Risk's Place in Decision Rules

Paul Weirich
Philosophy Department
University of Missouri-Columbia
Columbia, MO 65211, USA

What a delight to celebrate the achievements of Peter Gärdenfors! His monumental book, Knowledge in Flux, gives new direction to epistemology. His judicious collection of readings, compiled with Nils-Eric Sahlin, Decision, Probability, and Utility, is a standard text for many courses on decision theory. These are just two items from a long list of impressive contributions to philosophy. To highlight some of Gärdenfors's ideas, I will comment on an article from his collection on decision theory. He wrote it with his co-editor. It is entitled "Unreliable Probabilities, Risk Taking, and Decision Making." This article is a treasure of lucid and penetrating insights about rational decision making. I will review and reorganize some of those insights to show off their advantages when viewed from a new angle.

1. Unreliable Probabilities as a Source of Risk

Gärdenfors and Sahlin notice that people prefer betting on an event about which they have extensive information to betting on an event about which they know little, even if they assign the events the same probability. Ellsberg's Paradox provides a well known illustration of this phenomenon. Let me present a simplified version of the paradox. It involves your choice of gambles concerning two urns. Each urn contains 100 balls. Each ball is either red or green. You know that the first urn contains 50 red balls and 50 green balls. But you have no information at all about the percentages of red and green balls in the second urn. You will draw a ball at random from each urn. Suppose that you may choose between receiving $100 if you draw red from the first urn, and receiving $100 if you draw red from the second urn. Which do you prefer? Next, suppose that you may choose between receiving $100 if you draw green from the first urn, and receiving $100 if you draw green from the second urn. Which do you prefer? Most people prefer to take their chances with the first urn in both cases since they know more about it.

Straightforward applications of the principle to maximize expected utility make this pattern of preferences paradoxical. No probability assignment for states and utility assignment for money explain these preferences. The first preference requires that P(R1)U($100) > P(R2)U($100), and hence that P(R1) > P(R2). The second preference requires that P(G1)U($100) > P(G2)U($100), and hence that P(G1) > P(G2). But since P(R1) + P(G1) = P(R2) + P(G2), those inequalities cannot hold simultaneously. Despite seeming rational, it appears that the preferences imply a violation of the principle to maximize expected utility.

Gärdenfors and Sahlin devise an ingenious explanation of such anomalous preferences. The first step is to recognize that some probability assignments are more reliable than others. Although a person confronted with Ellsberg's Paradox may assign .5 as the probability of getting red from the first urn and also from the second urn, the first probability rests on knowledge of the percentage of red balls in the first urn, whereas the second probability rests on total ignorance of the percentage of red balls in the second urn. The second probability is just the result of applying some version of the principle of insufficient reason. Probability assignments that rest on relevant information are more reliable than those made in the absence of relevant information.

The second step is to recognize that the reliability of a probability assignment for events determining the outcome of a gamble affects the riskiness of the gamble. The less reliable the probability assignment, the more risky the gamble, other things being equal. Since people are generally averse to risk, they prefer betting on events for which they have reliable probability assignments to betting on events for which they have unreliable probability assignments, when other things are equal. In Ellsberg's Paradox, you are asked to choose between gambles that are equivalent except for the reliability of the probability assignments for events determining the gambles' outcomes. Being risk averse, people tend to choose the gambles for which they have reliable probability assignments. Such choices reduce risk.

This resolution of Ellsberg's Paradox rings true. Gärdenfors and Sahlin build on it. They use the paradox's resolution as the inspiration for a new, general procedure for making decisions. Their procedure starts with a representation of a decision maker's epistemic state of mind. Instead of representing it with a single probability assignment, they represent it with a set of probability assignments. The less extensive the agent's information, the more diverse his set of probability assignments.

Other theorists have used sets of probability assignments to represent an agent's epistemic state of mind. Gärdenfors and Sahlin, however, contribute novel instructions for obtaining that set. First, they consider all the probability assignments that the agent regards epistemically possible. An assignment is epistemically possible if the agent's information does not rule it out. Gärdenfors and Sahlin have in mind objective probability assignments, such as textbook assignments of probabilities to the outcomes of games of chance. These objective probability functions need not assign 1 to the true state of nature and 0 to all other states. They assign probabilities with respect to information practically attainable (326, n. 18). Thus, events concerning the past need not have objective probabilities that are 1 or 0 since verification of their occurrence or non-occurrence may be practically speaking impossible. In some decision problems the agent may know the objective probability assignment of the relevant states. In other decision problems he may know only that it is in a certain set of probability assignments. All the probability assignments he does not rule out are in his initial set of probability assignments.

Gärdenfors and Sahlin then say that the agent should rank probability assignments according to reliability. An assignment's reliability depends on the weight of evidence supporting it. Extensive and high quality support, coming, for example, from thorough and careful statistical testing, makes a probability assignment reliable. Flimsy and slipshod support, coming, for example, from guesswork, makes a probability assignment unreliable. Although Gärdenfors and Sahlin do not precisely define reliability, they encourage thinking of a probability function's reliability as its likelihood of being the correct objective probability function (319, n. 11).

The next step is the introduction of a threshold of reliability. It depends on the agent's attitude toward risk in the decision problem he faces. If willing to take some epistemic risk, the agent eliminates probability assignments with low reliability. If willing to take more epistemic risk, he eliminates probability assignments with reliability below a higher threshold. The more willing the agent is to take an epistemic risk, the higher the threshold of reliability. Only probability assignments with reliability above the threshold enter the set directing his decision. The agent's epistemic risk is the risk of excluding the objectively correct probability assignment. If it is excluded, then the set directing his decision does not contain the correct assignment.

Gärdenfors and Sahlin's decision rule says to identify each option's minimum expected utility with respect to the set of qualifying probability assignments. Then pick an option whose minimum expected utility is at least as great as any other option's minimum expected utility. They call this rule the maximin expected utility rule, MMEU. It was first introduced by Gärdenfors (1979: 169).

Notice that Gärdenfors and Sahlin's decision rule incorporates some aversion to pragmatic risk. It uses the minimum expected utility of an option with respect to the agent's set of qualifying probability assignments as a measure of the option's choiceworthiness. Using the minimum expected utility is a cautious policy. It reduces risk. The risk reduced, however, is the risk of an undesirable outcome rather than the epistemic risk of excluding the correct objective probability assignment. Their decision procedure mandates this reduction of pragmatic risk, whereas it leaves to the agent the degree of reduction of epistemic risk. According to it, the agent adopts for himself a degree of epistemic risk by adopting a threshold of reliability for probability functions. The agent is not free to adopt for himself a degree of pragmatic risk, however. The decision rule MMEU does this for him.

The traditional Bayesian rule to maximize expected utility uses the agent's subjective probability function for states rather than the correct objective probability function because typically the agent does not know the correct objective probability function. According to Gärdenfors and Sahlin, a subjective probability function arises from the reduction of a set of epistemically possible objective probability functions to a single probability function serving as an estimate of the correct objective probability function. One reduction that generates a single probability function is the adoption of a reliability threshold so high that just one epistemically possible objective probability function reaches it. Gärdenfors and Sahlin permit this reduction but think that in most cases it takes an extreme epistemic risk. Other reductions are also unappealing because they do not properly capture the relevance to rational decision making of the set of epistemically possible probability functions and their reliabilities. To Gärdenfors and Sahlin, the sensible policy is to forgo reduction to a single probability function.

Gärdenfors and Sahlin show that in the absence of relevant information, their decision rule MMEU is equivalent to the maximin rule; whereas given full information, their decision rule is equivalent to the rule to maximize expected utility (325-6). Hence their decision rule moves smoothly from decision under uncertainty (with total ignorance of the objective probability function) to decision under risk (with complete knowledge of the objective probability function). It fills in the gap between these two traditional categories for decision problems.

Since their rule reduces to the maximin rule in certain cases, it needs a judicious choice of states. To prevent inconsistency as partitions of states vary, the maximin rule must be restricted to partitions that yield basic outcomes, whose utilities are constant across finer partitions of states. Otherwise, if utilities of outcomes are computed in the usual way, the maximin rule may change its recommendation when the partition of states is made more or less fine. Gärdenfors and Sahlin's rule MMEU likewise needs a partition of states that yields basic outcomes. They say that their decision rule yields consistent results as the partition of states is made less fine (332, n. 21). But this is not so in all cases. Making the partition of states less fine reliably preserves the rule's recommendation only if states generate basic outcomes. I therefore presume that Gärdenfors and Sahlin make the implicit assumption that states generate basic outcomes. Besides making their decision rule consistent, this assumption also prevents unreliable probabilities from making utilities unreliable too, which would happen if outcomes were chancy rather than basic and had utilities dependent on probabilities.

How does the new decision rule MMEU stand with respect to the traditional rule to maximize expected utility? In some passages Gärdenfors and Sahlin reject the traditional rule because they say it involves a mistaken assumption that an agent's epistemic state of mind may be represented by a single probability function (313-16). It presumes that the agent knows the objective probability function for relevant states although this is not always the case. They propose to substitute MMEU for the mistaken traditional rule. However, in other passages they accept the traditional rule, recognizing its implicit restriction to cases in which a single probability function represents an agent's epistemic state of mind (322-26). They advance MMEU as the traditional rule's generalization for cases in which a person's epistemic state of mind is not represented by a single probability assignment, but by a set of assignments together with their reliabilities.

I prefer the more charitable interpretation of the traditional rule, and so view their rule as its generalization to a wider range of decision problems. The traditional rule, as they say, makes the assumption that a decision maker has selected a unique probability assignment. In realistic cases, the decision maker may not have settled on a unique probability assignment. Even with unlimited reflection, he may not be able, in a justified way, to narrow down the epistemically possible assignments to a single one. The traditional rule does not apply to such cases. It does not go wrong in them since it is restricted to cases where a unique probability assignment has been adopted. It does not make the mistaken assumption that in every decision problem the agent has settled on a unique probability assignment. Its assumption of a unique probability assignment is an idealization. The idealization is a convenient theoretical tool, useful for constructing partial explanations of rational decision making. It puts aside a factor complicating decision making, namely, imprecision about probabilities, in order to deal with other factors, such as the utilities of options' possible outcomes. A realistic decision theory must eventually remove the traditional rule's idealization and explain how to cope with imprecise probabilities. Gärdenfors and Sahlin take an interesting step in this direction. They admit it is not a final step. Their new decision rule MMEU also rests on certain idealizations. Removal of those idealization, for the sake of greater realism, requires further modifications (333).

2. Good's Rule

Gärdenfors and Sahlin compare their decision rule to similarly motivated rules of Wald, Levi, and Kyburg, and argue for their rule's advantages. Their scholarship on the literature pertaining to their decision rule is excellent. However, because Gärdenfors and Sahlin think that unreliable probabilities provide a reason to move beyond expected utility maximization, they do not consider how expected utility maximization might respond to their worries about risk. I will explore a way of responding to those worries without abandoning expected utility maximization.

My approach follows Good (1952: 114). It uses only one standard for rational decisions, expected utility maximization, but in cases where a set of probability functions represents an agent's epistemic state of mind, sanctions any option that maximizes expected utility under some probability assignment in that set. When no single probability assignment, but rather a set of probability assignments, represents an agent's epistemic state, Good's decision rule says to maximize expected utility with respect to any one of those assignments. It advances no other constraint on rational decisions. It proposes a necessary and sufficient condition for a rational choice. Let us explore this decision rule further and then compare it with Gärdenfors and Sahlin's decision procedure.

Good's rule begins with a set of probability assignments representing an agent's epistemic state of mind. Then it says simply that any option maximizing expected utility according to some probability assignment in the set is a rational choice. This rule seems to ignore Gärdenfors and Sahlin's points about risk. According to Gärdenfors and Sahlin, risk aversion rules out many of the options Good's rule sanctions since it first eliminates some probability assignments as unreliable and then eliminates some options optimal under some probability assignments because those options are not cautious about pragmatic risk. However, under some interpretations, Good's rule captures Gärdenfors and Sahlin's insights about risk.

First, Good's rule can heed Gärdenfors and Sahlin's main point about the initial set of probability assignments. The initial set's job is to represent the agent's epistemic state of mind.

According to Gärdenfors and Sahlin, this is done by listing the objective probability assignments that the agent regards as epistemically possible and then attaching a reliability rating to each probability assignment. However, let us take the initial set of probability assignments as a set of subjective probability assignments that together represent the agent's epistemic state of mind, in the spirit of Jeffrey (1992: 29, 69). The set of probability assignments is obtained by listing those compatible with constraints the agent places on his assignment of subjective probabilities. He may, for example, impose the constraint that an event's probability fall between .4 and .6. Then the set includes those assignments, meeting other constraints, that give the event a probability between .4 and .6. The set of subjective probability assignments may, but need not, express judgments about objective probabilities. Its production may bypass objective probabilities altogether. It does not need them to generate a representation of an agent's epistemic state of mind. A subjective probability function need not be an estimate of anything objective, but may be simply a representation of dispositions to bet and the like.

Our modification of the initial set of probability assignments generalizes applications of Good's rule, since the rule requires a set of subjective probability assignments, and since subjective probability assignments need not reduce to epistemically possible objective probability assignments, or estimates of objective probability assignments.

Using subjective probability assignments, the representation of an agent's epistemic state of mind may be simplified in some cases. Even if the agent considers many objective probability functions epistemically possible, and assigns them different reliability ratings, it may be possible to reduce the set of objective probability functions to a single subjective probability function that adequately represents the agent's epistemic state of mind for decision purposes. The reduction produces a single subjective probability function that provides objective probability estimates to use in the decision rule to maximize expected utility. Thus, under the subjective interpretation of the set of probability assignments directing decisions, there may be fewer cases in which that set contains multiple subjective probability assignments, and so fewer cases in which the agent must rely on a decision rule that accommodates imprecise probabilities.

Besides beginning with a set of probability assignments, Good's rule also begins with a set of utility assignments. These are subjective utility assignments that together represent the agent's evaluation of outcomes. The utility assignments may cover the options themselves as well as their outcomes. To qualify as a subjective utility assignment, an assignment of utilities to both options and their outcomes must meet certain coherence requirements, just as an assignment of degrees of beliefs to states must meet certain coherence requirements to qualify as a subjective probability assignment. A probability and a utility assignment are paired and together coherently quantify an agent's beliefs and desires.

Second, although Good's rule abandons Gärdenfors and Sahlin's reliability parameter, it can accommodate aversion to risk. In their examples Gärdenfors and Sahlin assume that an option's outcome is an amount of money (324). Instead of adopting this interpretation, let an option's outcome include everything that would ensue if the option were realized. Besides money won or lost, the outcome of a gamble then includes the risk undertaken in adopting the gamble. In general, let an option's outcome be the possible world that would be realized if the option were realized. An option's outcome, understood this way, comprehends the risk the option involves. Given this broad interpretation, an option's possible outcomes, together with a set of probability and utility assignments for them, comprehend everything relevant to the option's evaluation. When appraising the option's utility, expected utility calculations omit no relevant consideration.

Suppose an objective probability assignment is unreliable in Gärdenfors and Sahlin's sense. In our reconstruction of their framework, the unreliability of the assignment generates many subjective probability assignments corresponding to it. Given our broad interpretation of outcomes, if an option's outcome depends on states of the world for which the agent has diverse subjective probability assignments, then adopting the option involves undertaking a risk. If the agent is averse to risk, the utility of every possible outcome of the option is diminished because of that risk. The objective probability assignment's unreliability thus diminishes the option's expected utility. The option therefore falls in status with respect to Good's rule. The risk diminishes an option's expected utility under each subjective probability assignment in the agent's set. The risk makes it harder for the option to maximize expected utility under any of those assignments. In this way Good's rule accommodates Gärdenfors and Sahlin's insight about unreliable probabilities, risk, and rational decision. With outcomes taken broadly, Good's rule responds to Gärdenfors and Sahlin's point about the rationality of avoiding epistemic risk.

Take Ellsberg's Paradox. Since betting on the 50-50 urn is less risky than betting on the mystery urn, every outcome of betting on the mystery urn is more deeply discounted for risk than the corresponding outcome of betting on the 50-50 urn. Therefore the expected utility of a bet on the 50-50 urn exceeds the expected utility of the analogous bet on the mystery urn. Consequently, Good's rule provides an explanation of the usual preferences. That explanation agrees with Gärdenfors and Sahlin that the crucial factors are epistemic risk and aversion to it.

Good's rule presumes compliance with all pertinent procedural and substantive standards for the formation of probability and utility assignments. Probabilities and utilities are as precise as warranted. The agent has reflected as much as warranted. The justification for Good's rule is, then, that every consideration bearing on an option's relative choiceworthiness bears on its place in a preference ranking of options, and every such consideration affects utility comparisons of options, which in turn depend on the probability and utility assignments for options' outcomes. Nothing affecting an option's relative choiceworthiness does not affect the probability or utility of its possible outcomes. A reason for choosing an option, other than a tie breaking reason, is a reason for preferring it to other options, and so a reason for its having higher expected utility than they have, and so a reason pertaining to the probabilities or utilities of options' outcomes. Only options ruled out by such reasons are unchoiceworthy. Once options ruled out by probability and utility assignments have been dismissed, no considerations are left to guide a choice between remaining options. Each pair of probability and utility assignments in the set representing an agent's epistemic state of mind expresses every reason that bears on preferences between options and so every reason that bears on choice. Any reason to pass over an option is a reason to prefer other options, and so a reason to adopt a constraint that elevates their expected utilities above that option's expected utility. Such a constraint then bars all pairs of probability and utility assignments supporting the option from the set representing the agent's state of mind.

One question Gärdenfors and Sahlin leave unanswered is whether an option their rule recommends is to be preferred or at least indifferent to all other options. If it is, as I suspect, then to make options' utilities agree with their preference ranking, and to make options' utilities equal their expected utilities, the utilities of outcomes must be adjusted so that compliance with Gärdenfors and Sahlin's decision procedure is compatible with compliance with Good's rule. Good's rule then accommodates their concerns.

3. Comparison of Rules

How should we handle the risk arising from unreliable probability assignments? Gärdenfors and Sahlin offer a clever proposal. Using comprehensive outcomes, Good's rule offers an alternative proposal. The main difference between the proposals is that Gärdenfors and Sahlin's says more precisely how to handle risk. Good's rule imposes no particular method of assessing risk and no particular attitude toward risk. It is, in fact, compatible with an attraction to risk. Gärdenfors and Sahlin's decision procedure mandates an aversion to risk, especially in its second stage, which applies the decision rule MMEU. Does rationality require their treatment of risk?

The principle to maximize expected utility is part of a theory of instrumental rationality. Rules for handling risk are part of a noninstrumental theory of rationality. Such rules dictate attitudes toward risk. They declare contrary attitudes and goals irrational. Since the assessment and evaluation of risk is an unsettled area, I prefer working within the liberal framework provided by Good's rule. Perhaps Gärdenfors and Sahlin are right about risk, but the literature is filled with rival treatments of risk. Going beyond Good's rule and adopting a particular rule for handling risk is a risky business. I side with Gärdenfors and Sahlin on the importance of considering risk when making decisions, and agree completely with their insight that unreliable probabilities generate risk, but suspend judgment about their method of handling risk. Perhaps on some occasions attraction to risk is warranted. Perhaps on some occasions it is rational to be less cautious than the maximin expected utility rule, MMEU.

Also, the decision procedure that Gärdenfors and Sahlin advance is not completely general. They plan adjustments for an agent's aspiration levels (333-4). For example, aspiration levels explain shifts in reference points separating gains from losses. They think that this phenomenon, and related phenomena Kahneman and Tversky (1979) study, call for a more careful look at utilities.

For generality, their decision procedure needs expansion in other directions as well. It accommodates only risks arising from unreliable probabilities, but risks have other sources too. Their procedure needs enrichment to take account of risks that arise, for example, from the range of utilities of an option's possible outcomes. These risks arise in Allais's Paradox, and attention to them is critical for resolving the paradox, as Weirich (1986) argues.

Gärdenfors and Sahlin's approach to risk aversion takes an option's outcomes narrowly, and responds to risk by eliminating insufficiently reliable probability assignments and applying the maximin expected utility rule, MMEU. If outcomes are taken narrowly, factors besides risk, such as regret, are also left out. They must be handled by additional rules supplementing expected utility considerations. The decision procedure that emerges threatens to spin a web of supplementary rules, one rule for each factor excluded from an option's outcome. It will be difficult to coordinate these supplementary rules, since some excluded factors, such as risk and regret, are complementary. Supplementary rules for such complementary factors may conflict. Rather than advance a patchwork of supplementary rules, it is more attractive from a theoretical perspective to extend the expected utility framework, as Good's rule does.

Good's rule also helps with the general problem of incomparability. A pair of options is incomparable if the agent is unable to form a preference ranking of them. Options in a decision problem may be incomparable. The agent, even with reflection, may be unable to rank them because of an inability to assign probabilities to the states of the world determining an option's outcome, or because of an inability to assign utilities to the option's possible outcomes. The incomparability of options frustrates the decision procedure to pick an option at the top of a preference ranking of options. Because of incomparability, there may be no preference ranking of options. In that case, the decision procedure needs enlargement. Gärdenfors and Sahlin address incomparability arising from imprecise probabilities. Good's rule addresses incomparability arising from both imprecise probabilities and imprecise utilities. In effect, it asks only that the option picked be such that no other option is preferred.

Theorists are sometimes unsatisfied with Good's approach to incomparability. Suppose, for example, that a professor must decide whether to accept an administrative position. She is content with her current position, but sees opportunities as an administrator to improve her university. The two jobs may be incomparable for her. Perhaps she is not indifferent between them but yet does not prefer one to the other. In that case Good's rule says each is a rational choice. This lack of direction may seem unsatisfactory, however. The reason, I believe, is that the professor suspects that further reflection may reveal a method of comparing the jobs. If she vividly imagines being an administrator—the inflexible routine, the additional paperwork, the inevitable change in values—perhaps she will find a way to compare the two jobs. Dissatisfaction with Good's tolerance arises from the feeling that if only she could be more imaginative, resourceful, and reflective, she could make the comparison. If she is reconciled to the jobs' incomparability, then choosing either indeed seems to be a satisfactory resolution of the decision problem. Good's rule handles the qualms incomparability generates by insisting that before it is applied probability and utility assignments, and so the options' preference ranking, be made as precise as warranted.

4. A Defense of Comprehensive Outcomes

I have presented Good's rule as a way of accommodating risk without commitment to a method of appraising it. If the rule is to be a solid foundation for methods of evaluating risks, a number of issues must be addressed.

First, what is risk? We need to know more about it before we can say whether building it into an option's outcome makes sense. I take risk, as I think Gärdenfors and Sahlin do, as a feature of an option whose outcome is uncertain. Usually, for an option to be called risky, some possible outcome must be bad, but we may stretch ordinary usage by calling any option with an uncertain outcome risky if some possible outcome is either good or bad. Also, the uncertainty creating a risk relevant to rational decision making is the decision maker's uncertainty. Since this uncertainty is a subjective feature of the decision maker, dependent on his information, the type of risk of interest is subjective. It depends on the decision maker's subjective probability and utility assignments.

In economics it is common to define risk aversion in terms of the shape of an agent's utility curve for a commodity. Risk itself is not defined, just an attitude toward it. When the commodity is money, risk aversion amounts to the diminishing marginal utility of money. To distinguish the two, we define risk aversion as an aversion to subjective risk as just introduced. Making such distinctions is one motivation for making risk independent of attitudes toward it and including it in a gamble's outcome.

Second, is risk as introduced double counted by making it part of an option's outcome? Some theorists believe that the rule to maximize expected utility handles risk. Building risk into an option's outcome makes the rule double count risk, they think. But risk is surely part of the outcome of a risky option. And the argument for the rule to maximize expected utility requires taking outcomes comprehensively. Any feature of an option's outcome not included in the outcomes used for expected utility calculations makes those calculations an unreliable indicator of the option's utility and its choiceworthiness. Ellsberg's Paradox teaches this lesson.

Third, if the risk involved in an option depends on probability and utility assignments for possible outcomes, and possible outcomes comprehend the risk involved in an option, circularity threatens. When assigning utilities to an option's possible outcomes, one evaluates the risk the option involves, a feature of every possible outcome. But when assessing the risk the option involves, one considers its possible outcomes' utilities. In some cases these interlocking considerations create a disequilibrium that blocks an option's utility assignment. However, in normal cases, these considerations reach an equilibrium that permits an option's utility assignment. In the abnormal cases, Good's rule applies using the set of probability and utility assignments forming the unstable cycle of assignments.

Fourth, in applying Good's rule one considers whether an option maximizes expected utility under some assignments of probability and utility in the agent's set of assignments. A pair of probability and utility assignments in the set precisely specify probabilities and utilities, and so eliminate the risk arising from imprecise probabilities and utilities. Hence it seems that applications of Good's rule ignore risks arising from imprecision. This objection misunderstands Good's rule. The rule does not say to imagine that a pair of probability and utility assignments from the agent's set of assignment are the agent's actual assignments. It says only to compute expected utilities using those assignments to see which options they sanction. The utility assignment evaluates outcomes containing risks arising from imprecision. Use of the pair of probability and utility assignments does not remove from outcomes risks imprecision generates. The agent's set of probability and utility assignments represents a state of mind. No pair of probability and utility assignments from the set need represent a rational state of mind. They may, indeed, represent an incoherent state of mind, one in which precise probabilities are attached to outcomes whose utilities are reduced because the outcomes' probabilities are imprecise. Still, sets of such probability and utility assignments may represent a coherent state of mind, one arising from lack of information about outcomes.

Fifth, taking an option's outcome to be comprehensive so that it includes risk makes the rule to maximize expected utility hard to apply. Instead of evaluating a gamble by attaching utilities to possible monetary outcomes, one must evaluate complex outcomes involving money, risk, and every other possible component of the gamble's realization. Comprehensive outcomes are so complex, it is difficult to evaluate them in a way that facilitates an option's evaluation. My reply to this objection is that for the sake of decision theory, firm foundations are more important than rules of thumb easy to apply. If making the expected utility principle sound calls for taking outcomes comprehensively, then decision theory should do that. If the expected utility principle crumbles, little remains on which to build. From a theoretical perspective, preserving that principle is worth the practical difficulties its preservation causes. It is the cornerstone of decision theory.

An analogous situation arises in logic. The law of excluded middle seems doubtful when applied to vague sentences. One might abandon the law, but a more fruitful approach is to look for an interpretation that makes it reliable. Let it apply to sentences that have definite truth values. This approach does not solve the practical problem of settling whether a sentence has a truth value. But theoretical advances such as sentential logic are possible without solving those practical problems. Similarly, the expected utility principle is sound if interpreted as applying to an option's comprehensive outcomes. This interpretation aggravates the practical problem of assigning utilities to an option's outcomes. Nonetheless, it is attractive because it preserves the foundations of classical decision theory.

Sixth, some may object to applying Good's rule to a series of decisions. For example, if the probability of rain tomorrow ranges from .4 to .6, the rule authorizes paying $.6 for a gamble that pays $1 if it rains, and paying $.6 for a gamble that pays $1 if it does not rain. But if one makes both bets, one is sure to lose money. This objection rests on a misconception of applications of Good's rule to a series of decisions. One should apply it to multiple decisions in a way that is context sensitive. When making new bets, one should keep track of bets already made. Bets already made influence the possible net gains or losses from betting if one makes additional bets. One should evaluate new bets with respect to consequences for net wealth.

The foregoing points defend Good's rule against objections. I conclude that it is a sensible standard for decisions in cases where a set of probability and utility assignments represent an agent's state of mind. When outcomes are taken comprehensively, the rule gives risk its due without departing from an expected utility framework. I think it provides an attractive package for the important points about reliability, risk, and rationality that Gärdenfors and Sahlin forcefully bring to our attention.

References

Gärdenfors, P. 1979. ”Forecasts, Decisions, and Uncertain Probabilities.” Erkenntnis 14: 159-81.

Gärdenfors, P. 1988. Knowledge in Flux: Modeling the Dynamics of Epistemic States. Cambridge, MA: The MIT Press.

Gärdenfors, P. and N. Sahlin. 1982. ”Unreliable Probabilities, Risk Taking, and Decision Making.” Synthese 53: 361-86. Reprinted in P. Gärdenfors and N. Sahlin, eds., Decision, Probability, and Utility: Selected Readings, pp. 313-34. Cambridge: Cambridge University Press.

Good, I. J. 1952. ”Rational Decisions.” Journal of the Royal Statistical Society, Ser. B, 14: 107-14.

Jeffrey, R. 1992. Probability and the Art of Judgment. New York: Cambridge University Press.

Kahneman, D. and A. Tversky. 1979. ”Prospect Theory: An Analysis of Decision under Risk.” Econometrica 47: 263-91.

Weirich, P. 1986. ”Expected Utility and Risk.” British Journal for the Philosophy of Science 37: 419-42.