11/15/04

Plans and Decisions[1]

John L. Pollock
Department of Philosophy
University of Arizona
Tucson, Arizona 85721
pollock@arizona.edu
http://www.u.arizona.edu/~pollock

Abstract

Counterexamples are constructed for classical decision theory, turning on the fact that actions must often be chosen in groups rather than individually, i.e., the objects of rational choice are plans. It is argued that there is no way to define optimality for plans that makes the finding of optimal plans the desideratum of rational decision-making. An alternative called "locally global planning" is proposed as a replacement for classical decision theory. Decision-making becomes a non-terminating process without a precise target rather than a terminating search for an optimal solution.

1. The Optimality Prescription

How should I go about making rational decisions? That is the fundamental question of the theory of rational choice. Human beings, and any real cognitive agents, are subject to cognitive resource constraints. They have limited reasoning power, in the form of limited computational capacity and limited computational speed. This makes it impossible, for example, for them to survey all of the logical consequences of their beliefs, or to compare infinitely many alternatives. This is a fundamental computational fact about real agents in the real world, and I would suppose that it could not have been otherwise. An account of how a real agent should make decisions must take account of these limitations.

Theories of rational decision-making are sometimes taken to be theories about how ideal agents, immune to such cognitive limitations, should make decisions (Cherniak 1986; Skyrms 1980, 1984; Lewis 1981). One can, of course, choose to talk that way, but it is hard to see what that has to do with what we, as fallible human beings, should do. For instance, if a theory of ideal agents says that they should attend to all of the logical consequences of their beliefs, but we as human beings cannot do that, then the recommendations applicable to ideal agents are simply not applicable to us. *We* should do something else. As I use the term "the theory of rational decision-making", it is about what we, and other resource bounded cognitive agents, should do. I want to know how, given our cognitive limitations, we should decide what actions to perform.

[1] This work was supported by NSF grant no. IRI-IIS-0080888.

In other words, I want a theory of *real rationality* as opposed to a theory of *ideal rationality*. This distinction is widely recognized, but it often seems to be supposed that as philosophers our interest should be in ideal rationality. The rationality a human can achieve is mere "bounded rationality" — a crude approximation to ideal rationality. But we come to the study of rational decision-making with an initial interest in how we, and agents like us, should make decisions. This is the notion of rationality that first interests us, and it is the target of this paper.

Decisions must often be made in the face of uncertainty regarding both the agent's initial situation and the consequences of his actions. Most discussions of rational choice proceed against the background of classical decision theory, which is generally assumed uncritically. The basic ideas of classical decision theory can be stated simply. We assume that our task is to choose an action from a set **A** of *alternative actions*. The actions are to be evaluated in terms of their outcomes. We assume that the *possible outcomes* of performing these actions are partitioned into a set **O** of pairwise exclusive and jointly exhaustive outcomes. We further assume that we know the probability **PROB**(O/A) of each outcome conditional on the performance of each action. Finally, we assume a *utility-measure* **U**(O) assigning a numerical utility value to each possible outcome. The *expected-utility* of an action is defined to be a weighted average of the values of the outcomes, discounting each by the probability of that being true if the action is performed. The crux of classical decision theory is that actions are to be compared in terms of their expected-utilities, and rationality dictates choosing an action that is *optimal*, i.e., such that no alternative has a higher expected-utility. I will call this *the optimality prescription*.

It is my conviction that classical decision theory and the optimality prescription are seriously flawed in a number of essentially orthogonal respects.[2] One class of problems has given rise to several varieties of *causal decision theory*, which require **PROB** to be a kind of "causal probability".[3] I have discussed these problems elsewhere (my 2002, 2005), and constructed my own favored version of causal decision theory. But this is independent of the problem that will be discussed in this paper. A second problem that besets classical decision theory is the failure of "action omnipotence". In deciding on a course of action, an agent will often not know with certainty which actions he will be able to perform. This necessitates important changes to the way expected-utilities are computed. But again, this problem can be ignored for present purposes.[4] This paper raises a more fundamental problem for classical decision theory — one related to Savage's (1954) "small worlds problem". The problem is that actions cannot be chosen in isolation. In general, an action can only be evaluated as part of a more comprehensive package of actions — a plan. This leads ultimately to a plan-based decision theory, but the

---

2 See my forthcoming book, *Thinking about Acting: Logical Foundations for Rational Decision-making*.

3 See Gibbard and Harper 1978; Sobel 1978; Skyrms 1980, 1982, 1984; Lewis 1981. See my (2002) for my own preferred version of causal decision theory, and my (2005) for an updated version.

4 They are discussed at length in my (2003, 2005).

resulting theory differs from classical decision theory in major respects.

2. Decision-Theoretic Alternatives

2.1 Deciding-Whether and Deciding-Which

The fundamental problem of practical deliberation is what to do. Such deliberation comes in two forms. Sometimes we are deciding whether to perform a particular action. At other times we are deciding which of several actions to perform. A fundamental presupposition of decision theory is that deciding-whether reduces to deciding-which. Deciding-which is a matter of deciding between alternatives, and the alternatives are to be evaluated, in accordance with the optimality prescription, in terms of their expected utilities.

There is a trivial sense in which deciding-whether is a case of deciding which. In deciding whether to do A, we are deciding between doing A and not doing A. Let $\overline{A}$ be the action of *not* doing A. Then deciding whether to do A is a matter of deciding which of A and $\overline{A}$ to do. However, this decision cannot be made by comparing the expected utilities of A and $\overline{A}$. Jeffrey (1965) considers the example of deciding whether to bring red wine or white wine to a dinner party. This is an example of deciding-which. Suppose instead that we just consider whether we should bring red wine. Let this be the action A. Can we make this decision by comparing the expected utility of A with the expected utility of $\overline{A}$? No. In computing the expected utility of $\overline{A}$, we cannot assume that if we do not take a bottle of red wine then we will take a bottle of white wine. We might take nothing at all. There is no way to predict that we will take a bottle of white wine until we have solved the decision problem at least to the extent of determining that these are the only two viable alternatives. Thus the expected utility of $\overline{A}$ will be essentially 0. Hence applying the optimality prescription in this way would lead us to take a bottle of red wine. However, by the same reasoning, it is better to take a bottle of white wine than not to take a bottle of white wine, so this reasoning will also lead us to take a bottle of white wine. And we do not want to take both. Thus we cannot choose between A and $\overline{A}$ in this way.

Decision theory assumes that in deciding whether to perform A, we should consider what we might do instead, where this is more than just not doing $\overline{A}$. Thus in the wine case, we evaluate the action of taking red wine by comparing it with the action of taking white wine. These are considered "alternatives", and the optimality prescription proposes that we should choose A only if there are no better alternatives. But it doesn't tell us how to determine what actions are alternatives. Without that, the prescription cannot be applied to real cases and does not constitute a theory of rational choice.

It will be my contention that this entire approach to rational decision-making is fundamentally misguided. Deciding-whether questions cannot, in general, be reduced in this way to deciding-which questions. This will turn upon the claim that there is, in general, no way to make sense of the notion of a decision-theoretic alternative to an action that makes the optimality prescription a correct theory of rational decision-making. My argument for this claim will be rather involved. I will argue first that we cannot evaluate actions in isolation — they must be evaluated in combination with other actions, i.e., as parts of plans. Then I will argue that there is no way to select a fixed set of alternatives from the potentially infinite array of plans available, and even if there were there would be no reason to expect there to be optimal plans. Finally, I will argue that instead of looking for optimal plans, we should be looking for good plans, and I will attempt to make that notion precise.

2.2 Strong and Weak Competition

It might be supposed that the search for a set of alternatives has a trivial solution — any set of actions can constitute the alternatives — it is up to the decision maker. The alternative actions are just those the decision maker is considering performing. There are two obvious problems for this proposal. The first is that the decision maker may not be considering everything he should be considering. If I employ the optimality prescription to decide to take red wine without even considering the possibility of taking white wine, it does not follow that my decision is the right one. A second problem is that at any given time a decision maker may be considering many different actions and may choose to perform more than one of them. I may be considering going to lunch at noon, and reading a novel this evening. There is no reason why I should not do both. These actions are not in competition with each other.

The first problem is just the problem of finding the set of alternatives, but the second problem is more fundamental Apparently the actions in a set of alternatives must be, in some sense, "competing actions" that I must choose between. But what is it for actions to compete? It is frequently supposed that competing actions are those that cannot be performed together. Let us call these *strongly competing* actions. This is often (but not always[5]) built into the formal definition of a decision problem. It is often required that **A** is a set of actions that are pairwise logically incompatible (it is logically impossible to perform more than one of them) and exhaustive (it is logically necessary that you will perform at least one of them). In other words, **A** is a "partition of the action space".

However, common-sense decision problems do not generally involve choices between strongly competing actions. Jeffrey (1965) considers the example of deciding whether to take red wine or white wine to a dinner party. One could, of course, take both. These actions do not strongly compete. In fact, taking our ordinary descriptions of our decisions at face value, choices between strongly competing actions seem to be the exception rather than the rule. I face such problems as deciding whether to paint the shed or clean the house this afternoon, whether to fly to LA next week or the following week, whether to cook chicken for dinner or cook lamb chops,

---

[5] For instance, see Jeffrey (1983). It is also worth noting that in the literature on decision-theoretic planning, with the exception of MDP planning, the alternatives are never strong competitors.

and so on. In all of these cases I could do both. Still, in each case the two choices are in some sense in competition. I do not regard doing both as a serious option. We want alternatives to be actions we should, rationally, choose between. That is, we should choose one but not more than one. This can be the result of much weaker relations than strong competition. For instance, in Jeffrey's wine example, we must choose between red wine and white wine because taking both would be foolish. The cost would be twice as great but the payoff only marginally greater, so the expected utility of taking both is less than the expected utility of taking one. Thus rather than take both, we should take one, but that necessitates deciding which to take. We might capture this with a notion of weak competition — two actions *compete weakly* iff either they compete strongly or the expected utility of doing both is less than the expected utility of at least one of them.

An appeal to weak competition generates a theory with a somewhat different structure than the optimality prescription. The problem is that the optimality prescription assumes we have a set of alternative actions, and prescribes choosing an optimal member of the set. However, weak competition doesn't generate a set of alternatives. This is because weak competition is not transitive. $Action_1$ may compete weakly with $action_2$, and $action_2$ with $action_3$, without $action_1$ competing weakly with $action_3$. Thus if we simply pick an action and let the set of alternatives be the set of all actions competing weakly with the given action, it does not follow that other members of the set of alternatives will be in competition. It may be desirable to perform several of those alternatives together. For instance, suppose I am planning a wedding. Folk wisdom dictates that I should select something borrowed and something blue, but suppose it is undesirable to select two borrowed things or two blue things. If x and y are borrowed, and y and z are blue, then selecting x competes weakly with selecting y, and selecting y competes weakly with selecting z, but selecting x does not compete weakly with selecting z.

This problem does not depend upon taking weak competition as our competition relation. For instance, strong competition is not transitive either. However competition is to be defined, it seems that what the optimality prescription really ought to say is:

(OP) It is rational to decide to perform an action iff it has no competitor with a higher expected utility.

This is equivalent to talking about an optimal member of a set of alternatives only if the competition relation is transitive. There is no obvious reason to expect that to be the case, so I will henceforth assume that the optimality prescription takes the form of (OP).

2.3 Boolean Combinations

There is a way of reformulating these decision problems so as to obtain a set of pairwise strongly competing alternatives. Where $A_1,...,A_n$ are the actions we are choosing between, a *Boolean combination* of them is a specification of which A_i's to perform and which to not perform.

For example, where $\overline{A}$ is the action of *not* performing A, a Boolean combination might have the form A_1 & $\overline{A}_2$ & A_3 & ... . Different Boolean combinations are logically incompatible with each other, and the disjunction of them is logically necessary. The actions $A_1,...,A_n$ compete weakly just in case no Boolean combination with multiple positive constituents has as high an expected utility as some Boolean combination with just one positive constituent. Thus we can retain the original form of the optimality prescription if we apply it to the Boolean combinations of what are intuitively "alternative actions".

Although the appeal to Boolean combinations allows us to use strong competition as our competition relation, it is not clear that it really buys us anything. The problem is that the negative elements of a Boolean combination do not appear to make any contribution to the probabilities involved in computing expected utilities. When we discussed the comparison of A with $\overline{A}$, we noted that supposing we don't perform A leaves it pretty much wide open what actions we do perform and what else will happen in the world. That seems to be equally true for Boolean combinations in general. Supposing, for example, A_1 & $\overline{A}_2$ & A_3, will not usually give us any more useful information than just supposing A_1 & A_3. If this is right then applying (OP) to the individual actions is equivalent to selecting the optimal Boolean conjunction, but the fact that we get strong competition by looking at the Boolean combinations is really irrelevant to the decision-making. We get the same result without strong competition by just considering actions that compete weakly.

2.4 Universal Plans

Regardless of whether we understand competition in terms of weak competition, or strong competition, or something else, we still need an account of what competing alternatives we must compare an action to. For example, on any account of competition it will presumably be the case that A and $\overline{A}$ are competitors, but as we have seen, in deciding whether to perform A it is not sufficient to just compare its expected utility with that of $\overline{A}$. Conversely, even if B is not, intuitively, a competitor of A, $\overline{A}$&B competes strongly with A and, as we have seen, the expected utility of $\overline{A}$&B will normally be the same as the expected utility of B. In deciding whether to take red wine, we do not want the expected utility of that to have to surpass the expected utility of every unrelated action, like vacationing in Brazil next summer, but if $\overline{A}$&B were regarded as a competitor, it would have that effect. Thus there has to be something more to the set of alternatives than that it contains actions that compete with A. How should we select the set of alternatives?

One possibility is that all actions and combinations of actions should be regarded as potential competitors. Notice that the possible alternatives include more than just actions to be performed at the same time. For instance, I may want to choose between flying to LA next week or flying the following week. This suggests that $\mathbf{A}$ should consist of the Boolean combinations of all possible actions. On this approach, all actions, for all times, must be considered together in the

Boolean combinations. A Boolean combination becomes a complete prescription of what to do for the rest of one's life, and a decision problem is reduced to the problem of finding the optimal such course of action. We might call these "universal plans".

Savage (1954) toys with the idea that rational decisions should be between universal plans, but he rejects it for the obvious reason. It is computationally absurd to suppose I must plan the entire rest of my life in order to decide what to have for lunch. The real world is too complex for us to be able to construct or compare universal plans. No decision maker with realistic computational limitations could possibly govern its life by finding a plan prescribing the optimal action for every instant of the rest of its life. Most likely such plans would involve infinitely many individual actions (including all the not-doings), and finding an optimal plan would involve comparing infinitely many alternatives. Even if the agent only has to consider a finite number n of "atomic" actions, the number of Boolean combinations will be 2^n, and that number quickly becomes immense. Even for a number of actions that is implausibly small for describing an entire life (e.g., 300), 2^n will be a much larger number of Boolean combinations than a real agent can possibly survey and compare ($2^{300} \approx 10^{90}$).

As I have emphasized, my topic is rational decision-making by real agents, not ideal agents. We seek a theory of how a real decision maker should, rationally, go about deciding what actions to perform at any given time. A theory that requires a decision maker to do something that is impossible cannot be a correct theory of rationality. Real decision makers cannot construct and evaluate universal plans, so the theory of rationality cannot require them to.

2.5 The Search for Alternatives — Resurrecting (OP)

We have to decide what the reasonable alternatives should be before we can apply the optimality prescription. We need something in between taking $\mathbf{A}$ to consist of just A and $\overline{A}$ and taking it to consist of all possible universal plans. One proposal is to simply adopt (OP) as it stands. The suggestion would be that the competitors of a given action are all the actions that compete weakly with it. This does not seem to get us into the same sort of computational difficulties as the appeal to Boolean combinations did. If there are n actions that compete weakly with A, then those are all we have to consider in deciding whether to perform A. This may still be a rather large number, but not exponentially huge. However, the next section demonstrates that this simple proposal does not work.

3. Groups of Actions

The principle (OP) at least appears to evaluate actions by comparing them with other individual actions — those with which they compete weakly. The next step of my attack on the optimality prescription is an argument that we cannot, in general, make decisions in this way — by focusing on individual actions. I will argue that the proper objects of decision-theoretic evaluation are plans rather than individual actions. At this stage, we can give at least three reasons why this must be the case.

The first reason turns upon the observation that neither weak competition nor any other reasonable competition relation can be expected to be transitive. To illustrate the problem, in the "borrowed and blue" example, suppose choosing *z* (a blue thing) has a higher expected utility than choosing *y* (a borrowed and blue thing), and choosing *y* has a higher expected utility than choosing *x* (a borrowed thing). Thus (OP) implies that one ought to choose *z* (a blue thing), but it also implies that it is not reasonable to choose either *x* or *y* because both have competitors with higher expected utilities. Then we are left without a borrowed thing. On the contrary, it seems clear that if we choose *z* as our blue thing, then we ought to choose *x* as our borrowed thing. To get this result, we must consider choosing *x* and *z* as a package, and compare that with choosing *y* alone. So we cannot evaluate *x* in isolation. We have to look at groups of actions rather than single actions.

We can be led to this same conclusion by reflecting on the fact that we typically have a number of different decisions to make at more or less the same time. I may be deciding whether to go to the bank before lunch or after lunch, and also deciding where to go for lunch. This mundane observation again creates a problem for the optimality prescription because (OP) evaluates actions one at a time and has us choose them individually on the basis of their being optimal. The problem is that decisions can interact. Carrying out one decision may alter the probabilities and utilities involved in another decision, thereby changing what action is optimal. It could be that, prior to deciding where to go to lunch, because I am very hungry the optimal decision would be to postpone going to the bank until after lunch. But if I decide to have lunch at a restaurant far from the bank and I have other things to do in that part of town that could occupy me for the rest of the afternoon, this may make it better to go to the bank before lunch. Alternatively, because I am very hungry and want to eat before going to the bank, it might be better to choose a different restaurant. The point is that actions can interfere with one another, with the result that if several actions are to be chosen, their being individually optimal does not guarantee that the group of them will be optimal. This strongly suggests that the object of decision-theoretic evaluation should be the entire group of actions rather than the individual actions.

This same conclusion can be defended in a third way. Often, the best way to achieve a goal is to perform several actions that achieve it "cooperatively". In this case we must choose actions in groups rather than individually. To illustrate, suppose my objective is to transport a ton of silver and a ton of gold from one location to another. I have a one-and-a-half-ton truck. I could fit both the gold and the silver into the truck at the same time and transport them on a single trip, but in doing so I would risk damaging the truck springs. The actions I am considering are to transport the gold on a single trip, to transport the silver on a single trip, and to transport both on a single trip. We can imagine the probabilities and utilities to be such that the action with the highest expected utility is that of transporting both on a single trip, even though that risks damaging the springs. However, if I have time to make two trips, that might be a better choice. That is, I should perform *two* actions, transporting the gold on one trip and the silver on another, rather than performing any of the single actions I am considering. This illustrates again that actions cannot always be considered in isolation. Sometimes decision-theoretic choices must be between groups of actions, and the performance of a single action becomes rational only because it is part of a group of actions whose choice is dictated by practical rationality.[6]

The last two examples illustrate two different phenomena. In the first, actions interfere with each other, changing their execution costs and hence their expected utilities from what they would be in isolation. In the second, actions collaborate to achieve goals cooperatively, thus changing the expected utilities by changing the probabilities of outcomes. These examples might be viewed as cases in which it is unclear that actions even have well-defined expected utilities in isolation. To compute the expected utility of an action we must take account of the context in which it occurs. If the expected utilities are not well-defined, then the optimality prescription cannot be applied to these decision problems. Alternatively, if we suppose that the expected utilities of the actions in isolation are well-defined, then what is important about these examples is that in each case we cannot choose the group of actions by choosing the individual actions in the group on the basis of their expected utilities. In the first example, the expected utility of the group cannot be computed by summing the expected utilities of the actions in the group. In the second example, the members of the group would not be chosen individually on their own strength. Rather, a pairwise comparison of actions would result in the action of transporting both the gold and silver on a single trip being chosen, and that is the intuitively wrong choice. In these examples, it is the group itself that should be the object of rational choice, and the individual actions are only derivatively rational, by being contained in the rationally chosen group of actions.

Groups of actions, viewed as unified objects of rational deliberation, are *plans*. The simplest plans are linear sequences of actions. In general, plans can be viewed as "recipes" for action. I will say more about the logical structure of plans below. For now, the important point is that the actions in a plan may be good actions to perform only because they are part of a good plan. It appears that the only way to get the optimality prescription to make the right prescription in the above examples is to apply it to plans rather than individual actions. For instance, the reason we should transport the gold alone on a single trip is that doing so is part of the plan of making two trips, and that plan is better than the plan of transporting both the gold and silver on a single trip. The plan of making two trips has a higher expected utility than the plan of transporting

---

[6] Faced with this example, decision-theorists sometimes complain that by entertaining the possibility of making two trips I am changing the decision problem. But they are missing the point that I am discussing decision-making in the real world. We do not come to the problem with a predetermined list of alternatives. Part of the problem of deciding what to do concerns choosing the right set of alternatives.

both the gold and silver on a single trip, and that is the basis upon which it is chosen.

Traditionally, choices were supposed to be between individual actions, but now we have seen that rational choices must often be made instead between plans, and the individual actions in the plans become only derivatively rational by being prescribed by a rationally chosen plan. How then do we choose between plans? The obvious proposal is to simply apply the optimality prescription to plans rather than actions. *Simple plan-based decision theory* would propose that we choose between competing plans in terms of their expected utilities. Savage (1954) seems to suggest that plans can be chosen in this way, and most work on decision-theoretic planning in AI is based upon this idea (for example, see Blythe and Veloso 1997; Boutelier et al 1999; Haddawy and Hanks 1990; Ngo, Haddawy and Nguyen 1998; Onder and Pollack 1997, 1999; Onder, Pollack and Horty 1998; and Williamson and Hanks 1994).

4. Choosing Between Plans

Just as for actions, we need not choose between plans unless they are in some sense in competition. If two plans are not in competition, we can simply adopt both. So to construct a plan-based theory of rational choice, we need an account of when plans compete in such a way that a rational choice should be made between them. Competing plans should be plans that we must choose between, rather than adopting both. A sufficient condition for this is that executing one of the plans makes it impossible to execute the other one, i.e., the plans *compete strongly*. However, it is clear that we often want to choose between plans that compete in much weaker ways. For example, plans can contain conditional steps telling us to do something only if something else is the case. Consider two route plans. One might say, "Take Speedway unless you encounter road construction. If you do encounter road construction, take Grant instead." The second plan might say, "Take Speedway unless you encounter road construction. If you do encounter road construction, take Broadway instead." Such plans can prescribe different courses of action in some circumstances (if you encounter road construction) but not in others (if you don't encounter road construction), so although they are intuitively competitors, they are not strong competitors. Just as for actions, we might try to capture this in terms of weak competition. Let us say that two plans *compete weakly* iff either they compete strongly or the plan that results from merging the two plans into a single plan has a lower expected utility than at least one of the original plans. It might be proposed, then, that two plans are competitors iff they compete weakly, and accordingly:

(PB) It is rational to adopt (decide to execute) a plan iff it has no weak competitor with a higher expected utility.

To evaluate (PB), let us first reflect briefly on the nature of the plans a realistic decision maker must construct and evaluate. We have seen that they cannot be universal plans. They have to be plans of more limited scope. What is the nature of these less-than-universal plans? This is most easily understood by reflecting on the fact that, over the course of her life, a decision maker is not faced with a single fixed planning problem. First, her beliefs will change as she acquires experience of her environment and as she has time for further reasoning. This will affect what solutions are available for her planning problems. Second, as our decision maker acquires more knowledge of her environment, her goals may change. The upshot is that the planning problems she faces evolve over time. We cannot expect her to redo all of her previous planning each time she acquires new knowledge or new goals, so planning must produce lots of *local plans*. These are small plans of limited scope aiming at disparate goals.

If (PB) is to work, it must work when applied to local plans. Those are the kinds of plans that real decision makers construct and decide to adopt. However, there are two simple reasons why (PB) cannot possibly be correct when applied to local plans. The simplest reason is that there are infinitely many of them. Plans are logical entities of potentially unbounded complexity. (PB) would have us survey and compare all possible local plans in order to determine whether they compete with a given plan and, if they do, to determine whether they have a higher expected utility. But this is an impossible task. No real agent can consider all possible competitors to a given plan, so he cannot make decisions in accordance with (PB).

The cardinality problem is devastating enough, but it is worth noting that there is a second problem (taken from my 1992). Even if we could somehow survey and compare an infinite array of plans, (PB) would not yield rationally correct decisions. (PB) is simply wrong as a theory of rational choice. This arises from the fact that for any plan there will almost always exist a competing plan with a higher expected utility. To illustrate, suppose again that I am choosing between roasting chicken and barbecuing lamb chops for dinner. Suppose the former has the higher expected utility. This implies that the plan of barbecuing lamb chops for dinner is not rationally adoptable, but it does not imply that the plan of roasting chicken for dinner is adoptable, because some other plan with a higher expected utility may compete with it. And we can generally construct such a competing plan by simply adding steps to the earlier competing plan. For this purpose, we select the new steps so that they constitute a subplan aimed at achieving some valuable unrelated goal. For instance, we can consider the plan of barbecuing lamb chops for dinner and then later going to a movie. This plan still competes with the plan of roasting chicken for dinner, but it has a higher expected utility. Thus the plan of roasting chicken for dinner is not rationally adoptable. However, the competing plan is not rationally adoptable either, because it is trumped by the plan of roasting chicken for dinner and then later going to the same movie.

It seems clear that given two competing plans P_1 and P_2, if the expected utility of P_1 is greater than that of P_2, the comparison can generally be reversed by finding another plan P_3 that pursues unrelated goals and then merging P_2 and P_3 to form P_2+P_3. If P_3 is well chosen, this will have the result that P_2+P_3 still competes with P_1 and the expected utility of P_2+P_3 is higher than the expected utility of P_1. If this is always possible, then there are no optimal plans and simple plan-based decision theory implies that it is not rational to adopt any plan.

In an attempt to avoid this problem, it might be objected that P_2+P_3 is not an appropriate object of decision-theoretic choice, because it merges two unrelated plans. However, recall the third example used to motivate the application of decision theory to plans rather than actions — the example of transporting a ton of gold and a ton of silver. The plan we wanted to adopt in preference to transporting either the gold, the silver, or both on a single trip, was the plan to transport the gold on one trip and the silver on another trip. This plan is constructed by merging two unrelated plans for achieving unrelated goals. If we are not allowed to construct such merged plans, (PB) will not produce the intuitively correct prescription in this example.

The inescapable conclusion is that the rational adoptability of a plan cannot require that it have a higher expected utility than all its competitors. The problem is that plans can have rich structures and can pursue multiple goals, and as such they are indefinitely extendable. We can almost always construct competing plans with higher expected utilities by adding subplans pursuing new goals. Thus there is no appropriate set of alternatives to use in defining optimality, and hence no way to define optimality so that it is reasonable to expect there to be optimal plans. Consequently, simple plan-based decision-theory fails.

The failure of simple plan-based decision theory is of fundamental importance, so let me recapitulate the three reasons it fails. First, it is unlikely that there will, in general, be such things as optimal plans. Second, because plans cannot generally be compared just by comparing their expected utilities, optimality may not even be desirable. Finally, even in those unusual cases in which there are optimal plans and optimality is desirable, finding them will be computationally intractable. So deciding whether to adopt a plan cannot turn upon it's being optimal.

5. When is a Plan a Good One?

The question we ultimately want to answer is, "How should a cognizer go about deciding whether a plan should be adopted?" There will be no point at which a decision maker has exhausted all possibilities in searching for plans. Despite this, decision makers must take action. They cannot wait for the end of a non-terminating search before deciding what to do, so their decisions about how to act must be directed by the best plans found to date — not by the best possible plans that *could* be found. The upshot is that plan adoption must be defeasible. Decision makers must work with the best knowledge currently available to them, and as new knowledge becomes available they may have to change some of their earlier decisions. If a better plan is found later, that should supplant the plan adopted initially.

For this account to work, we need a notion of "better plan" that can be used in deciding whether a plan should be adopted. I have argued that this cannot be cashed out as one plan merely having a higher expected utility than a second. To get a grip on this notion of one plan being better than another, let us think about plan adoption in rational decision makers. First consider the limiting case in which a decision maker has no background of adopted plans, and a new plan is constructed. Should the new plan be adopted? The basic insight of the optimality prescription is that what makes a course of action (a plan) good is that it will, with various probabilities, bring about various value-laden states, and the cost of doing this will be less than the value of what is achieved. This can be assessed by computing the expected utility of the plan. In deciding whether to adopt the plan, all the decision maker can do is compare the new plan with the other options currently available to him. If this is the only plan the decision maker has constructed, there is only one other option — do nothing. So in this limiting case, we can evaluate the plan by simply comparing it with doing nothing.

Things become more complicated when one has already adopted a number of other plans. This is for two reasons. First, the new plan cannot be evaluated in isolation from the previously adopted plans. Trying to execute the previous plans may affect both the probabilities and the utilities employed in computing the expected utility of the new plan. Clearly, the expected utility of the new plan must be computed "in the context of the decision maker's other plans". But what does that require? Roughly, the probabilities and utilities should be conditional on the situation the decision maker will be in as a result of having adopted and tried to execute parts of the other plans. However, there isn't just one possible situation the decision maker might be in, because the other plans will normally have their results only probabilistically.

The second reason it becomes more complicated to evaluate a new plan when the decision maker already has a background of adopted plans is that the new plan can affect the value of the old plans. If an old plan has a high probability of achieving a very valuable goal but the new plan makes the old plan unworkable, then the new plan should not be adopted. Note that this is not something that is revealed by just computing the expected utility of the new plan.

We have seen that normal planning processes produce local plans. How should the decision maker decide whether to adopt a new local plan? The decision must take account of both the effect of previously adopted plans on the new plan, and the effect of the new plan on previously adopted plans. We can capture these complexities in a precise and intuitively appealing way by defining the concept of the decision maker's *master plan*. This is the result of merging all of the plans the agent has adopted but not yet executed into a single plan.

Don't confuse the master plan with a universal plan. The master plan simply merges a number of local plans into a single plan. Each local plan talks about what to do under certain circumstances, so the resulting master plan talks about what to do under every circumstance mentioned by any of the individual local plans. But this is still a very small set of circumstances relative to the set of all possible world-states. If none of the local plans have anything to say about what to do in some new previously unconsidered situation, then the master plan doesn't either. But by definition, a universal plan must include a prescription for what to do in every situation. If we have *n* local plans each making *m* prescriptions of the form "If *C* is true then do *A*", the master plan will contain $m \cdot n$ prescriptions. But supposing the conditions *C* are all logically independent of each other, a universal plan for the state space generated by just this limited vocabulary will contain $2^{m \cdot n}$ prescriptions. For example, if the agent has thus far adopted 30 ten-step plans, the master plan will include 300 prescriptions, but a universal plan would have to consider at least 2^{300} (i.e., 10^{90}) prescriptions, and probably many orders of magnitude more.

Although master plans are totally different beasts from universal plans, they share an important property — master plans can be meaningfully compared in terms of their expected utilitys. We can think of the master plan as the agent's tool for making the world better. The expected utility of the master plan is the agent's expectation of how good the world will be if he adopts that as his master plan. Thus one master plan is better than another iff it has a higher expected utility. Equivalently, rationality dictates that if an agent is choosing between two master plans, he should choose the one with the higher expected utility.

It may at first occur to one that the objective should be to find an optimal master plan. But that cannot be right, for two familiar reasons. First, it is unlikely that there will be optimal master plans that are smaller than universal plans. If a master plan leaves some choices undetermined, it is likely that we can improve upon it by adding decisions regarding those choices. But as we have seen, it is not possible for real agents to construct universal plans, so that cannot be required for rational choice. Second, even if there were optimal master plans, realistically resource-bounded agents could not be expected to find them. So rationality cannot require finding optimal master plans.

These points are fairly obvious, and yet they completely change the face of decision-theoretic reasoning. Planning and plan adoption must be done defeasibly, and actions must be chosen by reference to the current state of the decision maker's reasoning at the time he has to act rather than by appealing to the idealized but unreachable state that would result from the decision maker completing all possible reasoning and planning. Decision makers begin by finding good plans. The good plans are "good enough" to act upon, but given more time to reason, good plans might be supplanted by better plans.[7] The decision maker's master plan evolves over time, getting better and better, and the rules for rationality are rules directing that evolution, not rules for finding a mythical endpoint. We might put this by saying that a rational decision maker should be an *evolutionary planner*, not an optimizing planner.

---

[7] This is reminiscent of Herbert Simon's (1955) concept of "satisficing", but it is not the same. Satisficing consists of setting a threshold and accepting plans whose expected-utilities come up to the threshold. The present proposal requires instead that any plan with a positive expected utility is defeasibly acceptable, but only defeasibly. If a better plan is discovered, it should supplant the original one. Satisficing would have us remain content with the original.

6. Locally Global Planning

We can regard planning as aiming at finding a master plan with a higher expected utility than our current one. If the only way an agent had of finding a master plan with a higher expected utility was to plan all over again from scratch and produce a new master plan essentially unrelated to its present master plan, the task would be formidable. Performing the requisite planning would at the very least be slow and complex, making it difficult for the agent to respond to emergency situations. And if the agent's master plan is sufficiently complex, the agent's inherent computational limitations may make the task impossible. It does not take a very large problem to bog down a planning procedure. The reader unfamiliar with the AI literature on planning may not appreciate the severity of this problem. A few years ago, the very best AI planners could solve toy problems described in terms of 53 independent variables, where the solution was a plan of 105 steps (Weld 1999). In what was considered state of the art performance, BLACKBOX (Kautz and Selman 1998) was able to solve such a problem in 6 minutes on a fast computer. Typically, the master plan will be significantly larger than these toy problems. Furthermore, if every planning problem requires the construction of a new master plan, then every little planning problem becomes immensely difficult. To plan how to make a sandwich for lunch, I would have to replan my entire life.

Obviously, humans don't do this. Normal planning processes produce local plans, not entire master plans. The only way resource-bounded agents can efficiently construct and improve upon master plans reflecting the complexity of the real world is by constructing or modifying them incrementally. When trying to improve his master plan, rather than throwing it out and starting over from scratch, what an agent must do is try to improve it piecemeal, leaving the bulk of it intact at any given time. This is where local plans enter the picture. The significance of local plans is that they represent the building blocks for master plans. We construct master plans by constructing local plans and merging them together.

Earlier, we encountered the purely logical problem of how to evaluate a newly constructed local plan, given that we must take account both of its effect on the agent's other plans and the effect of the agent's other plans on the new plan. We are now in a position to propose a preliminary answer that question. The only significance of local plans is as constituents of the master plan. When a new local plan is constructed, what we want to know is whether the master plan can be improved by adding the local plan to it. Thus when a new plan is constructed, it can be evaluated in terms of its impact on the master plan. We merge it with the master plan, and see how that affects the expected utility of the master plan.

The upshot of all this is that a theory of rational choice becomes a theory of how to construct local plans and use them to systematically improve the global plan — the master plan. I call this *locally global planning*. As a first approximation, we might try to formulate locally global planning as follows. Let us define the *marginal expected utility* of the local plan P to be the difference its addition makes to the master plan *M*:

$$\mathbf{MEU}(P,M) = \mathbf{EU}(M+P) - \mathbf{EU}(M).$$

If the marginal expected utility is positive, adding the local plan to the master plan improves the master plan, and so in that context the local plan is a good plan. Furthermore, if we are deciding which of two local plans to add to the master plan, the better one is the one that adds more value to the master plan. So viewed as potential additions to the master plan, local plans should be evaluated in terms of their marginal expected utilities, not in terms of their expected utilities simpliciter.

It is not quite accurate to say that it is rational to adopt a plan iff its marginal expected utility is positive. This is for two reasons. First, adding a new plan may only increase the expected utility of the master plan if we simultaneously delete conflicting plans. For example, suppose I have adopted the plan to barbecue lamb chops for dinner. Then I remember that I have chicken in the refrigerator, and so I construct the new plan of roasting chicken for dinner. I cannot improve the master plan by simply adding the latter local plan to it. That would result in my making two dinners but eating only one, and so would lower the expected utility of the master plan rather than raising it. To improve the master plan I must simultaneously delete the plan to barbecue lamb chops and add the plan to roast chicken.

Second, plans may have to be added in groups rather than individually. Recall again the example of transporting the gold and silver to a common destination in a truck. The plan to deliver the gold and silver on a single trip, by virtue of achieving both goals (and taking account of the possible damage to the truck), had a higher expected utility than any single plan with which it competes, e.g., the plan to deliver the gold without delivering the silver. What is better than adopting the plan to deliver them both on a single trip is adopting the two separate plans to deliver the gold on one trip and deliver the silver on another trip.[8] So suppose I first adopt the plan to deliver both the gold and the silver on a single trip. Then it occurs to me that I could make two trips. The change I should make to the master plan at that point involves deleting the plan to deliver the gold and silver on a single trip, and adding two other plans — the plan to deliver the gold on one trip and the plan to deliver the silver on another trip.

In general, a change to the master plan may consist of deleting several local plans and adding several others. Where *M* is a master plan and *C* a change, let $M\Delta C$ be the result of making the change to *M*. We can define the *marginal expected utility of a change C* to be the difference it makes to the expected utility of the master plan:

$$\mathbf{MEU}(C,M) = \mathbf{EU}(M\Delta C) - \mathbf{EU}(M).$$

---

[8] Of course, it could first merge the two plans into a single plan and then add the single plan, but that would be pointless.

The *principle of locally global planning* can then be formulated as follows:

> **Locally Global Planning**
>
> It is rational to make a change *C* to the master plan *M* iff the marginal expected utility of *C* is positive, i.e., iff **EU**($M\Delta C$) > **EU**(*M*).

This is my proposal for a theory of rational decision-making. The theory has two parts: (1) it is rational to perform an action iff it is prescribed by a rationally adopted master plan; and (2) a master plan is adopted rationally iff it is the result of incremental updating in accordance with the principle of locally global planning. I propose this as a replacement for the optimality prescription. It captures the basic insight that rational decision makers should guide their activities by considering the probabilities and utilities of the results of their actions, and it accommodates the observation that actions must often be selected as parts of plans and the observation that optimality cannot be defined in such a way that practical deliberation can be viewed as a search for optimal solutions to practical problems. A decision maker should be an evolutionary planner, not an optimizing planner. The principle of locally global planning tells us how evolutionary planning should work. It involves a fundamental change of perspective from prior approaches to rational choice, because decision-making becomes a non-terminating process without a precise target rather than a terminating search for an optimal solution.

7. Ongoing Work — Constructing a Planning Algorithm

The principle of locally global planning constitutes the framework of a theory of rational decision-making for realistically resource-bounded agents. Many of my objections to the optimality prescription turned on its computational infeasibility. But why should we expect locally global planning to be any more computationally feasible? This is a large issue, and I cannot address it satisfactorily in this paper. It will be taken up at more length in my (2005). But it may be useful to give at least a brief account of why I think we can expect this general approach to planning and decision-making to lead to incremental improvements to the master plan. This can be justified if we make four defeasible assumptions. I will refer to these as the *pivotal planning assumptions*:

Assumption 1: The process of constructing "crude local plans" produces plans that will normally have positive expected utilities.

Assumption 2: Ordinarily (but certainly not always), the expected utility of the result of merging two plans will be the sum of the expected utilities of the two plans.

Assumption 3: Computationally feasible reasoning procedures will reveal those cases in which the second assumption fails.

Assumption 4: There will be "repair techniques" that can often be used to modify either the local plans or the master plan in such a way as to remove the destructive interference leading to the failure of the second assumption without having to replan from scratch.

Given the pivotal planning assumptions, the planning agent can begin the construction of the master plan by constructing a single local plan having a positive expected utility, and take that to be the master plan. Then the agent can systematically construct further local plans with positive expected utilities, and on the basis of the second assumption it can be assumed defeasibly that each time one of them is merged with the existing master plan, the result will be a master plan with a higher expected utility. On the basis of the third assumption, rational investigation will enable the agent to discover those cases in which the defeasible assumptions fail. This amounts to discovering destructive interference. The fourth assumption tells us that it will often be possible to refine the local plan and/or the master plan so as to avoid the destructive interference, thus leading to a modification of the original plans which, when merged, produces a master plan with a higher expected utility than the original master plan. In this way we avoid having to replan from scratch. By proceeding in this way, a rational agent can systematically evolve progressively better master plans.

But why should we accept the pivotal planning assumptions? They will be defended at length in my (2005), but here I can give only a very brief (and no doubt unsatisfying) sketch of why I think they are true. I take it that the first assumption is independently plausible, so I will say nothing further about it. The second pivotal planning assumption is perhaps the most controversial. It turns on two subsidiary assumptions. The first is that it is defeasibly reasonable to expect the probabilities of outcomes of acts to remain unchanged when they are embedded in larger contexts. This is a kind of probabilistic principle of indifference. Such principles have played important roles in the foundations of probability theory, and the principle that is employed here is that of "non-classical direct inference", investigated at length in my (1990). The second subsidiary assumption is a similar one regarding utilities. It is assumed defeasibly that the utility of an outcome does not change when it is embedded in a larger context. Clearly, this often fails. The utility of ketchup combined with hamburger is not the sum of the individual utilities. There are lots of small-scale interactions, but on a large scale we tend to get independence. For example, the utility of having a hamburger for lunch is not affected by whether I vacation in Brazil this summer. One way of defending the general defeasible assumption is explored in my (2001), although it will be defended somewhat differently in my (2005). Given these two subsidiary assumptions, it follows that we can defeasibly expect the expected utility of a plan to not change when it is added to a master plan, and hence it is defeasibly reasonable to expect its contribution to the master plan to be its expected utility in isolation.

The third and fourth assumptions are analogues of assumptions made in classical deterministic planning. For some kinds of planning algorithms, they are provable.[9] Their defense for decision-theoretic planning will depend upon the details of the planning algorithm, and that has yet to be determined, but I presume that they are at least plausible. However, until the planning algorithm is developed, they must be regarded as a promissory note. Current research in the OSCAR Project[10] is aimed at the construction of an implemented decision-theoretic planning algorithm that satisfies these four assumptions.[11] The details of constructing such an algorithm are going to be complex, and they are the focus of ongoing research.

8. Conclusions

I have argued that the optimality prescription founders on an uncritical appeal to alternatives. The optimality prescription would only be reasonable if rational choices were made from small precompiled sets of alternatives. Pursuing the question of what makes actions alternatives led us to the more fundamental observation that, in general, actions cannot be chosen in isolation. Actions can both interfere with each other, and collaborate to achieve goals cooperatively. To accommodate this, actions must be chosen as parts of plans.

We cannot save the optimality prescription by adopting a simple plan-based decision theory — according to which it is rational to adopt a plan iff it is an optimal plan from a set of alternatives. The problems for simple plan-based decision theory are three-fold. First, it is unlikely that there will, in general, be such things as optimal plans. Second, because plans cannot generally be compared just by comparing their expected utilities, optimality may not even be desirable. Finally, even in those unusual cases in which there are optimal plans and optimality is desirable, finding them will be computationally intractable because it would require surveying and comparing infinitely many plans.

The upshot is that rational deliberation cannot be expected to produce optimal plans. A decision maker should be an evolutionary planner rather than an optimizing planner. An evolutionary planner finds good plans, and replaces them by better plans as they are found. The concept of a "good plan" and a "better plan" were analyzed in terms of master plans, with the result that the objective of rational deliberation should be to find an acceptable master plan and to be on the continual lookout for ways of improving the master plan.

Real decision makers will not be able to construct master plans as the result of single planning exercises. Master plans are too complex for that. The master plan must instead be constructed incrementally, by engaging in local planning and then merging the local plans into the master plan. The result is the theory of locally global planning.

---

[9] See my (1999) for a further discussion of this.

[10] http://www.u.arizona.edu/~pollock/

[11] For other work on decision-theoretic planning algorithms, see Blythe and Veloso (1997), Boutelier et al (1999), Haddawy and Hanks (1990), Ngo, Haddawy and Nguyen (1998), Onder and Pollack (1997, 1999), Onder, Pollack and Horty (1998), and Williamson and Hanks (1994).

References

Blythe, Jim, and Manuela Veloso

1997 "Analogical replay for efficient conditional planning", *AAAI9*.

Boutelier, Craig, Thomas Dean, and Steve Hanks

1999 "Decision theoretic planning: structural assumptions and computational leverage", *Journal of AI Research* **11**, 1-94.

Cherniak, Christopher

1986 *Minimal Rationality*. Cambridge: MIT Press.

Gibbard, Alan and William Harper

1978 "Counterfactuals and two kinds of expected utility", in *Foundations and Applications of Decision Theory*, ed. C. A. Hooker, J. J. Leach and E. F. McClennen, Reidel, Dordrecht, 125-162.

Haddawy, Peter, and Steve Hanks

1990 "Issues in decision-theoretic planning: symbolic goals and numeric utilities", *Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, 48-58.

Lewis, David

1981 "Causal decision theory", *Australasian Journal of Philosophy* **59**, 5-30.

Ngo, L., P. Haddawy, H. Nguyen

1998 "A modular structured approach to conditional decision-theoretic planning", *AIPS'98*, 111-118.

Onder, Niluger, and Martha Pollack

1997 "Contingency selection in plan generation", *ECP97*.

1999 "Conditional, probabilistic planning: a unifying algorithm and effective search control mechanisms", AAAI 99.

Onder, Niluger, Martha Pollack, and John Horty

1998 "A unifying algorithm for conditional, probabilistic planning", *AIPS1998 Workshop on Integrating Planning, Scheduling, and Execution in Dynamic and Uncertain Environments*.

Pollock, John

1990 *Nomic Probability and the Foundations of Induction*, Oxford University Press.

1992 "New foundations for practical reasoning", *Minds and Machines*, **2**, 113-144.

1999 "The logical foundations of goal-regression planning in autonomous agents", *Artificial Intelligence* **106**, 267-335.

2001 "Evaluative cognition", *Nous*, **35**, 325-364.

2002 "Causal probability", *Synthese* **132** (2002), 143-185.

2003 "Rational choice and action omnipotence", "Rational choice and action omnipotence", *Philosophical Review* **111**, 1-23.

2005 *Thinking about Acting: Logical Foundations for Rational Decision-making*. In preparation. Currently available on my website at http://www.u.arizona.edu/~pollock/.

Savage, Leonard

1954 *The Foundations of Statistics*, Dover, New York.

Simon, Herbert

1955 "A behavioral model of rational choice", *Qart. J. Economic*s, **59**, 99-118.

Skyrms, Brian

1980 *Causal Necessity*, Yale University Press, New Haven.

1982 "Causal decision theory", *Journal of Philosophy* **79**, 695-711.

1984 *Pragmatics and Empiricism*, Yale University Press, New Haven.

Sobel, Howard

1978 *Probability, Chance, and Choice: A Theory of Rational Agency*, unpublished paper presented at a workshop on Pragmatism and Conditionals at the University of Western Ontario, May, 1978.

Weld, Daniel
1999 “Recent advances in AI planning”, *AI Magazine* **20**, 55-68.
Williamson, Mike and Steve Hanks
1994 “Beyond symbolic goals: a step toward utility-directed planning”, *AAAI Spring Symposium on Decision-Theoretic Planning*, 272-278.