that such an agency first eliminate any real causes for distrust and then embark on a long-term strategy of reassurance. But that will not help the SDM in the immediate situation. We know of no remedy for this problem.

People are not very good intuitive scientists. Their views of facts and of possibilities may be inappropriately influenced by cognitive biases and heuristics (Kahneman, Slovic, and Tversky 1982), leading them to a mistaken view of the world. Moreover, such effects can be highly resistant to change (Fischhoff 1982).

Of course, experts and the SDM are also subject to the influence of cognitive biases and heuristics. That is one good reason for using a formal decision-aiding technique like decision analysis. It is far more difficult to disregard relevant base-rate information when one is reminded of it by Bayes's theorem. Similarly, SDMs, like other people, may fail to appreciate the burgeoning growth of exponential functions (Wagenaar and Sagaria 1975). But statisticians (or hand-held calculators) can remedy this failing.

In keeping with our insistence, in this chapter, of holding the SDM to the most difficult case, let us suppose that a vivid, dramatic, but minor accident has artificially elevated (via the availability heuristic) the public's assessment of the probability of disaster for liquid natural gas storage facilities. The accident and resulting change in public views comes just as the SDM is about to issue a new safety regulation governing liquid natural gas storage. The public demands stringency. The experts have studied the situation extensively and have assured the SDM that the regulation need not be as severe as the public demands; the SDM believes the experts' analysis. Newspaper coverage and pubic hearings fail to change the public's beliefs. Let us not suppose that the SDM can say, "Well, it's only money," and issue the stronger regulation. Instead, suppose that the stronger regulation, because of its greater cost, will prevent the implementation of an additional forthcoming safety regulation, one that is expected to save many lives.

What should the SDM do? A dyed-in-the-wool decision analyst may argue that the only thing to do is to analyze this larger problem, looking at both regulations and including such attributes as the effects of the anger and mis-

trust engendered if the SDM goes against public opinion. This seems sensible, if perhaps prohibitively expensive. But is it really ethical to incorporate, into one's decisions, the public's wrongheadedness? We are not sure.

It is not only facts and possibilities that the public may be misinformed about. They may also be misinformed as to what risk-reduction policies can accomplish. Although specific risks can be reduced, at some cost, Keeney (1988) has noted that the possibility of zero risk is an illusion. Strictly speaking, risk cannot be eliminated; it can only be transferred or delayed. The contrary view, that perfect safety is an attainable goal, can be a source of conflict between the SDM and the affected group, one which the SDM may have little ability to change. Again, we are left with an unanswered question. Is it right for SDMs to alter their decisions in order to appear to be pursuing the illusion of perfect safety?

What If Individual and Societal Perspectives Differ?

In many cases, both the societal and the individual perspective are valid; both can be included in the analysis. Okrent (1987), for example, proposed standards for nuclear power that place separate limits on the risk to whole populations and the risk to any one individual. But what about seat belts? The SDM sees that the costs and consequences of auto accidents extend beyond the immediate victims. Moreover, requiring seat-belt usage nationwide would save hundreds of lives annually. Such savings would likely outweigh the inconvenience of wearing seat belts and may even outweigh the loss of personal freedom entailed by a national law. Should the SDM also include in the analysis the fact that, from an individual perspective, the risk of a fatal accident seems too small to warrant protective action?

Do People Really Want What They Say They Want?

The best way to discover the preferences of the affected group is to ask; answers are generally forthcoming to questions about values. But do people really want what they say they want? Should assertions of preference always be taken at face value?

These are difficult questions on ethical as well as technical grounds. Certainly, it is dangerous to say that a decision maker might know better than someone else what that person really wants; we are rightfully suspicious of people disregarding our wishes "for our own good." Nevertheless, we explore below several skeptical cautions.

First, people's willingness to answer the questions we researchers put to them does not always ensure that they are expressing well-understood and deeply held values, particularly when the questions concern a rare and emotionally laden event. Consider, for example, a patient facing the choice of a radical mastectomy. There may be no way for her to understand the multiplicity of pain, psychological distress, and physical impairment that results from such surgery. How bad is it, really, to be permanently unable to lift one's arm above one's shoulder? We who can do that now do not know.

Second, the way a question is worded can affect the answer. Difficulties may arise especially with emotionally laden terms. With such terms, people may be so impelled by symbolic meanings that they are blinded to differences in how others are using the term. For example, *compensation* may mean *bribery* to some people; this connotation leads them to different values than if the same concept were thought of as *fair payment*.

Other examples of the effect of question framing are myriad (Hogarth 1982). The most troubling ones are those for which it cannot be said that further reflection will reveal the true preference. Some of these involve risk aversion (e.g., Tversky and Kahneman 1981; Hershey and Schoemaker 1985); we have argued above that, in these cases, the SDM should take neither frame, using risk neutrality instead. Other examples cannot be so easily resolved. For example, McNeil et al. (1982) showed that formally equivalent but different ways of presenting information about the probability distribution, over time, of death from a particular kind of cancer led both doctors and patients to have different preferences for surgery versus radiation treatment. Yet neither description can be viewed as a more natural or "correct" way to present probabilistic information than the other.

Third, it is sometimes problematic for a researcher to detect, accurately, the variable that the respondents were attending to. A discouraging cautionary tale has been told by Wagenaar, Keren, and Lichtenstein (1988). They replicated an experiment by Hammerton, Jones-Lee, and Abbott (1982) in which respondents were given a cover story involving an infectious disease on an island with 100 residents and a choice between one of the islanders dying for sure versus each islander having a 1 in 100 chance of dying. Hammerton, Jones-Lee, and Abbott reported that a large majority of subjects preferred the one sure death, which they interpreted as evidence of risk aversion. Wagenaar, Keren, and Lichtenstein found significant differences with apparently minor changes of wording. More disturbing, they found enormous differences when the cover story was changed to one concerning children who had been taken hostage by terrorists (while the underlying structure of the choice remained the same). Respondents then overwhelmingly chose the risk-equitable option (in which each child has a 1 in 100 chance of dying while in hostage). It became clear (after running 1,366 subjects and 11 story variations) that, whatever it was that the subjects were reacting to, it was not the conflict between risk aversion and risk equity. In the face of these results, SDMs seeking guidance on people's preferences should be cautious in generalizing from the research literature.

Finally, expressed preferences are cheap, but are people willing to put their money where their mouths are? It does not necessarily follow that, if I say that I prefer coffee over tea, I will order coffee when it costs twice as much as tea. The risk-perception literature has delineated several attributes (e.g., catastrophe avoidance) that people rely on when describing activities or technologies as risky and when judging the need for risk reduction. SDMs could make use

of this information by incorporating these attributes into their analyses. But doing so means that attaining a more desired state on one of these attributes will, usually, entail giving up a bit of some other attribute, perhaps money or lives. It is critically important to examine people's willingness to make such trade-offs. Slovic, Lichtenstein, and Fischhoff (1984) explored the importance of catastrophe avoidance relative to lives lost. Their subjects were asked to choose one of two possible safety measures for a factory. One would reduce the incidence of single-fatality accidents, thereby saving about 30 lives a year. The other would reduce the probability of a multiple-fatality accident (or, in another form, the number of lives lost thereby), thus reducing the expected lives lost by 27 per year. A substantial proportion of the subjects chose the former, suggesting that they did not view catastrophe avoidance as more important than saving three lives. This suggestion needs more direct confirmation. SDMs should be hesitant to incorporate the public's strongly felt desires until the public is asked to put their money where their mouths are.

CONCLUSIONS

Our purpose in writing this chapter is to encourage others to join us in thinking about the ethical problems faced by SDMs in making decisions when their views differ from those of the affected public. We feel compassion for the SDM, struggling to do the right thing in difficult circumstances. We want our SDM to be intellectually well armed with an understanding of decision analysis sufficient to ensure that complex social problems can be viewed from a broad, consistent perspective. We also want our SDM to have a backbone, able to go against (while never ignoring) public desires, and a heart, caring for and respecting (but not always acquiescing to) public views.

This is a tall order, we admit. These are tough problems, unavoidably so. We recognize that the world of laboratory experiments, which is at the root of both behavioral decision theory and risk perception, is a far safer and simpler place than the world faced by SDMs. Yet we also believe that the link between experimental and policy settings is strong.

In an important sense, what the SDM knows has got to be enough; the problems do not allow the luxury of indecision. But neither do they permit complacency; as students of decision making, we are able to tell SDMs far less than they would like to know. One implication is that more research is needed, research about the consequences of SDMs' actions as well as about what people want their SDMs to do. Another implication is that decision theorists should be encouraged to seek out real-world decision-making settings, as rich sources of inspiration and as challenging settings to explore the validity of laboratory-based findings. Finally, we need to learn more about the ethical bases for decision making so that SDMs can feel more secure in knowing when to lead, when to follow, and when to punt.

REFERENCES

Curley, S. P., F. Yates, and R. A. Abrams. 1986. Psychological sources of ambiguity avoidance. *Organizational Behavior and Human Decision Processes* 38:230–56.

Edwards, W., and D. von Winterfeldt. 1987. Public values in risk debates. *Risk Analysis* 7:141–58.

Einhorn, H., and R. Hogarth. 1981. Behavioral decision theory: Processes of judgment and choice. *Annual Review of Psychology* 32:53–88.

———. 1987. Decision making under ambiguity. In *Rational choice: The contrast between economics and psychology,* ed. R. Hogarth and M. Reder, 41–66. Chicago: University of Chicago Press.

Ellsberg, D. 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* 75:643–69.

Fischhoff, B. 1982. Debiasing. In *Judgment under uncertainty: Heuristics and biases,* ed. D. Kahneman, P. Slovic, and A. Tversky, 422–44. New York: Cambridge University Press.

Fischhoff, B., S. Lichtenstein, P. Slovic, S. Derby, and R. Keeney. 1981. *Acceptable risk.* New York: Cambridge University Press.

Fischhoff, B., P. Slovic, and S. Lichtenstein. 1980. Knowing what you want: Measuring labile values. In *Cognitive processes in choice and decision behavior,* ed. T. Wallsten, 117–41. Hillsdale, N.J.: Erlbaum.

Hammerton, M., M. W. Jones-Lee, and V. Abbott. 1982. Equity and public risk: Some empirical results. *Operations Research* 30:203–7.

Henrion, M., and B. Fischhoff. 1986. Assessing uncertainty in physical constants. *American Journal of Physics* 54:791–98.

Hershey, J. C., and P. J. H. Schoemaker. 1985. Probability versus certainty equivalence methods in utility measurement: Are they equivalent? *Management Science* 31:1213–31.

Hogarth, R., ed. 1982. *New directions for methodology of social and behavioral science: Question framing and response consistency.* San Francisco: Jossey-Bass.

Kahneman, D., P. Slovic, and A. Tversky, eds. 1982. *Judgment under uncertainty: Heuristics and biases.* New York: Cambridge University Press.

Keeney, R. L. 1980. Equity and public risk. *Operations Research* 28:527–34.

———. 1988. Facts to guide thinking about life-threatening risks. In *Proceedings of the IEEE 1988 International Conference on Systems, Man, and Cybernetics.* Oxford: Pergamon.

Keller, L. R., and R. K. Sarin. 1988. Equity in social risk: Some empirical observations. *Risk Analysis* 8:135–46.

Kunreuther, H. 1987. The failure of EIL coverage: Gridlock in environmental insurance. *Environment* 29, no. 1:18–20, 31–34.

Kunreuther, H., and P. R. Kleindorfer. 1986. A sealed-bid auction mechanism for siting noxious facilities. *American Economic Review* 76:295–99.

Lichtenstein, S., and P. Slovic. 1971. Reversals of preference between bids and choices in gambling decisions. *Journal of Experimental Psychology* 89:46–55.

Loomes, G., and R. Sugden. 1982. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal* 92:805–24.

MacGregor, D., and P. Slovic. 1986. Perceived acceptability of risk analysis as a decision-making approach. *Risk Analysis* 6:245–56.

Machina, M. J. 1987. Choice under uncertainty: Problems solved and unsolved. *Economic Perspectives* 1:121–54.

McNeil, B. J., S. G. Pauker, H. C. Sox, Jr., and A. Tversky. 1982. On the elicitation

of preferences for alternative therapies. *New England Journal of Medicine* 306: 1259–62.

Morgan, M. G., H. K. Florig, I. Nair, and G. L. Hester. 1987. Power-frequency fields: The regulatory dilemma. *Issues in Science and Technology* 3:81–91.

Okrent, D. 1987. The safety goals of the U.S. Nuclear Regulatory Commission. *Science* 236:296–300.

Samuelson, P. A. 1950. Probability and the attempts to measure utility. *Economic Review* (Tokyo) 1:167–73.

Savage, L. J. 1972. *The foundations of statistics.* New York: Dover.

Slovic, P. 1987. Perception of risk. *Science* 236:280–85.

Slovic, P., B. Fischhoff, and S. Lichtenstein. 1977. Behavioral decision theory. *Annual Review of Psychology* 28:1–39.

————. 1979. Rating the risks. *Environment* 21, no. 3:14–20, 36–39.

Slovic, P., and S. Lichtenstein. 1983. Preference reversals: A broader perspective. *American Economic Review* 73:596–605.

Slovic, P., S. Lichtenstein, and B. Fischhoff. 1984. Modeling the societal impact of fatal accidents. *Management Science* 30:464–74.

————. 1988. Decision making. In *Handbook of experimental psychology,* vol. 2, *Learning and cognition,* ed. R. C. Atkinson, R. J. Herrnstein, G. Lindzey, and R. D. Luce, 673–738. New York: Wiley.

Slovic, P., and A. Tversky. 1974. Who accepts Savage's axiom? *Behavioral Science* 19:368–73.

Tversky, A. 1969. Intransitivity of preferences, *Psychological Review* 76:105–10.

Tversky A., and D. Kahneman. 1981. The framing of decisions and the psychology of choice. *Science* 211:453–58.

————. 1986. Rational choice and the framing of decisions. *Journal of Business* 59, no. 4, pt. 2:S251–S278.

Wagenaar, W. A., G. Keren, and S. Lichtenstein. 1988. Islanders and hostages: Deep and surface structures of decision problems. *Acta Psychologica* 67:175–89.

Wagnaar, W. A., and S. D. Sagaria. 1975. Misperception of exponential growth. *Perception and Psychophysics* 18:422–26.