Report Number: WUCS-89-09

1989-02-23

Two Heuristic Functions for Decision

R. P. Loui

This paper investigates a different foundation for decision theory in which successive model refinement is central. The idea is to modify utility so that it can sometimes be calculated for an outcome without considering all the relevant properties that can be proved of the outcome, and without considering the utilities of its children. We build partially ordered heuristic utility functions. We treat the analysis of personal decision trees like heuristic search of game trees (taking expectations instead of doing minimax). Analysis of decision then becomes a process of constructing and evaluating defeasible arguments for decision. This leads to an... Read complete abstract on page 2.

Two Heuristic Functions for Decision

R. P. Loui

Complete Abstract:

This paper investigates a different foundation for decision theory in which successive model refinement is central. The idea is to modify utility so that it can sometimes be calculated for an outcome without considering all the relevant properties that can be proved of the outcome, and without considering the utilities of its children. We build partially ordered heuristic utility functions. We treat the analysis of personal decision trees like heuristic search of game trees (taking expectations instead of doing minimax). Analysis of decision then becomes a process of constructing and evaluating defeasible arguments for decision. This leads to an iteratively improving computation of decision, or what Dean and Boddy have dubbed an "anytime algorithm" for decision. An axiomatization of this idea is simple in an existing system of defeasible reasoning. As a special case of defeasible reasoning, computing defeat among decision trees is also simple. The axioms for preference that lead to metric utility can be retained if we take the defeasibility to be a result of the epistemic problem of individuating objects of value. We say nothing yet about the specification of actual search strategies for particular forms of heuristic utility functions, though it is clearly a matter for further research.

TWO HEURISTIC FUNCTIONS FOR DECISION

Ronald P. Loui

WUCS-89-09

March 1989

Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899

To be given at the AAAI Spring Symposium on Limited Rationality, Stanford, March 1989.

Two Heuristic Functions for Decision

R. P. Loui *
Dept. of Computer Science and
Dept. of Philosophy †
Washington University
St. Louis, MO 63130
loui@ai.wustl.edu

February 23, 1989

Abstract

This paper investigates a different foundation for decision theory in which successive model refinement is central.

The idea is to modify utility so that it can sometimes be calculated for an outcome without considering all of the relevant properties that can be proved of the outcome, and without considering the utilities of its children. We build partially ordered heuristic utility functions. We treat the analysis of personal decision trees like heuristic search of game trees (taking expectations instead of doing minimax). Analysis of decision then becomes a process of constructing and evaluating defeasible arguments for decision. This leads to an iteratively improving computation of decision, or what Dean and Boddy have dubbed an "anytime algorithm" for decision.

An axiomatization of this idea is simple in an existing system of defeasible reasoning. As a special case of defeasible reasoning, computing defeat among decision trees is also simple. The axioms for preference that lead to metric utility can be retained if we take the defeasibility to be a result of the epistemic problem of individuating objects of value. We say nothing yet about the specification of actual search strategies for particular forms of heuristic utility functions, though it is clearly a matter for further research.

---

*I have benefitted from the comments and direction of many on this topic, including Len Schubert, Mark Drummond, Judea Pearl, and Dana Nau; Eric Horvitz, Ben Grosof, and Stuart Russell; Teddy Seidenfeld, Mike Wellman, Greg Cooper and Ned McClennen; Rich Thomason and Hector Geffner; Guillermo Simari; Jon Doyle and Jack Breese; Mike Fehling and Kathryn Laskey; Paul Weirich; Henry Kyburg and John Pratt. Emergency typesetting with grs.

†Part of this work was completed on postdoctoral support at Stanford University.

Two heuristic functions are exhibited: one based on linear additivity of each factor's individual contribution, and one based on projecting from similar outcomes, the utilities of which have been declared.

1 INTRODUCTION: What's the Basic Idea?

1.1 Philosophy: What's Wrong with Classical Theory That's Fixed Here?

TOO MANY UTILITIES TO WRITE DOWN.

In the real world, there is no such thing as a final outcome or final description of an outcome. If you die, how did you leave the world? If you go broke, who got your money? If you win the ball game, how many relief pitchers will be too tired to play in tomorrow's game? The more realistic the model of the domain, the more factors are taken into account, the more combinations of contingencies are considered. If we think of partitions on the space of outcomes of the world, more detail means more relevant distinctions, which means more individuable outcomes whose utilities must be determined. If we think of decision trees which branch on outcomes of uncontrolled events and also branch on conditional action, then more detail means the horizon at which we can attribute value to an outcome becomes farther away. In the extreme, if indeed there is no such thing as a final outcome, the search for final outcomes, or quiescent states, never terminates. If there is no such thing as a final description of an outcome, then the list of things to try to prove about results of actions is endless (and so are some of the searches for proofs). The more detailed the model, the more our faith in the analysis, but we still want to be able to do decision analysis in reasonable time.

Decision theory remains our only effective tool for trading one desideratum against another when planning action without omniscience about the world.

But how are we to apply decision theory when computational resources are dear? To how much detail and depth in our decision analysis are we prepared to commit when it is unclear how much time we will have to study our model?

Savage, who invented modern decision theory, had the wrong idea for the construction of decision models (L. Savage, *Foundations of Statistics*, Dover, 1950):

> In application of the theory, the question will arise as to which [description of the] world to use . . .. If the person is interested in the only brown egg in a dozen, should that egg or the whole dozen be taken as the world? It will be seen . . . that in principle no harm is done by taking the larger of two worlds as a model of the situation.

Essentially, Savage is suggesting that decision trees be taken to be as tall and as bushy as can be imagined. Meanwhile, utility valuations for outcomes cannot be assigned in arbitrary combinations to nodes in the tree. Decision theory requires as a coherence constraint that the utility of a node equal the appropriate expectation of its children's utilities. This paralyzes those who want to specify their utility functions; their assigments of parents' utilities must wait until they have considered the children. No value assignment without search.

Example.

I am deciding to drive or fly to Detroit for IJCAI. Properties that contribute to my valuations include { *cost-exceeds-$200, time-exceeds-eight-hours, increased-engine-wear, passed-thru-gary-indiana, enjoyed-most-of-the-trip* }. Of course, there are more relevant properties, and the longer I think, the more I can invent, all of which legitimately affect my preference, *e.g.*, { *got-to-see-ann-arbor, met-someone-hot-enroute, had-to-take-ijcai-buses, talked-to-people-on-buses, saw-inflight-movie, zoomed-past-lots-of-ford-probes-and-honda-accords, had-to-get-towed*, . . .}.

The departure of this paper from classical theory is that (1) I want to be able to attribute value to outcomes based on these properties, even though I admit that *got-to-visit-ann-arbor* is only relevant to me because I care about its possible cases: *got-to-see-janie-in-ann-arbor*, or not; and *talked-to-people-on-buses* is only relevant to me because the case, *talked-to-levesque-on-bus*, has utility to me, and the case, *talked-to-some-other-dog-on-bus*, has disutility to me. Ultimately, of course, the only relevant property is how I feel, and even properties such as *seeing-janie-in-ann-arbor* and *talking-to-levesque-on-bus* may not directly fix that property. Moreover, (2) I do not want to commit to saying that my valuation based on the coarse property is equal to the appropriately weighted average of valuations based on the refined properties, because it might not be (I don't know until I do the more refined analysis). And if I presume that it is, I will have no incentive to deepen the analysis, that is, return to this node in the decision tree and study its children.

Transforming the initial situation into outcomes are events and conditional actions such as { *change-oil-before-go, change-oil-enroute, car-breaks-down, bad-heat, find-cheap-flight, busy-interstates, rent-car-in-detroit, no-time-to-stop-in-ann-arbor, seek-out-levesque, seek-out-honda-accords-to-pass, ac-clutch-rattles-again, car-stolen-in-detroit*, . . .}. For a given sequence of actions applied to the initial state, s_0, and events occurring, we have an outcome state about which we can do theorem-proving. For instance,

$$\begin{aligned}&\mathrm{T}(\textit{increased-engine-wear},\\ &\quad < \textit{bad-heat} \mid \sim\textit{change-oil-enroute} \mid \textit{drive} \mid\\ &\quad \sim\textit{change-oil-before-go} \mid s_0 >)\end{aligned}$$

should be familiar to adherents to the situation calculus. Note that it is not obvious what properties do hold in a state, that is, what are all the ramifications of actions and events. Actions and events change conditional probabilities, too, so one might show

$$\text{T}(\ \text{Prob}(\textit{passed-lots-of-probes-and-accords}) = .8,$$
$$< \textit{seek-out-accords-to-pass} \mid \textit{drive} \mid s_0 >),$$

which is essentially a statement about the children of $<$*seek-out-accords-to-pass* | *drive* | $s_0$$>$ (namely, that the child in which passing is done has probability .8).

Assuming all the probabilities can be calculated or represented, there is the problem of calculating or representing the utility function. If there are n boolean properties to distinguish outcomes, there are 3^n utilities to represent (outcomes are not possible worlds; properties can be reported to hold, reported to fail to hold, or not be reported). Classical theory just supposes the existence of a mapping from outcomes to reals; it has nothing to say about its representation in a language, notation, or shorthand. The only relation an outcome bears to another outcome, which imputes a relation of their respective utilities, is the descendant relation. But AI planners describe outcomes, that is, situations, with sentences in a language of first-order predicates. Two outcomes may differ only on the truth of one atomic formula, of those formulae whose truth values are reported. If the difference is unimportant, the utilities of these two outcomes ought to be close. We should exploit such regularities; we should exploit the logical infrastructure of descriptions of outcomes of the world.

WHAT KIND OF SHORTHAND?

Assuming that utility is known and needs only to be specified, there may be a compact representation. The expected utility rule is one impetus of compaction: given the utilities of leaves and assuming the calculability of probabilities, none of the interior nodes' utilities need be given; all can be calculated. From a descriptive standpoint, there may be agents whose reported preferences do not satisfy this regularity. On the other hand, for many agents, preferences might satisfy additional regularities.

We are not necessarily advocating using regularity to trade space for time. Descriptions of utilities can be pre-processed so that they need not be calculated on-line. We are concerned with getting automated decision analysis started in worlds with too many outcomes, irrespective of how the computation proceeds after its start.

Described in the next section are two shorthands for compact representation of utility. There may be others. The adequacy of a shorthand depends on the utility function being represented; some utility functions will still require long descriptions in all convenient notations. The usefulness of each shorthand below is plausible only because each makes use of defeasible rules: rules that admit of explicit exceptions.

A DIFFERENT PICTURE OF DECISION ANALYSIS.

As soon as defeasibility is introduced, a new possibility emerges. Utilities for outcomes can be calculated defeasibly, as provisional values, given the part of the description and the depth of the tree currently manageable. Since valuation is defeasible, it may not be the same as the valuation when more information is taken into account. However, under computational limitation, the defeasible calculation may serve as a heuristic evaluation of an outcome, perhaps to be improved with time, but happy to be had at this resource-limited moment.

In this picture, decision analysis looks like heuristic game playing, with expected utility used instead of minimax to induce parent values from children. This view is commonly held among computer scientists, but with no explication of the nuances.

Game playing has the notion of a final outcome, which we may want to resist in the case of real-world decision analysis. There is often such a thing as a complete game tree, whereas decision trees can always be made more extended. Game playing also uses only one heuristic function, applied to nodes at varying depths. Here, we have a variety of heuristic functions, each of which applies to a node at any depth. The more information about a node that is used in the defeasible calculation of utility, the better the heuristic evaluation of that node. In chess-playing, it is as if we could do a major piece count, or else a count of all pieces, or an analysis of board control, or a count of all pieces and an analysis of board control, and so on.

Defeasible arguments for decision are constructed, and they interfere with each other, defeat each other, and justify their conclusions, just like any defeasible arguments (see J. Pollock, "Defeasible Reasoning," *Cognitive Science 12*, 1987, or R. Loui, "Defeat Among Arguments," *Computational Intelligence 3*, 1987; the idea is very much like Touretzky's competition among paths in inheritance networks, D. Touretzky, *The Mathematics of Inheritance Systems*, Pitman, 1986). At a time, there may be an inclination to choose a particular act, tempered by reason: an argument that says a different act is better. At a later time, a different argument could be found that justifies the inclination. Perhaps the new argument defeats the old one. Later still, a third argument may be found: an argument against the inclination, which interferes with the existing argument. When there is no reason to say one argument defeats the other, we fall back on our inclination. In this way, the dialectic for decision yields an answer at all times: the more time for deliberation, the more the answer is tutored by reasoning. If more reasoning time increases expected performance, in whatever way that may be measured, then this strategy is an anytime algorithm in the sense of T. Dean and M. Boddy ("An Analysis of Time-Dependent Planning," *AAAI* 1988).

1.2 Two Heuristic Functions.

DEFEASIBLY LINEAR-ADDITIVE UTILITY.

In decision theory, one regularity on utility functions that is explored is multi-attribute utility. When attributes can be identified that contribute to utility proportionally, the description of the utility function can be made exponentially more compact. The utility of an n-vector of attribute values $< \ldots p_i \ldots >$ is

$$\sum \{ \; k_i \; u_i(p_i) \; ; \; 0 \leq i < n \} \; ,$$

where u_i is the utility function for the i-th attribute, and k_i is the i-th attribute's weight.

A similar linear-additive model can be given for utility, where each atomic formula specifies an attribute for utility. Attribute values, p_i, can be -1, 0, and 1, respectively, for formula falsehood, formula undecidedness, and formula truth. Alternatively, attributes $\vdash p$ and $\vdash \sim p$ can be taken to be two different attributes, with values 1 and 0. The latter alternative is more likely useful since it's unclear that the value of $\vdash p$ will frequently be the negation of the value of $\vdash \sim p$ (of course, they never co-occur). In either construal of properties as attributes, the multiplicative constants, k_i, are all that need to be determined.

It is implausible to suppose utility functions to be perfectly linear-additive for atomic formulae, that there will be no cancellation or reinforcements of individual contributions when attributes combine. So the rule can be made defeasible.

Example.

What is the utility of an outcome, s, when

T(*cost-exceeds-$200* & *time-exceeds-8-hours* &
~had-to-get-towed & *met-someone-hot-enroute* &
zoomed-past-lots-of-probes-and-accords, s)
i.e.,
T(p_0 & p_1 & $\sim p_2$ & p_5 & p_{10}, s) ?

Let basic contributions to utility be $contr(p_0) = -200$; $contr(p_1) = -150$; $contr(\sim p_2) = 10$; $contr(p_5) = 30$; $contr(p_{10}) = 40$. Cost exceeding \$200 coupled with time exceeding eight hours is particularly annoying, so there is a declared exception to additivity: $contr(p_0 \;\&\; p_1) = -400$. Time exceeding eight hours is ameliorated by meeting someone enroute, so declare $contr(p_1 \;\&\; p_5) = -40$, not the apparent sum, $-150 + 30$.

There are now several arguments for the utility of s, based on the combined contributions of p_0, p_1, $\sim p_2$, p_5, and p_{10}. It is unclear which is the most specific argument for the combined contribution of p_0 & p_1 & p_5, since we can sum -400

with 30, or else sum –200 with –40. It is however clear that an argument based on just the fact that p_{10} holds in s is defeated because it is less specific. And there would be no defeat of the several arguments that the utility of s is at least as small as –240.

This regularity of utility functions is formalized in the next section.

REFERENCE CLASSES FOR UTILITY.

A second compact representation for utility takes some outcomes' utilities to be declared, and other outcomes to be valued by interpolating, or by considering similarity to exemplar outcomes whose utilities are given. Utilities are declared for a non-exhaustive set of exemplar outcomes. One possibility is to consider the utility of an outcome, s, of which

$$\mathrm{T}(\ p_0 \ \&\ p_1,\ s),$$

to be determined by the average utility of exemplars that manifest p_0 and p_1: (let *if-declared*(e) be true when e is an exemplar, and *ud*(e) be the declared utility for e, a partial function on outcomes)

$$\begin{gathered}\sum \{\ \mathit{ud}(e) : \mathit{if\text{-}declared}(e) \ \&\ \mathrm{T}(\ p_0 \ \&\ p_1,\ e\ )\ \} \\ \text{divided by} \\ \#\ \{\ e : \mathit{if\text{-}declared}(e) \ \&\ \mathrm{T}(\ p_0 \ \&\ p_1,\ e\ )\ \}\ .\end{gathered}$$

The more specific the knowledge about s, the smaller the subset of exemplars used to project the utility of s. The idea is to use only the most similar exemplars. If in fact, we know p_2 to hold in s as well,

$$\mathrm{T}(\ p_0 \ \&\ p_1 \ \&\ p_2,\ s),$$

then

$$\begin{gathered}\{\ e : \mathit{if\text{-}declared}(e) \ \&\ \mathrm{T}(\ p_0 \ \&\ p_1 \ \&\ p_2,\ e\ )\ \} \\ \text{is a subset of} \\ \{\ e : \mathit{if\text{-}declared}(e) \ \&\ \mathrm{T}(\ p_0 \ \&\ p_1,\ e\ )\ \}\ ,\end{gathered}$$

and using the average *ud*(.) among this set provides a more specific argument for the utility of s. The problem is that the former set may be empty; there may be no examplar sharing as many properties with s as we know about s. In this case, reference sets such as

$\{\ e : \mathit{if\text{-}declared}(e) \ \&\ \mathrm{T}(\ p_0 \ \&\ p_1,\ e\ )\ \}$, and
$\{\ e : \mathit{if\text{-}declared}(e) \ \&\ \mathrm{T}(\ p_0 \ \&\ p_2,\ e\ )\ \}$, and
$\{\ e : \mathit{if\text{-}declared}(e) \ \&\ \mathrm{T}(\ p_1 \ \&\ p_2,\ e\ )\ \}$,

of non-comparable specificity, compete to determine the utility of s. The idea of reference sets here is analogous to Reichenbach's idea of the reference class for probability (H. Reichenbach, *Theory of Probability*, Berkeley, 1949).

Example.
What should be the utility of an outcome, s, when

T(*cost-exceeds-$ 200* & *time-exceeds-8-hours* &
~*had-to-get-towed* & *met-someone-hot-enroute* &
zoomed-past-lots-of-probes-and-accords, s)

i.e.,
$T(p_0 \,\&\, p_1 \,\&\, {\sim}p_2 \,\&\, p_5 \,\&\, p_{10}, s)$?

Let the stock of declared utilities be very small,
if-declared $= \{e_0, e_1, e_2, e_3\}$

$T(p_0 \,\&\, p_1 \,\&\, {\sim}p_5 \,\&\, p_7, e_0)$
$T({\sim}p_2 \,\&\, p_{10}, e_1)$
$T(p_0 \,\&\, {\sim}p_2 \,\&\, p_8 \,\&\, p_{10}, e_2)$
$T(p_1 \,\&\, {\sim}p_2 \,\&\, p_5 \,\&\, p_6 \,\&\, {\sim}p_{10}, e_3)$
$ud = \{< e_0, -600 >, < e_1, 40 >, < e_2, -100 >, < e_3, -10 >\}$.

There are again several arguments for the utility of s, based on the average of exemplars sharing various combinations of the properties p_0, p_1, ${\sim}p_2$, etc. Projecting from $\{e : T({\sim}p_2 \,\&\, p_{10}, e)\}$, that is, using the average of $ud(e_1)$ and $ud(e_2)$, is better than projecting from $\{e : T({\sim}p_2, e)\}$, that is, using the average of $ud(e_1)$, $ud(e_2)$, and $ud(e_3)$. It is unclear which argument is best among the projections from sets of exemplars satisfying, respectively, $p_0 \,\&\, p_1$, ${\sim}p_2 \,\&\, p_{10}$, $p_0 \,\&\, p_{10}$, and $p_1 \,\&\, {\sim}p_2 \,\&\, p_5$. But again, the arguments for $u(s)$ being in $[-600, 40]$ are undefeated, and this may be enough for comparison of alternatives.

This heuristic is also formalized in the next section.

The interest here is not in the bounding of utility, though that is an unavoidable consequence of arguments that interfere without defeat. When arguments disagree, we are supposed to direct the search for an argument that resolves the disagreement. There may be situations in which there is no resolution of the disagreement; neither proving more properties about s, nor examining its children resolves the disagreement. These are defects in the knowledge about utility; it could be that knowledge provided about tradeoffs among desiderata is very sparse. The more interesting case however, is when plenty is known about utility, but there is not enough time to make use of it all. This forces abstraction – deliberately ignoring some properties – and forces heuristic evaluation of utility.

WHAT CLASSICAL THEORY LACKS THAT IS SHARED HERE.

What is common between these two compact representations of utility is that they allow utility to be calculated for any outcome, no matter how incomplete the description of that outcome, without requiring coherence with expected utility. There are two ways of coping with an incomplete description: (a) the classical way, which is to introduce the probabilities of each case, and (b) to ignore the part of the description that is missing.

In the example above, if utility were not represented for an outcome with exactly the description given, namely, $T(\ p_0\ \&\ p_1\ \&\ p_3\ \&\ \sim p_{10},\ s\ )$, then classical theory would force us to consider those children in the decision tree that in fact have represented utility valuations, and the probabilities of those children. One such child might be s', where

$$T(\ p_0\ \&\ p_1\ \&\ p_3\ \&\ p_4\ \&\ \sim p_5\ \&\ p_6\ \&\ \sim p_{10},\ s'\ )\ ,$$

with probability calculable:

$$\text{Prob}(\ p_4\ \&\ \sim p_5\ \&\ p_6 \text{ given } p_0\ \&\ p_1\ \&\ p_3\ \&\ \sim p_{10}),$$

and with $u(s')$ represented, say $\vdash u(s') = 133$.

In my proposal we can just ignore p_4, p_5, and p_6 until there is time to come back and do justice to these factors in the analysis. We will use expected utility calculations when we can do them. Not all valuations are expectations of ultimate good feeling; some are heuristic evaluations based on heuristically valuable properties.

2 JUSTIFICATION: How Can Such a Strategy Be Justified?

2.1 Defeasible Reasons for Utilities.

Both of these heuristics have simple axiomatizations in an existing system of defeasible reasoning.

We suppose there is a metalanguage in which $\succ\!\!-$ is a 2-relation on sentences, the relation of one sentence being a reason for another. $\ulcorner Px \urcorner$ is the Quine quotation of Px. $\forall x$ is the object-language universal quantifer; (x) is the metalanguage universal quantifier.

Reasons can be grouped into arguments, and arguments can be contrary, they can disagree, interfere, and defeat each other, and they can sometimes justify their conclusions. Several systems have roughly this form (*cf.* H. Geffner, "On the Logic of Defaults," *AAAI* 1988; D. Nute, "Defeasible Reasoning," in J. Fetzer, ed., *Aspects of AI*, Kluwer, 1988). I will presume my own (*op.cit.*, Loui); in particular, I presume that an argument can defeat another because of superior evidence.

Let $sent(f)$ be the sentential form of a function, f: that is, if $f = \{< a, b >, < c, d >\}$ then $sent(f) =$ "$f(a) = b$ & $f(c) = d$".

Heuristic 1.

1.1. $(x)(y)(P)(Q)$.
$\ulcorner contr(P) = x \ \& \ contr(Q) = y \urcorner \ > \!\!\!- \ \ulcorner contr(P \ \& \ Q) = x + y \urcorner$

1.2. $(P)(c)(s)$.
if $c \subseteq contr$ *then* $\ulcorner \mathrm{T}(P,\ s) \ \& \ sent(c) \urcorner \ > \!\!\!- \ \ulcorner u(s) = contr(P) \urcorner$

1.3. $(P)(E)(s)(k)$.
$\ulcorner \mathrm{T}(P,\ s) \ \& \ \mathrm{T}(\mathrm{Prob}(E) = k,\ s) \urcorner \ > \!\!\!-$
$\ulcorner u(s) = \mathrm{u}(<E \ |s>)(k) + u(<\sim E \ |s>)(1 - k) \urcorner$

Heuristic 2.

2.1. $\vdash \forall P$.
$ucalc(P) = \sum \{ \ ud(e) : \textit{if-declared}(e) \ \& \ \mathrm{T}(\ P,\ e \ ) \ \} \ /$
$\# \{ \ e : \textit{if-declared}(e) \ \& \ \mathrm{T}(\ P,\ e \ ) \ \}$.

2.2 $(P)(c)(s)$.
if $c \subseteq ucalc$ *then* $\ulcorner \mathrm{T}(P,\ s) \ \& \ sent(c) \urcorner \ > \!\!\!- \ \ulcorner u(s) = ucalc(P) \urcorner$

2.3. $(P)(E)(s)(k)$.
$\ulcorner \mathrm{T}(P,\ s) \ \& \ \mathrm{T}(\mathrm{Prob}(E) = k,\ s) \urcorner \ > \!\!\!-$
$\ulcorner u(s) = u(< E \ |s >)(k) + u(<\sim E \ |s >)(1 - k) \urcorner$

In each case, the first axiom presents the structure of the shorthand for utility. The second axiom says that the more properties taken into account, the better. The third axiom says that expected utility calculations transmit value from children to parents.

Since disjunctive reasoning is not easy in these non-monotonic systems, lower bound arguments are constructed using

1.4. $(P)(c)(s)(x)$.
if $c \subseteq contr$ *then*
$\ulcorner \mathrm{T}(P,\ s) \ \& \ sent(c) \ \& \ contr(P) \geq x \urcorner \ > \!\!\!- \ \ulcorner u(s) \geq x \urcorner$

and the obvious analogue 2.4; similarly for upper bounds.

1.3 and 2.3 are not entirely satisfactory. They say that an argument that considers children of a node is no less specific than an argument that treats that node as a leaf. This is true even if some property p, which holds in the parent, and was used in its evaluation, was not taken into account in evaluating the children. The difficulty is that p holding in the parent does not entail p holding in children, because of ramifications of actions and events. Just as it is the burden of the control strategy to try to prove important properties first, or in pairs when properties are well-known tradeoffs, it should be the control strategy that attempts to inherit properties from parents to children.

Example.

Suppose heuristic 1. One argument for the utility of s, when T(p_1 & p_2, s), is based on *contr*(p_1) = 12. Another argument is based on *contr*(p_1) = 12; *contr*(p_2) = 5; defeasibly, *contr*(p_1 & p_2) = 17. "T(p_1 & p_2, s) & *contr*(p_1) = 12 & *contr*(p_2) = 5 & *contr*(p_1 & p_2) = 17" is a reason for "$u(s)$ = 17". The argument based on this reason is more specific than the best argument for "$u(s)$ = 12".

2.2 Metric Utility.

What does it mean to say that $u(s) = 17$, defeasibly? If metric utility is to make sense, we must either provide axioms of "defeasible preference" that guarantee representation in the reals, or else provide a translation of the current view into the existing theory of preference.

The latter is easiest. Let "$u(s) = 17$, defeasibly" mean that s is identified, defeasibly, with an object of value of utility 17. Objects of value continue to be totally ordered, and satisfy the axiom of independence. However, outcomes are no longer objects of value. We produce defeasible arguments that a particular outcome is a particular object of value, but we could be wrong. There could be a better argument that the outcome that results from applying a_1 in s_0, <a_1 | ssub0>, should be identified with a different object of value. One argument says that the outcome is the object of value, "p_1, *ceteris paribus*", and the other argument says the very same outcome is a different object of value, "p_1 & p_2, *ceteris paribus*". Since *contr* is defeasible, we need to make a few more distinctions among objects of value. The object of value, "$\{p_1, p_2\}$, *ceteris paribus*" is different from the object of value "$\{p_1$ & $p_2\}$, *ceteris paribus*". The former was built using *contr*(p_1) and *contr*(p_2), the latter was built using *contr*(p_1 & p_2).

Like many alleged deviations of preference that can be mapped into classical preference, this mapping produces unintuitive objects of value (*cf.* the distinction between objects of value *alive* and *alive-for-sure* in R. Jeffrey, "Risk and Human Rationality," Boston Philosophy of Science Series, 1986). But it guarantees that real-valued utility is meaningful. u("$\{p_1, p_2\}$ *ceteris paribus*") means that the unit-valued outcome is preferred as much as a lottery for ["$\{p_1, p_2\}$, *ceteris paribus*" at probability 1/17, and the zero-valued outcome otherwise]. Whether s is the object of value "$\{p_1, p_2\}$, *ceteris paribus*" is an epistemic matter. In fact, heuristic 1 can be made to look very much like presuming $\sim p$

whenever p has not been proved, that is, taking outcomes to be their default objects of value — clearly a claim about epistemics.

Some decision theorists entertain the possibility that utilities can change over time because preferences change, probabilities are conditioned on more introspection, or a different model is deemed appropriate. The present view aims to say what is the structure of such changes. What is the structure in terms of how the agent's regard for s changes with more computation?

3 COMPUTATION: What Has to be Done to Make it Run?

3.1 Simpler than Defeasible Reasoning.

Although general defeasible reasoning is used to axiomatize the heuristics, actual computation with these heuristics is far simpler than general defeasible reasoning. In general, it is not easy to determine whether an argument defeats another argument. Here, if sentences such as p_0 & p_1 & p_2 are forward-chained, the check for specificity is just a sublist check (at least in the propositional case). Consider the relevant parts of *contr*, or *ucalc*, and order them according to specificity, (*e.g.*, $\{<p_1, 12>, <p_2, 5>\}$ is less specific than $\{<p_1 \& p_2, 17>\}$. The interesting, mutually interfering but undefeated arguments correspond to the maximal elements in the order.

Any time a set of mutually interfering but undefeated arguments disagree over the value of $u(s)$, there will be analogous arguments each of which says that $u(s)$ falls within the min and max of the disagreeing values. All of these new arguments will be undefeated in the set of all arguments so far mentioned. So having found the undefeated arguments, bounds are readily constructed. Act a_1 is preferred to act a_2 just in case $\textit{lower-bound}(u(<a_1 \mid s>)) > \textit{upper-bound}(u(<a_2 \mid s>))$.

Finally, suppose one argument corresponds to a decision tree with a certain amount of theorem-proving done at the leaves, and another argument corresponds to a different decision tree with different theorem proving at the leaves. Then the first argument is better than the second whenever the first is everyhwere at least as deep as the second, and the theorems proved at corresponding nodes are at least as plentiful, and somewhere the first tree goes deeper or somewhere the first tree uses more of the theorems proved.

3.2 Control Strategy and Heuristic.

Choice of control strategy is important. In the discussion above, it does not matter how arguments are constructed and in what order properties of a state are attempted to be proved. Nevertheless, we intuit that "performance" can be improved by choosing the right control strategy.

So far, I have justified the preference of more specific decision arguments by appealing to our preference of more specific arguments in general. Die-hard decision theorists might write an expression for the expected value of basing judgement on more properties or more extensive trees. They might claim that there are cases in which specificity is misleading, taking the heuristic value farther from some "true" value. Just as game-tree search can have pathologies with depth, we can have pathologies with both depth and theorem-proving.

One problem is how we are to make sense of pathology when there is no conception of final state. We cannot guarantee good control until we know what is good and bad control.

Another problem is a problem shared by all defeasible reasoning. We should attempt to construct arguments based on what undefeated arguments we have and what interrelations they bear. If there is a single undefeated argument, we should try to produce a counter-argument. If there are mutually interfering undefeated arguments, we should try to resolve the disagreement. So far, no one has described control of such dialectic.

Still another problem is how to use knowledge of tradeoffs among desiderata to avoid proving properties that will not help to discriminate between leading choices of action (*cf.* M. Wellman, "Formulation of Tradeoffs in Planning under Uncertainty," MIT Ph.D. Thesis, 1988).

4 ADDITIONAL BRIEF COMMENTS.

When planning turns to uncertain worlds with known risks, what is the significance of search? No sequence of acts can be ignored; most sequences result in a state that has some chance of jumping into a goal state, with fortuitous outcomes of events. If utility is not 0-1, every state satisfies desiderata to some degree. Decision theory says that all paths to all states should be contemplated, and the path leading to the best states with the highest probabilities is normally preferred. In fact, paths are misnomers, because each sequence of actions can be compiled into a single compound, conditional act, and any sequence of events can be turned into a single event. When planning under known risks, all trees have height 2.

This work attempts to save AI planning from trivial subsumption under decision theory. Compound action is an idealization because there are too many irrelevant actions and it requires search to determine relevant action. Mapping outcomes to utilities is an idealization because determining ramifications of actions requires theorem-proving. Utility is an idealization because it presumes that some descriptions of outcomes will never be refined, that is, their children never analyzed, or else that their children once analyzed will yield values that cohere with their parent's original value. One can think of incoherence due to refined analysis in terms of probability shifts over time, but that masks a computation whose strucutre deserves explication. Classical decision theory simply ignores computation.

In this formalism, AI planning is like constructing arguments for action with a depth-first control strategy for forming decision trees. A reason for executing this plan is that it achieves certain desirable properties under default assumptions of which events occur. An analysis of the tree that contains this path and some of its alternatives would produce a better argument, whether for or against adopting the original plan.

There are two senses in which Herb Simon's idea of satisficing (*The Sciences of the Artificial*, MIT Press, 1969.) can be interpreted in the current proposal. First, we do not adopt a plan because it is optimal, but because it is the best according to the reasons we have constructed. Second, it could be that a reason for doing an act, a_1, is that a_1 achieves our aspiration level, with no reference to alternative acts. Acting under routine risks may also provide reasons for action: that a_1 is no worse than my usual gamble, is reason for doing a_1. Jon Doyle (personal communication) has suggested yet another kind of reason: that I cannot resolve the choice between a_1 and a_2, is reason for doing a_1. Although I have studied quantitative defeasible reasons for choosing actions in this paper, the appeal of embedding decision analysis in defeasible reasoning is clearly more general.

It is also straightforward to extend the discussion to include considerations of specificity among defeasible arguments for probability calculations (*cf.* B. Grosof, "Non-monotonicity in Probabilistic Reasoning," in *Uncertainty in AI v. II*, J. Lemmer, ed., North-Holland, 1987; H. Kyburg, *Probability and the Logic of Rational Belief*, Wesleyan, 1961; J. Pollock, "A Theory of Direct Inference," *Theory and Decision 15*, 1983). So defeasible reasons for acts' expected utilities are based on defeasible reasons for probabilities, defeasible reasons for utilities, and defeasibly defoliated decision trees.

The Bayesian decision theorist's picture of the world is not perfect. Expressive language is not cheap. The process of constructing successively better decision models is not a shrouded mystery, but a process with structure. Taking decision theory to be a part of defeasible reasoning places it on more satisfying foundations and ought finally to bring together the AI work on planning and decision theory's venerable concept of expected utility.