(version of 11/15/04)

Rational Decision-Making In Resource-Bounded Agents[1]

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

Abstract

The objective of this paper is to construct an implementable theory of rational decision-making for cognitive agents subject to realistic resource constraints. It is argued that decision-making should select actions indirectly by selecting plans that prescribe them. It is also argued that although expected values provide the tool for evaluating plans, plans cannot be compared straightforwardly in terms of their expected values, and the objective of a realistic agent cannot be to find optimal plans. The theory of Locally Global planning is proposed as a realistic alternative to standard "maximizing" theories of rational decision-making.

1. Rational Cognition

Rational agents think about the world, evaluate various aspects of it, reflect upon how they might make it more to their liking, and act accordingly. Then the cycle repeats. This is the *doxastic-conative* loop, diagrammed in figure one. The defining characteristic of cognitive agents is that they implement the doxastic-conative loop by thinking about the world and acting upon it in response to their deliberations. Both human beings and the autonomous rational agents envisaged in AI are cognitive agents in this sense.

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

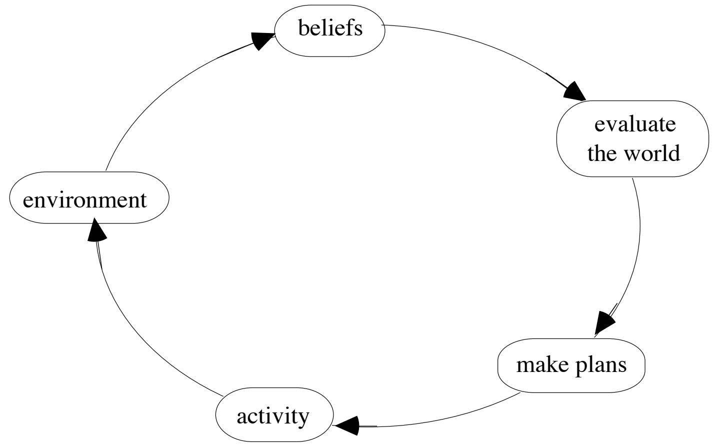

Figure 1. The Doxastic-Conative Loop

The cognition of a cognitive agent can be divided roughly into two parts. *Epistemic cognition* is that kind of cognition responsible for producing and maintaining beliefs. *Practical cognition* evaluates the world, adopts plans, and initiates action. We can further divide practical cognition into three parts: (1) the evaluation of the world as represented by the agent's beliefs, (2) the selection of actions or plans aimed at changing it, and (3) the execution of the plans.

A theory of rationality is a theory about how a cognitive agent should perform these cognitive tasks. Just as cognition divides roughly into epistemic cognition and practical cognition, so rationality divides roughly into epistemic rationality and practical rationality. Epistemology studies epistemic rationality, and I have written about that extensively elsewhere.[2] The focus of this paper is practical rationality. I want to know how a cognitive agent should go about deciding what actions to perform. An answer to this question constitutes a *theory of rational choice*. So this is a book about rational choice.

My principal concern is with human decision making. I want to know how we, as human beings, should decide what actions to perform. However, idiosyncratic features of human psychology sometimes obscure the logic of rational decision making, and we can often clarify the issues by focusing more broadly on rational decision making in any cognitive agent, human or otherwise. Humans are the most sophisticated cognizers we currently know about, but we can usefully ask how *any* cognitive agent should go about deciding how to act. The results of this investigation should be as applicable to the construction of artificial rational agents in AI as to human beings. The advantage of taking this broader perspective is that it can sometimes be argued that purely computational considerations illuminate issues in the theory of rational choice, showing that theories motivated by thinking specifically about human beings cannot be

[2] See particularly my (1986, 1995) and Pollock and Cruz (1998).

correct for any cognitive agents, and so in particular they cannot be correct for human beings.

2. Ideal Rationality and Real Rationality

How should I, and cognitive agents like me, go about making rational decisions? That is the fundamental question of the theory of rational action. 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 action 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 "theory of rational action", 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. 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 surely 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 this is what I am calling "real rationality". We might try to illuminate real rationality by taking it to be some kind of approximation to ideal rationality, but still our original interest is in real rationality.

Although theories of ideal agents are not directly about how real agents should solve decision problems, a plausible suggestion is that the rules of rationality for real agents should be such that, as we increase the reasoning power of a real agent, insofar as it behaves rationally its behavior will approach that of an ideal rational agent in the limit. This is to take theories of ideal rationality to impose a constraint on theories of real rationality. We can make this suggestion more precise by distinguishing, as I have elsewhere (1986, 1995), between "justified" choices and "warranted" choices. A *justified* choice is one that a real agent could make given all of the reasoning it has performed *up to the present time* and without violating the constraints of rationality. A *warranted* choice is one that would be justified if the agent could complete all possibly relevant reasoning. Two characteristics of real agents make this distinction important. First, for any cognitively sophisticated agent, reasoning is non-terminating. There will never be a point at which the agent has completed all the reasoning that could possibly be relevant to a decision. But agents have to act. They cannot wait for the completion of a non-terminating process, so decisions must be made on the basis of the reasoning that has been done so far. In other words, real agents must act on the basis of justified choices rather than waiting until they know that a choice is warranted. Second, it is characteristic of the reasoning of a real agent that almost all of its conclusions are drawn *defeasibly*. That is, the reasoning to date can make the conclusion justified, but acquiring additional information or performing additional reasoning may rationally necessitate the agent's changing its mind.[3]

For an agent that reasons defeasibly, we can characterize a warranted choice as one that, at some stage of its reasoning, the agent could settle on and never subsequently have to change its mind no matter how much additional reasoning it might perform. This can be made more precise by talking about "stages of reasoning". The agent starts from some initial epistemic situation, and then at each stage of reasoning it either draws a new conclusion or retracts a previous conclusion. A conclusion (or choice) is warranted iff there is a stage such that (1) it is justified at that stage, and (2) it remains justified at all subsequent stages of reasoning.[4] Although warranted choices could never be overturned by further reasoning, note that a warranted choice might still have to be retracted in the face of new information.

The warranted choices are those an ideal agent that was able to perform all relevant reasoning would make on the basis of the information currently at its disposal. One might suppose that warranted choices are those we want an agent to make. The difficulty is that a real agent cannot complete all the reasoning that might possibly be relevant to a decision. As remarked above, reasoning is a non-terminating process. Eventually the agent has to act, so we cannot require that it act only on the basis of warranted choices. The most we can require is that the agent perform a "respectable amount" of reasoning, and then base its choice on that. So a real agent acts on the basis of justified choices that might not be warranted.

In some cases it would actually be irrational for a real agent to make the warranted choice. For instance, suppose *P* and *Q* are logically equivalent, but the agent has not yet performed enough reasoning to know this. Suppose the agent has good reason to accept a bet that *P* is true at 2:1 odds. Suppose that choice is not only justified, but also warranted. Suppose, however, the

[3] For the most part, it will be unimportant in this book exactly how defeasible reasoning works. I have, however, discussed it at length elsewhere. See my (1995, 2002), and Pollock and Cruz (1998).

agent has no basis for assessing the probability of *Q*. That is, it has no justified beliefs about the probability of *Q*. Then it would be irrational for the agent to accept a bet that *Q* is true at 2:1 odds. That choice would not be justified. But it would be the warranted choice, because if the agent performed *enough* reasoning it would discover that *Q* is equivalent to *P* and hence has the same probability.

Theories of ideal agents are theories of warrant. What is the relationship between theories of justification and theories of warrant? If we suppose that there are always warranted choices to be made in any decision problem, then it will be true automatically that if the agent has sufficient memory and lives long enough, the set of justified choices approaches the set of warranted choices in the limit. This is because, by definition, a warranted choice is one that become justified if the agent does enough reasoning (reaches a late enough stage of reasoning), and there will be a stage in its reasoning at which it is justified and will never subsequently become unjustified on the basis of further reasoning. A real agent (with enough memory and a long enough life span) may never know that it has found a warranted choice, but it will nevertheless find one eventually.

It might be suggested that the behavior of an ideal agent is the target at which real agents should aim, and hence theories of real rationality can be evaluated in terms of whether they approach the correct theory of ideal rationality in the limit. More precisely, a theory of real rationality, viewed as a theory of justified choice, implies a theory of warrant. We can think of a theory of ideal rationality as a theory of what the correct theory of warrant should say. The suggestion would then be that a theory of justified choice (real rationality) is correct iff its implied theory of warrant describes the behavior of an ideal rational agent (given some theory of what ideal rationality requires).

For epistemic cognition, real rationality and ideal rationality might be related in some such fashion, but it will turn out that there can be no such connection in the case of practical cognition. The set of justified choices will only converge to the set of warranted choices if there are always warranted choices to be made. But it will emerge below that there may often be no warranted choices for real agents living in the real world. It could be that no matter how good a solution the agent finds for a decision problem, given enough time to reason there is always a better solution to be found. I will argue that this need not be an untoward result. The supposition that there must always be warranted choices turns on a misunderstanding of the logical structure of practical cognition — it assumes that decision problems always have optimal solutions. If they do not, then theories of warrant would seem to be irrelevant to theories of justified decision making.

[4] There are two different concepts of warrant here. For a discussion of their interconnections, see chapter three of my (1995).

So our target is a theory of real rationality — a theory of how real agents, with all their cognitive limitations, should make decisions about how to act. A theory of ideal rationality might conceivably be relevant to the construction of such a theory, somehow imposing constraints on it, but a theory of ideal rationality by itself cannot solve the problem of producing a theory of real rationality.

3. Classical Decision Theory

Rational 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 $\mathbf{A}$ 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 $\mathbf{O}$ of pairwise exclusive and jointly exhaustive outcomes. We further assume that we know the probability $\mathbf{PROB}(O/A)$ of each outcome conditional on the performance of each action. Finally, we assume a *utility-measure* $\mathbf{U}(O)$ 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:

$$\mathbf{EU}(A) = \sum_{O \in \mathbf{O}} \mathbf{U}(O) \cdot \mathbf{PROB}(O/A).$$

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-value. 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.[5] One class of problems has given rise to several varieties of *causal decision theory*, which require **PROB** to be a kind of "causal probability".[6] 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.[7] 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

---

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

[6] 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.

[7] They are discussed at length in my (2003, 2005).

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 resulting theory differs from classical decision theory in major respects.

4. Decision-Theoretic Alternatives

Classical decision theory is formulated in terms of a given set of "alternative" actions. It is a theory about how to choose between the alternatives. But it doesn't tell us how to determine what actions are alternatives. Without that, the theory cannot be applied to the real world.

It might be supposed that any set of actions can constitute the alternatives — it is up to the agent. The alternative actions are just those the agent is considering performing. An obvious problem for this proposal is that at any given time an agent may be considering many different actions and may choose 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 cannot do both. These actions are not in competition with each other.

Apparently the actions in a set of alternatives must be "competing actions" that I must choose between. What is it for actions to compete? It is sometimes supposed that competing actions are those that cannot be performed together. Let us call these *strongly competing* actions. But strong competition is too strong. 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 example, suppose I am considering baking a cake this afternoon, and also cleaning house. I could do both, but if I did I would be so tired that the expected utility of performing both actions is less than the expected utility of performing just one. Then surely, they are in competition and should be considered alternatives. 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 does not generate a theory with quite the same structure as classical decision theory. The problem is that classical decision theory 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 other actions together. For instance, suppose I am planning a wedding. Folk wisdom dictates that I should to select something borrowed and something blue, but 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 can be captured by talking about 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 classical decision theory takes this form. The problem I want to focus on is then that of characterizing competition between actions rather than that of characterizing the set of alternative actions.

5. 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.[8]

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 cooperate to achieve goals collaboratively. 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 what we call *plans*.

[8] 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.

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 decision theory to make the right prescription in the above example is to apply it to plans rather than individual actions. 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 both the gold and silver on a single trip, and that is the basis upon which it is chosen.

To recapitulate, classical decision theory is about how to choose between competing actions. To apply this theory, we need an account of what determines the scope of a choice. In deciding whether to perform an action, what is the range of competing actions we must compare it to? 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 extend classical decision theory by applying it 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).

6. Actions and Plans

Before pursuing the details of the proposal that we extend classical decision theory by applying it to plans, a digression is in order. I will argue that on a sophisticated view of actions, classical decision theory is already committed to this. Let us begin by noting that the distinction between actions and plans is not a clean one. Actions can be very plan-like. Consider the action of *making a cup of tea*. How do you make a cup of tea? Let us say that it consists of the following: (1) putting water in the teapot, (2) heating the water, (3) retrieving a tea bag, (4) placing the tea bag in a teapot, (5) pouring boiling water into the tea pot, (6) letting the concoction sit for several minutes, (7) pouring the tea from the teapot into a cup. The latter is a plan, but it is hard to see what the difference is between the action of *making a cup of tea* and this seven-step plan.

To clarify the terminology, I will say that *acts* are individual spatio-temporally located performances (I mean to include mental acts here). I will take *actions* to be act-types. In rational decision-making it is actions we are deliberating about. That is, we are deciding what type of act to perform.

I often perform one act *by* performing another. Goldman (1970) called this "level generation". High level acts are performed by performing one or more lower level acts. E.g., I make a cup of tea by performing the sequence of acts described above. In turn, I put water in the teapot by picking it up, putting it under the tap, turning on the water, waiting until the water reaches the appropriate level in the pot, turning off the tap, and setting the teapot down on the counter. We can progress to lower and lower levels of acts in this way, but eventually we will reach acts like *grasp the handle of the teapot* or *raise my arm* that I can perform "directly" — without performing them by performing some simpler act. The literature on action theory defines *basic acts* to be acts that are not performed by performing another act. I will define a *potentially basic action* to be be an act type that *can* be instantiated by a basic act. An example is *wiggle your finger*. Note, however, that potentially basic actions can typically be performed in nonbasic ways as well. For example, I can wiggle my finger by grasping it with my other hand and moving it up and down.

When the act instantiating an action is not a basic act, you perform the action by doing something else. The "something else" is typically a sequence of lower-level actions, and that sequence constitutes a plan. Roughly following the terminology used in "hierarchical task decomposition planning" in the artificial intelligence literature, let us say that a plan *decomposes* an action iff one can perform the action by executing the plan. Perhaps the main difference between an action and a plan that decomposes it is that there is generally more than one way to perform the action, i.e., actions can have multiple decomposition plans, allowing them to be performed in different ways on different occasions. In fact, for some actions there appears to be an unbounded set of decomposition plans. Consider the action of *traveling from Tucson to LA*. There are infinitely many ways of doing that if we consider all the different means of transportation and all the different routes that can be taken. Of course, on any given occasion we only perform the action by executing a single one of the decomposition plans.

The set of decomposition plans available for a particular action is open-ended. It is not fixed by the meaning of the action term. Rather, the agent can invent new ways of performing an action. For example, if I want to turn on the light, and I am sitting down and unable to reach the light switch from my chair, but there is a ski pole on the table beside me, I might turn on the light by poking the switch with the ski pole. This is a new decomposition plan that I have never previously considered.

What determines the range of possible decomposition plans for an action? The meaning of an action term often specifies the act type partly in terms of what it accomplishes. For example, the action of *making a cup of coffee* results in my having a cup of coffee. So does the action of *buying a cup of coffee*, so there is more to the action than the goal it aims at achieving. An action is also characterized in part by constraints on how its goal is to be achieved. *Making a cup of coffee* differs from *buying a cup of coffee* in that when I make a cup of coffee, I must brew it myself, whereas when I buy a cup of coffee I must acquire it by purchasing it from someone else. These are constraints on the decomposition plans, but the constraints are not usually sufficient to determine a unique decomposition plan. For example, I can make a cup of coffee using a percolator, or a French press, or pouring hot water into toddy syrup, etc.

What this all suggests is that at least for an important class of actions, the identity of an action is determined by (1) a goal-type, and (2) a constraint on what plans for achieving goals of that type can constitute decomposition plans for the action. Very high-level actions, like *save the* *world*, impose virtually no constraints on their decomposition plans. They are characterized almost exclusively in terms of their goal-types. This is quite different from the way in which philosophers have traditionally thought of actions. They have viewed actions as being much more concrete, although that may in part reflect a conflation of acts and actions.

You can perform a potentially basic action in either of two ways — by performing a basic act that instantiates the action, or by executing a decomposition plan for the action. For actions that are not potentially basic, you can only perform them by executing a decomposition plan. So for most actions, performing them consists of executing a decomposition plan. There is a class of "small world" cases in which we can choose between actions directly, and not on the basis of their being contained in more extensive plans that are adopted in their own right. In these cases, we evaluate actions in terms of their expected utilities. However, it should now be apparent that the expected utility of trying to execute an action will be dependent upon *how* one tries to execute it, i.e., on one's choice of a decomposition plan. I might fly from Tucson to LA via Phoenix or via Cincinnati. The execution costs of these decomposition plans will differ significantly, producing different expected utilities for the plans. So comparing the expected utilities of actions must always be done against some background assumptions about what decomposition plans may be employed. The agent need not have decided precisely what decomposition plan to use, but the range of decomposition plans under consideration must be sufficiently constrained to allow meaningful estimates of expected utilities. The expected utility of the action must be determined by the candidate decomposition plans. The only obvious way to do this is to compute expected utilities for the decomposition plans, and then identify the expected utility of the action with that of its candidate decomposition plans (or perhaps either an average or the minimum of the expected utilities of its candidate decomposition plans). If this is right, then choosing actions by comparing their expected utilities is the same thing as choosing decomposition plans by comparing their expected utilities. So it seems that a sophisticated view of actions commits classical decision theory to comparing plans in terms of their expected utilities.

7. Choosing Between Plans

Classical decision theory is a theory of how to choose between competing actions. We found that to apply the theory, we need an account of what makes actions competitors. That led us to the discovery that actions must often be chosen in groups, i.e., as parts of plans, rather than individually. In other words, rational choices must often be between plans rather than individual actions. It was suggested that we can solve this problem by reconstruing classical decision theory so that it applies to plans rather than actions.

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. We have, in effect, just pushed the problem of characterizing the relation of competition up one level, from actions to groups of actions.

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 observe that a cognitively sophisticated autonomous agent operating in a complex environment is not faced with a single fixed planning problem. First, its beliefs will change as it acquires experience of its environment and as it has time for further reasoning. This will affect what solutions are available for its planning problems. Second, as it acquires more knowledge of its environment, its goals may change. We cannot expect the agent to redo all of its previous planning each time it 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 the "extendability" objection, 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. However, this plan is constructed by merging two unrelated plans for achieving unrelated goals. If we are not allowed to construct such merged plans, decision theory 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.

8. Universal Plans

There is a way of trying to save plan-based decision theory from the extendability objection. The argument that led to the conclusion that plans cannot be selected for adoption just by comparing their expected utilities turned upon always being able to extend a plan by merging it with a subplan for achieving an additional goal. For plans as ordinarily conceived, this assumption is unproblematic. But there is one way of avoiding the argument — consider only "universal plans". These are plans prescribing what the agent should do for all the rest of its existence. Universal plans cannot be extended by adding subplans for new goals. Because universal plans include complete prescriptions for what to do for the rest of the agent's existence, any two universal plans will make different prescriptions, and so will strongly compete. It seems initially quite plausible to suppose that universal plans can be compared in terms of their expected utilities, and that a universal plan is rationally adoptable iff it is optimal, i.e., iff no other universal plan has a higher expected utility.

Savage (1954) toys with the idea that rational decisions should be between universal plans, but he rejects it for the obvious reason. A plan is said to have a *finite time horizon* if it only takes account of states of the world up to some limited time point. In constructing universal plans, it is impossible to justify the restriction to a finite time horizon. For any time chosen, there is presumably some possibility that the agent will continue to exist beyond that time. Cutting the planning off at a time when states that are value-laden are still occurring has the consequence that there is no guarantee that an optimal universal plan for this truncated state space will be part of an optimal plan for a larger state space. Accordingly, universal plans must be of infinite duration, and so can prescribe infinitely many actions. However, real agents cannot compute infinite plans.

It was remarked above that theories of rationality are often taken to be theories of ideally rational agents. According to this point of view, if ideally rational agents have to have infinite memory and computing power, and real agents cannot, then so much the worse for real agents — they are necessarily irrational. One can, of course, choose to employ the term "theory of rationality" in this way, but so conceived it is hard to see what a theory of rationality is good for. At the beginning of the paper, I introduced theories of rational action somewhat differently — they are theories of how a real agent should, rationally, go about deciding what actions to perform at any given time. A theory that requires an agent to do something that is logically impossible cannot be a correct theory of rationality. Real agents cannot compute infinite plans, so the theory of rationality cannot require them to.

9. Evolutionary Planning

Thus far I have drawn two main negative conclusions: (1) actions cannot be chosen just by comparing their expected utilities — we often have to choose actions derivatively by choosing plans that prescribe them; but (2) plans cannot be chosen by comparing their expected utilities either. The upshot is that neither classical decision theory nor simple plan-based decision theory is a correct theory of rational action. 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.

Still, it is hard to let go of the intuition that rational agents should somehow decide what to do by considering the probabilities and utilities of the outcomes of their actions or attempted actions. I am convinced that this vague intuition is correct, but a concrete theory of rational action based upon this intuition is going to have to look rather different from classical decision theory. The rest of the paper will be devoted to formulating a new kind of decision theory that tries to capture this fundamental intuition in another way. In this section I will try to diagnose where classical decision theory goes wrong, and then in the next section I will propose an alternative theory of rational action that avoids these difficulties.

Both classical decision theory and simple plan-based decision theory try to do two things. First, they try to characterize what is the "objectively best thing to do". Given the actual probabilities and utilities of outcomes, they attempt to pick out a unique action or plan (or a set of equally good actions or plans) that is objectively best. These are the *optimal actions* or *optimal plans*. Second, they tell us that rationality requires an agent to choose what he rationally believes to be an optimal action or optimal plan out of all the alternatives available. My main objection to classical decision theory is that actions cannot be chosen in isolation. My main objection to simple plan-based decision theory is that the only optimal plans will be universal plans,[9] and real agents cannot construct or choose between optimal plans.

The question we ultimately want to answer is, "How should a cognizer go about deciding whether a plan should be adopted?" Theories of rationality often assume that an agent can complete all relevant reasoning before deciding how to act. But outside of toy problems, that will never be the case. Assuming that the agent's reasoning about the world involves at least full first-order logic, and more likely some defeasible reasoning about its environment, that reasoning will not produce a recursive set of conclusions and so will in general be non-terminating.[10] Even if the agent is only engaging in deterministic planning, if it has to reason about its environment to detect threats to causal links then the set of threats will not generally be recursive, and I showed in my (1999) that this makes the set of ⟨*problem*,*solution*⟩ pairs not even

---

[9] There is no guarantee that there will even be optimal universal plans. It could be that for each universal plan there is a better one.

[10] I showed in my (1995) that even if the set of reasoning schemas available to a defeasible reasoner are "well behaved", the set of conclusions will only be Δ_2 in the arithmetic hierarchy.

recursively enumerable. In general, reasoning will be non-terminating. There will be no point at which an agent has exhausted all possibilities in searching for plans. Despite this, agents 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. Agents 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.

This point is fairly obvious, and yet it completely changes the face of decision-theoretic planning. The objective cannot be to find optimal plans. First, non-terminating reasoning may produce better and better plans without limit, so there may be no optimal plans. Second, even if there were optimal plans the agent would have no way of knowing it has found one until all of the non-terminating reasoning is completed. Planning and plan adoption must be done defeasibly, and actions must be chosen by reference to the current state of the agent's reasoning at the time it has to act rather than by appealing to the idealized but unreachable state that would result from the agent completing all possible reasoning and planning. Agents 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.[11] The agent's set of adopted plans evolves over time, getting better and better, and the rules for rationality are rules directing that evolution, not rules for finding a mythical endpoint. Accordingly, a decision-theoretic planner should implement rules for continually improving the set of adopted plans rather than implementing a search for the endpoint, i.e., a search for optimal plans. We might put this by saying that a decision-theoretic planner should be an *evolutionary planner*, not an optimizing planner. An evolutionary planner will systematically direct an agents "entire life" rather than just discrete segments aimed at local goals.

10. Locally Global Planning

An evolutionary planner tries to find the best plans it can given the limited amount of time and computational resources it has to devote to the task. There is no way to search for optimal plans, so it will search for good plans. The good plans are "good enough" to act upon, but given more time to reason the evolutionary planner may find better plans and adopt them in place of the original good plans. At each stage the best plans found to date are adopted defeasibly, but they will be withdrawn later if better plans are found.

Let us look more closely at the basis upon which plans are defeasibly adopted. First consider the limiting case in which an agent has no background of adopted plans, and a new plan is constructed. Should the new plan be adopted? The basic insight of classical decision theory is

[11] 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.

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 agent can do is compare the new plan with the other options currently available to it. If this is the only plan the agent 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. This is the same thing as asking whether its expected utility is positive (where now the marginal expected-value is computed by comparing the expected-value of trying to execute the plan with the expected-value of doing nothing at all).

Things become more complicated when the agent 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. For example, if the new plan calls for the agent to perform a feat of strength, the probability of the agent being able to do that may be fairly high. But if other plans the agent has adopted will result in the agent being very tired at that time, then the probability of being able to perform the feat of strength may be lower. So the probabilities can be affected by the context provided by the agent's other plans. The same thing is true of the values of goals. Suppose the new plan is a plan for eating a sumptuous meal. In the abstract, this may have a high value, but if it is performed in a context in which the agent's other plans include participation in a pie-eating contest immediately before dinner, the value of the sumptuous meal will be seriously diminished. Execution costs can be similarly affected. If the new plan prescribes transporting an object from one location to another in a truck, this will be more costly if a previous plan moves the truck to the other side of town.

Clearly, the expected utility of the new plan must be computed "in the context of the agent's other plans". But it is not entirely clear what that means. Roughly, the probabilities and utilities should be conditional on the situation the agent 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 agent 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 agent 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 agent 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 agent's *master plan*. This is simply the result of merging all of the agent's adopted plans into a single plan. A master plan is a *global plan* for achieving all (or as many as possible) of the agent's goals simultaneously.

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 utilities. 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 is natural to think of rationality as directing the search for a master plan. Rational decision-making cannot require the agent to construct an optimal master plan. This is for two, now familiar, reasons. First, there is no reason to expect there to be an optimal master plan. Master plans are by definition finite. Given any finite master plan, there is always the possibility (in fact, generally the likelihood) that given enough time for further planning the agent could construct additional local plans that, when merged with the master plan, would increase its expected utility. The only way to avoid this is to take planning to the limit and construct universal plans, but that is impossible. Second, even if there were an optimal master plan, it is computationally unreasonable to expect resource-bounded agents to be able to find it.

Rather than searching for optimal master plans, rationality dictates that a master plan is defeasibly adoptable iff it has a positive expected utility, and it remains adoptable iff no master plan with a higher expected utility has been found. Rational deliberation is aimed at finding better and better master plans.

If the only way we had of finding a master plan with a higher expected utility than our current one was to plan all over again from scratch and produce a new master plan essentially unrelated to our present master plan, the task would be so formidable as to be computationally impossible. 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:

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

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 adoptable. 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.

The upshot of all this is that rational choice becomes a theory of how to construct local plans and use them to 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:

> It is rational for an agent to adopt a plan iff its marginal expected utility is positive, i.e., iff adding it to the master plan increases the expected utility of the master plan.

However, for two reasons, this formulation is inadequate. The first reason is that adding the 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.

The second reason the preceding formulation is inadequate is that 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 my 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-value 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. So suppose that 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:

$$\textbf{MEU}(C,M) = \textbf{expected utility}(M\Delta C) - \textbf{expected utility}(M).$$

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

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

The principle of locally global planning forms the core of a theory of rational decision-making. This 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 classical decision theory. It captures the basic insight that rational agents 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 makes no sense for plans outside of toy examples. A decision-theoretic planner should be an evolutionary planner, not an optimizing planner. The principle of locally global planning tells us how an evolutionary planner should work.

11. 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.[12] 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[13] is aimed at the construction of an implemented decision-theoretic planning algorithm that satisfies these four assumptions.[14] The details of constructing such an algorithm are going to be complex, and they are the focus of ongoing research.

12. Conclusions

I have argued that the optimality prescription flounders on an uncritical use of alternatives. The prescriptions of classical decision theory would only be reasonable if rational choices were made from small precompiled sets of alternatives. Pursuing the question of what makes actions

---

12 See my (1999) for a further discussion of this.

13 http://www.u.arizona.edu/~pollock/

14 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).

alternatives led us to the more fundamental observations that, in general, actions cannot be chosen in isolation. Actions can both interfere with each other, and cooperate to achieve goals collaboratively. 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 a good master plan and to be on the continual lookout for ways of improving the master plan.

Real agents will not be able to construct master plans as the result of single planning exercises. The 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.

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.
Gibbard, Alan and William Harper
1978 "Counterfactuals and two kinds of expected value", in *Foundations and Applications of Decision Theory*, ed. C. A. Hooker, J. J. Leach and E. F. McClennen, Reidel, Dordrecht, 125-162.
Goldman, Alvin
1970 *A Theory of Human Action*, Princeton University Press.
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.