

An Incompatibility Between Preferential Ordering and the Decision-Theoretic Notion of Utility

Marek J. Druzdzel

Decision Systems Laboratory
School of Information Sciences
and Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
marek@sis.pitt.edu

Abstract

A class of preferential orderings in non-monotonic logics assumes that various extensions of a model (possible worlds) can be ordered based on both their likelihood and desirability. I suggest that there is a basic incompatibility between this qualitative notion of preference and the decision-theoretic notion of utility. I demonstrate that while reasoning and decision making in the former can focus on a single state, it is meaningless in expected utility theory to say that a state or a set of states is important for a decision. This, I believe, is thought-provoking as it poses the question whether a qualitative formalism should be compatible with its quantitative counterpart or whether it can afford to be at odds with it. I discuss the difference between normative and cognitive utility and the implications of this difference for work on user interfaces to systems based on probabilistic and decision-theoretic methods.

Introduction

One way of looking at models of uncertain domains is that these models describe a set of possible states of the world, only one of which is true. A state can be succinctly defined as an element of the Cartesian product of sets of outcomes of all individual model's variables.¹ This view is explicated by the logical Artificial Intelligence (AI) approaches to reasoning under uncertainty — at any given point various extensions of the current body of facts are possible, one of which, although unidentified, is assumed to be true. Decision theory, the most widely accepted quantitative formalism for decision making under uncertainty, assigns each state of the model two numbers: a measure of uncertainty in this state (its probability) and a measure of its desirability (its utility). Optimal decisions are those that maximize the expected utility of the states that may result from the corresponding actions.

To accommodate the notion of preferences in addition to uncertainty, researchers in non-monotonic log-

ics proposed preferential ordering of possible worlds. Shoham (1987) explicated this notion and argued that differences among non-monotonic logics can be reduced to differences in preference for various possible non-monotonic inferences. This ordering supposedly captures both uncertainty and preferences.

In this paper, I suggest that there is a basic incompatibility between the qualitative notion of utility implied by preferential ordering and the decision-theoretic notion of utility. I demonstrate that while reasoning and decision making in the former can focus on a single state, it is meaningless in expected utility theory to say that a state or a set of states is important for a decision. This, I believe, is thought-provoking as it poses the question whether a qualitative formalism should be compatible with its accepted quantitative counterpart or whether the two can be at odds with each other. I discuss the apparent difference between normative and cognitive notions of utility and the implications of this difference for work on user interfaces to systems based on probabilistic and decision-theoretic methods.

Utility in Decision Theory

The foundations for the decision-theoretic notion of utility have been laid out by Von Neumann and Morgenstern (von Neumann & Morgenstern 1944). Starting from the qualitative notion of preferences, they proposed a procedure for deriving a quantitative measure of utility that is essentially based on observing human betting behavior. They demonstrated that given certain weak assumptions, from among the large class of functions reflecting the ordinal properties of utility

$$a \succ b \iff U(a) > U(b) , \quad (1)$$

where " $a \succ b$ " means " a is preferred to b ," there exists a subclass of utility functions that has the desirable property of combining in exactly the same way as mathematical expectation, i.e.,

$$U(\alpha a + \beta b) = \alpha U(a) + \beta U(b) , \quad (2)$$

where a and b are any gambles and α and β are non-negative numbers such that $\alpha + \beta = 1$.

¹There is a richness of terms used to describe states of a model: extension, instantiation, possible world, scenario, etc. Throughout this paper, I will attempt to use the term *state of a model* or briefly *state* whenever possible.

Von Neumann and Morgenstern have also proposed a procedure based on preference between gambles for eliciting this utility. Consider three events A , B , and C such that the order of the individual preferences is $A \succ B \succ C$. The decision maker is offered the choice




Figure 1: Utility measurement

between B for sure and a gamble in which A can be won with probability α and C with probability $1 - \alpha$ (see Figure 1). By assumption, there exists such value α for which the decision maker is indifferent between the two options and, therefore, assigns both options the same utility value.

$$U(B) = U(\alpha A + (1 - \alpha)C)$$

Given the values of $U(A)$ and $U(C)$, the utility of B is then equal to

$$U(B) = \alpha(U(A) - U(C)) + U(C).$$

The values of $U(A)$ and $U(C)$ are assigned arbitrarily. It has become customary to choose A and C to be the most desirable and the least desirable outcome and assign them the values of 0 and 100 respectively. It can be verified that $\forall B U(B) \in [0, 100]$. In other words, all utility values are confined to the interval between the utility of the least and the most desirable outcomes.

The reader is referred to an excellent exposure of the relation between the cardinal utility proposed by von Neumann and Morgenstern and the economic notions of utility offered by Savage (Savage 1972, Section 5.6). Savage demonstrates that both subjective probability and von Neumann and Morgenstern type of utility are a natural consequence of a set of axioms that can be viewed as maxims of rational behavior under uncertainty. He also gives these theories a normative interpretation.

Utility constrained by (1) and (2) is shown by von Neumann and Morgenstern (1944) to be determined up to an increasing linear transformation. If U is a utility function, then any increasing linear transformation of U

$$\phi(U) = \omega_0 U + \omega_1,$$

where ω_0, ω_1 are constants, with $\omega_0 > 0$, is also a utility function encoding the same preferences. In other

words, there may exist many mappings of utility into real numbers, but all these mappings will be linear transformations of each other. Presence of the positive term ω_0 in the transformation ϕ amounts to changing the scale. The constant factor ω_1 moves the zero point of the scale. Since ϕ is a meaningless transformation, it follows that although utility is a cardinal measure, it has no unit and no zero point. The fact that utility in this formalization is determined up to an increasing linear transformation follows from the desire not to overconstrain the formalism above its empirical foundations: this form of utility is sufficient for decision making under uncertainty.

A practical consequence of the decision-theoretic formalization is that it is possible to compare the relative desirability of several events against each other, but it would be meaningless to say, for example, that one event is twice as desirable as another event. A useful parallel to this measure, originally pointed out by von Neumann and Morgenstern, is the measure of temperature (let us forget the absolute Kelvin scale for a moment). There are many scales possible (take the most popular: Celsius and Fahrenheit), each having a separate unit and a different zero point. Although it is possible to make relative comparisons of temperature, it is meaningless to say that a day A was twice as warm as a day B . I would like to refer interested readers to a review paper by Stevens (1959), who makes several very interesting points about measurement of utility and measurement in general.²

Individual States in Decision-Theoretic Reasoning

Preferential ordering in non-monotonic logics assumes that a state s_1 is preferred to a state s_2 if s_1 precedes s_2 in the ordering. This notion of preference takes into account both the uncertainties and the desirabilities of s_1 and s_2 . I am taking here the simplest possible but also the most intuitive interpretation of this ordering: Of two equally desirable states s_1 and s_2 , s_1 precedes s_2 in the ordering if s_1 is more likely than s_2 . Of two equally likely states s_1 and s_2 , s_1 will precede s_2 if it is more desirable than s_2 . Admittedly, there are other interpretations possible. One such interpretation, suggested by a reviewer, is that s_1 will precede s_2 if its expected utility is higher than that of s_2 . I do not consider this interpretation, as I see no evidence for nor any indication of how such ordering is or may be accomplished without a computationally complex process that considers uncertainty and utility of individual states. If the ordering is based purely on utility, as another reviewer suggested, it is of practical value only in reasoning under certainty, i.e., where there is only one extension of the facts possible. (This, in turn, does not require non-monotonic logics.) With

²I would like to thank an anonymous reviewer for a pointer to this excellent paper.

any amount of uncertainty, i.e., with more than one state possible, the ordering would have to be sensitive to the likelihoods of these states, as argued convincingly by Shoham (1987).

It turns out that because of the nature of decision-theoretic utility and, in particular, because of its invariance to increasing linear transformation, an ordering on states is from the point of view of decision theory provably meaningless. Suppose that we have identified a set S_1 of states that are important in the sense of covering most of the expected utility of a decision option d . Let S_2 be another set of states that contributes less than S_1 to the expected utility of d , and therefore is judged to be inferior to S_1 in explaining d . In other words,

$$\Delta EU(S_1|d) - \Delta EU(S_2|d) > 0$$

An argument in favor of d cannot be made on the ground of the selected set of states S_1 instead of the set S_2 for the following reason. If we have

$$\sum_{s_i \in S_1} Pr(s_i|d)U(s_i|d) - \sum_{s_j \in S_2} Pr(s_j|d_2)U(s_j|d_2) > 0$$

then an increasing linear transformation of utility

$$\phi(U) = \omega_0 U + \omega_1, \quad \omega_0 > 0.$$

transforms this equation into the following

$$\begin{aligned} & \sum_{s_i \in S_1} Pr(s_i|d)\phi(U(s_i|d)) \\ & - \sum_{s_j \in S_2} Pr(s_j|d)\phi(U(s_j|d)) = \\ & = \sum_{s_i \in S_1} Pr(s_i|d)(\omega_0 U(s_i|d) + \omega_1) \\ & - \sum_{s_j \in S_2} Pr(s_j|d)(\omega_0 U(s_j|d) + \omega_1) = \quad (3) \\ & = \omega_0 \left(\sum_{s_i \in S_1} Pr(s_i|d)U(s_i|d) \right. \\ & \quad \left. - \sum_{s_j \in S_2} Pr(s_j|d)U(s_j|d) \right) \\ & \quad + \omega_1 \left(\sum_{s_i \in S_1} Pr(s_i|d) - \sum_{s_j \in S_2} Pr(s_j|d) \right) \end{aligned}$$

Now, suppose that

$$\sum_{s_i \in S_1} Pr(s_i|d) > \sum_{s_j \in S_2} Pr(s_j|d).$$

We choose w_0 and w_1 so that they satisfy the following constraint

$$\omega_1 < \omega_0 \frac{\sum_{s_j \in S_2} Pr(s_j|d)U(s_j|d) - \sum_{s_i \in S_1} Pr(s_i|d)U(s_i|d)}{\sum_{s_i \in S_1} Pr(s_i|d) - \sum_{s_j \in S_2} Pr(s_j|d)}.$$

If, on the other hand,

$$\sum_{s_i \in S_1} Pr(s_i|d) < \sum_{s_j \in S_2} Pr(s_j|d),$$

we choose w_0 and w_1 as follows

$$\omega_1 > \omega_0 \frac{\sum_{s_j \in S_2} Pr(s_j|d)U(s_j|d) - \sum_{s_i \in S_1} Pr(s_i|d)U(s_i|d)}{\sum_{s_i \in S_1} Pr(s_i|d) - \sum_{s_j \in S_2} Pr(s_j|d)}.$$

In both cases, we have that

$$\begin{aligned} & \omega_0 \left(\sum_{s_i \in S_1} Pr(s_i|d)U(s_i|d) - \sum_{s_j \in S_2} Pr(s_j|d)U(s_j|d) \right) \\ & + \omega_1 \left(\sum_{s_i \in S_1} Pr(s_i|d) - \sum_{s_j \in S_2} Pr(s_j|d) \right) < 0, \end{aligned}$$

or, in other words

$$\sum_{s_i \in S_1} Pr(s_i|d)\phi(U(s_i|d)) - \sum_{s_j \in S_2} Pr(s_j|d_2)\phi(U(s_j|d_2)) < 0$$

which shows that, according to the very same criterion, S_2 is superior to S_1 in justifying d . Of course, it may also happen that

$$\sum_{s_i \in S_1} Pr(s_i|d) = \sum_{s_j \in S_2} Pr(s_j|d),$$

in which case, S_1 will remain superior to S_2 under any transformation of utility. This is, however, of little practical significance — we get the absurd result that even though S_1 is superior to S_2 , it can be shown inferior to any subset of S_2 .

Figure 2 illustrates this problem graphically. In one of the possible utility transformations (the upper horizontal axis), the rectangles A, B, C, D, E , and F are ordered according to their size. The area of a rectangle is proportional to the product of its probability and its utility. In a different utility transformation (the lower horizontal axis, corresponding to a lower value of the zero point of the utility function), this order is no longer valid. Note that the rectangle F , for example, which was one of the smallest rectangles under the first transformation, becomes the largest rectangle under another transformation.

We see that the ordering among individual states is meaningless because it is not invariant to a meaningless transformation of utility. Following the same procedure, it can be shown that for any single state s_1 there exists a state s_2 that, under some utility transformation, contributes more to the expected utility of the decision than s_1 . Similarly, if two sets of states S_1 and S_2 are chosen to justify superiority of a decision d_1 over a decision d_2 , it can be demonstrated that there exist utility transformations under which these two sets of states actually justify something exactly opposite, namely that d_2 is superior to d_1 .




Figure 2: Sensitivity of selection of states to utility transformations.

A useful analogy for the decision-theoretic inference that might help the reader to understand the problem is the following. Suppose that we want to determine whether one building is warmer than another building by averaging the temperature in each building over all the rooms (taking the sizes of the rooms into consideration). It is easy to verify that this operation is invariant to the temperature scale used (i.e., it will produce the same answer whether one uses the Celsius scale or the Fahrenheit scale). The moment, however, that we decide to approximate the task by taking only those rooms into consideration that contribute the most to the sum of averages, which rooms will be considered depends on the scale used. For temperatures around freezing point of water (around 0°C , 32°F), even very large rooms will contribute little to the average temperature when using the Celsius scale. These rooms will, because of their size, contribute a lot when Fahrenheit scale is used. In other words, which rooms are important for the difference in temperature between the two buildings is not invariant to a change in temperature scale.

Cognitive and Normative Utility

Normative utility expresses only a relative valuation of outcomes and is insensitive to their absolute values. What the lowest and the highest utility values mean is relative with respect to all possible outcomes of a decision problem. The utilities of the least desirable and the most desirable of all possible outcomes are as-

signed in an arbitrary way in the elicitation process. If the utility function over all outcomes in a decision problem is confined to an interval $[0, 100]$, there is no way to tell whether the lower bound of this interval, zero, stands for an extremely undesirable, catastrophic event (like death) or just a relatively less desired outcome (such as winning only \$1,000 instead of the top \$100,000 prize in a TV trivia game). Similarly, the upper bound of this interval, 100, can mean a spectacularly desired outcome (a free vacation for two on the Bahamas), but it may also mean simply the lesser evil in a difficult dilemma (such as becoming disabled instead of dying).

While the relative valuation such as encoded in the utility function is sufficient for normative decision making, it seems to be insufficient for explaining it. It seems that a system generating explanations of decision-theoretic inference based on the state method needs to have a way of identifying those states that contribute significantly both in a positive and in a negative way to the expected utility of a decision and should also have a rough idea of the absolute desirability of outcomes.

The theoretical argument outlined in the previous section has demonstrated the impossibility of pointing out which of the many possible states are important for a decision. In other words, it has shown that speaking of the contribution or the importance of a state for the decision is, with respect to the expected utility theory, theoretically meaningless. This result seems to be counterintuitive. After all, it seems, that reference to states is not uncommon in explanations and justifications of decisions made by humans. Consider a decision whether or not to build a dam in a region affected by floods. It is not uncommon to hear justifications like: “Even though the possibility of a large flood is remote, the fact that it may literally wipe out all houses along with their inhabitants makes us very reluctant not to build a dam,” or “A large flood has not happened in the area for over 50 years — think of all the money that could be saved by not building a dam and could be invested in other programs in the region.” The first sentence refers to a very undesirable course of action that involves not building a dam, a large flood, and huge losses in terms of human life and property and is clearly a state with a certain load of valuation. The second sentence refers to a state involving not building a dam, no large flood in the area, and enjoying the money saved on the dam project on some other communal need, in short, a rather desirable course of events. According to the theoretical argument in the previous section, both states in isolation should be meaningless with respect to the decision. But it seems that these states are not meaningless for people. Both sentences are perfectly meaningful: we understand that the speaker refers to an unlikely and highly undesirable outcome in the first case and to a very likely and desirable outcome in the second case.

Although these sentences clearly do not refer to differences in expected utility, but rather to possible courses of events and desirability or undesirability of outcomes, we seem to understand them and even tend to agree about their importance for the decision. What makes these sentences meaningful in isolation, if, given the nature of utility, they should make no sense? There is something that seems to go beyond decision-theoretic utility in human reasoning: it seems that valuations of outcomes have more meaning for people than relative comparisons. Catastrophic and spectacularly desirable outcomes seem to be making their way to our attention, even though their contribution to the decision should be judged only in the context of all other possible outcomes.

A look at the Equation 4 shows that the reason for the trouble is the term ω_1 , which, as explained earlier, is related to the lack of a zero point in utility.³ Does this mean that utility, as we perceive it in the real world, has a zero point? The different constructs that have been introduced in our language in the context of decision making, such as *desirable* and *undesirable* or *gains* and *losses* suggest existence of an outcome with respect to which some outcomes are favorable and some are not.⁴ Use of these words is certainly context dependent, but they are used and they are meaningful for people. It suggests that we do have decision problem dependent zero points and we refer to them in our language. It is a robust experimental finding that most people act differently in the domain of gains and in the domains of losses: they exhibit a risk averse behavior in the first and risk-seeking behavior in the second (Tversky & Kahneman 1981). Decisions in the domain of losses seem to be more difficult for people to make. There are even two different terms for each of the two in the language — decision involving gains is usually called a choice, decision involving losses is called a dilemma.

Similarly, it seems that the magnitude of stakes plays an important role in human decision making. If we return to the temperature scale analogy: it is true that there was no physical reason for fixing a zero on the scale (before the existence of an absolute zero was postulated) and there was no particular reason for fixing the unit. Still, an isolated temperature data, say 0°C, tells a lot — it is the temperature at which water freezes and a temperature that we start to experience as unpleasant when we go outside. We know what 100°C is, as we know how it feels to put a finger in

a pot of boiling water. Similarly, we know that 25°C outside means a nice summer day. Although the scale is in a sense arbitrary, we do know well what most points on that scale mean. Returning to the utility measurement, units are arbitrary, but still most people know how it feels to lose \$1,000, how it feels to get ill, and can somehow imagine how it might feel to win \$1,000,000 in a lottery. Also, there seems to be always a reference point with respect to which we perceive some outcomes as positive and some negative. It is like preparing a bath and experiencing temperatures above 37°C as warm and below that as cold. This reference point is not constant: when setting the thermostat, one might treat temperatures close to 20°C in a special way.

My conjecture is that in the case of utility the outcome with respect to which we perceive another outcome as a loss or as a gain is the *status quo* (with respect to the elements affected by the decision process). If we are ill and make a decision about which treatment to undergo, the zero point may be the *status quo* and the anticipated developments given no medical treatment. An improvement in health due to a treatment has a positive value, while getting worse (or even dying) due to a treatment is normally associated with a negative value. If we consider an investment, we usually call yields below the average interest rate in the “bank around the corner” to be losses, while yields higher than that to be gains. We should keep in mind that such zero points, if they exist, are unnecessary for normative decision making and lack formal foundations in the expected utility theory. They have also counterintuitive implications: once a zero point has been fixed, we should be able to provide relative valuations of outcomes, such as “outcome *a* is twice as desirable as outcome *b*’.

Conclusion

This paper suggested a basic incompatibility between the qualitative notion of utility implied by ordering imposed on possible worlds in non-monotonic logics and the decision-theoretic notion of utility. I demonstrated that while states are compared and a single state in the former may be the focus of reasoning and decision making, it is meaningless in expected utility theory to say that a state is important for a decision. This seems to be also at odds with the cognitive notion of preference, where concentrating on individual states is a common practice. A practical consequence of the decision-theoretic formalization is that it is possible to compare the relative desirability of several events against each other, but it would be meaningless to say, for example, that one event is twice as desirable as another event.

Problems like the one identified above deserve attention, as ideally a qualitative scheme should be an abstraction of its quantitative counterpart. For example, building qualitative physics that ignores the quan-

³Note that an increasing linear transformation with $\omega_1 = 0$, i.e., only changing the scale, will never produce this effect. On the other hand, a transformation involving $\omega_0 = 1$ and $\omega_1 \neq 0$, i.e., changing the zero point will produce this effect.

⁴I am concerned here with one particular usage of these words. They may also be used as pure valuation without a decision making context “being ill is undesirable” or “it is good to be rich.”

titative laws of physics may lead to wrong inferences and resulting disappointment. Qualitative notions of decision-theoretic utility, used in economics long before the development of the expected utility theory, relax the requirements on the possible utility measures and extend the class of equivalent utility functions to any monotonic transformation of a given utility function. Non-monotonic logics are essentially qualitative schemes and one might expect them to be similarly unrestrictive towards utility. It seems that the problem identified in this paper will disappear if the only transformations on utility allowed are those monotonic transformations that do not change the zero point. But this implies something that may be undesirable and counterintuitive: If there were a quantitative formalism at their foundations, it would allow to make statements like “state s_1 is twice as desirable as state s_2 .”

Acknowledgments

While I am solely responsible for possible shortcomings of this paper, I am indebted to Herb Simon for inspiring discussions on the topic of utility and its role in decision making and anonymous reviewers for useful suggestions and interesting comments. Support for this work was provided by the National Science Foundation under Faculty Early Career Development (CAREER) Program, grant IRI-9624629, and by ARPA’s Computer Aided Education and Training Initiative under grant N66001-95-C-8367 to University of Pittsburgh.

References

- Savage, L. J. 1972. *The Foundations of Statistics (Second Revised Edition)*. New York, NY: Dover Publications.
- Shoham, Y. 1987. Nonmonotonic logics: Meaning and utility. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence, IJCAI-87*, 388–393.
- Stevens, S. 1959. Measurement, psychophysics, and utility. In Churchman, C. W., and Ratoosh, P., eds., *Measurement: Definitions and Theories*. New York: John Wiley & Sons, Inc. 18–63.
- Tversky, A., and Kahneman, D. 1981. The framing of decisions and the psychology of choice. *Science* 211(4481):453–458.
- von Neumann, J., and Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton University Press.