Nondeterministic Action and Dominance: Foundations for Planning and Qualitative Decision

Richmond H. Thomason
Intelligent Systems Program
University of Pittsburgh
Pittsburgh, PA 15260
U.S.A.
thomason@isp.pitt.edu

John F. Horty
Philosophy Department and UMIACS
University of Maryland
College Park, MD 20742
U.S.A.
horty@umiacs.umd.edu

1. Introduction

A common sense dominance argument (1) divides possible outcomes into two or more exhaustive, exclusive cases, (2) points out that in each of these alternatives it is better to perform some action than not to perform it, and (3) concludes that this action is best unconditionally.

Although such arguments are often used, and are convincing when they are used, they are invalid. A classic illustration of the invalidity is the argument for cold-war disarmament.[1]

The informal argument: Either there will be a nuclear war or there won't. If there won't be a nuclear war, it is better for us to disarm because armament would be expensive and pointless. If there will be a nuclear war, we will be dead whether or not we arm, so we are better off saving money in the short term by disarming. So we should disarm.

The argument can be formalized using a payoff matrix like the following one.

The payoff matrix:

	Disarm	Arm
War	-50	-100
Peace	50	0

The fallacy, of course, depends on the assumption that the action of choosing whether to arm or disarm will have no effect on whether there is war or not. As Jeffrey goes on to show, if the probability matrix is

.1	.8
.9	.2

then arming has the greater expected utility.

In the context of a qualitative approach to decision theory, dominance arguments will play a central role, since we can no longer rely on numerical assignments of utilities to actions.

[1]We have taken the example, and the payoff matrix, from [Jeffrey 1983, pp. 8–12].

It therefore becomes crucial to distinguish valid from invalid dominance arguments without resorting to numerical probabilities.

A form of independence that depends on counterfactual conditionals can be used here: in this qualitative sense, a condition B is independent of an action a that is performed if either (1) B is true, and the conditional 'B (even) if a were not performed' is true, or (2) B is false, and the conditional '$\neg B$ (even) if a were not performed' is true.

Thus, the task of developing a qualitative theory that does justice to dominance seems to require an account of temporal counterfactuals. In the following paper, we show how to develop such an account within a generalization of the STRIPS formalism for deterministic, single-agent planning, and explain how it can be used to provide a formalization of dominance.

2. Motivation and background

We begin with the classic STRIPS approach to actions and reasoning about time.[2] This approach assumes from the very start that outcomes are entirely determined by a single agent's actions. Thus, though there may be a role for plan evaluation in which the outcomes of different plans are compared according to the utility of the direct consequences of actions, there is really no role here for the reasoning about risk that is the core of decision theory. However, the need for true decision-theoretic reasoning becomes essential as soon as actions are allowed to have nondeterministic consequences—and, of course, the need for such a generalization has often been noted in the planning literature.[3] As soon as this generalization is made, a simple single-agent strategy (i.e., a STRIPS-like plan consisting of a series of actions) corresponds to a set of fully specified outcome states, or, equivalently, of sets of histories in branching time. So this generalization of STRIPS has to be unified with an account of how the utility of sets of histories can be compared.

Assuming that we know the utilities of histories, the problem then becomes how to extend these point utilities to utilities on sets. Classical decision theory provides a way to do this through the definition of expected utility—which, of course, assumes that a probability distribution over histories is available. A decision-theoretic formalism for planning could simply import these probabilities: either directly, or in a modified form using "orders of magnitude," as in [Tan & Pearl 1994] and [Goldszmidt & Pearl forthcoming]. This alternative is being explored by Judea Pearl and others.

Here, we follow a more radically qualitative approach, which assumes that we are given only a linear preference ordering on histories, and seeks to extend this ordering to a partial ordering over sets of histories. This "utilities lifting" problem is discussed or alluded to, for instance, in [van Fraassen 1972], [Jennings 1974], [Jennings 1985], [Wellman 1988], and [Horty 1994]. But, as far as we can tell, there has never been a systematic attempt to develop a solution to the problem that does justice to the very robust common-sense intuitions that people have for assessing judgments of preference over sets of outcomes—intuitions that seem in many cases to be largely independent of any precise estimate of the probabilities of outcomes. The literature based on classical decision theory occasionally alludes to these arguments, as in [McClennen 1990] and [Luce & Raiffa 1957, Chapter 13], but the arguments do not seem to have been examined extensively in that literature.

[2] See [Fikes & Nilsson 1971], [Lifschitz 1990].

[3] See, for instance, [Allen 1990].

On the other hand, Wellman's study of dominance-driven planning in [Wellman 1988] shows that dominance relations can be very useful in uncertain domains where exact probabilities are not readily available. And the more theoretical approaches to qualitative preference that have been developed in the recent AI literature are also very relevant to our project.[4]

Two simple approaches emerge from the discussions of the problem; the most common is to say that set P is preferred to Q in case for all $h \in P$ there is a better $h' \in Q$. The other, which goes back at least to [Friedman & Savage 1948], says that P is preferred to Q in case for all $h \in P$ and $h' \in Q$, h is at least as good as h' and some $h \in P$ is better than some $h' \in Q$.

Both of these accounts suffer from flaws that are pretty glaring. The first implies that a lottery with a large prize is better than one with a smaller prize, regardless of the odds. The second fails to imply that accepting an outright gift of a large prize is better than not accepting it, as long as there is an independent chance on any given day that I might suffer from heart failure. (The history in which I accept the gift and have heart failure is, presumably, worse than the one in which I do not accept the gift and do not have heart failure.)

Motivated by these flaws, our project attempts to create a more adequate account of dominance that (1) in reaction to the first problem, introduces abnormalities over outcomes (and therefore appeals to ideas from nonmonotonic logic), and (2) in reaction to the second problem, seeks to provide a definition of dominance that takes into account the relevant causal relations of histories. In this paper, we concentrate on the second of these tasks. This study contains no discussion of abnormalities or nonmonotonicity, but is intended to provide a monotonic theory of dominance from which a more adequate nonmonotonic theory can then be developed.

Probably the chief insight of our paper (which as far as we know is entirely new) is that the notion of *action* that is so important in the STRIPS approach provides a very useful basis for providing the necessary causal structure for an adequate dominance definition. The main goal of the material that follows is to present this insight and to articulate it in the form of a formal theory. In the course of developing that theory, we also present a formalism for planning and action that takes into account concurrent and nondeterministic actions. In keeping with our policy of developing the monotonic theory first, this account does not seek to deal with the frame problem or to introduce nonmonotonicity in any way. In this way, it differs from other generalizations of the STRIPS approach, such as [Lin & Shoham 1992] and [McCarthy 1995]. Our account introduces into the formalism the notion of the *possible outcomes* of an indeterministic action, and generalizes the dynamic assumptions of STRIPS by assuming that a successor state is uniquely determined by the actions that are performed in the initial state and their outcomes. This idea seems to provide a very natural and useful generalization of STRIPS. We then show how causal notions that can be added to the action-driven temporal models of this theory can yield a conditional distance relation.

[4]These include [Doyle, Shoham, & Wellman 1991], [Doyle & Wellman 1991], [Doyle & Wellman 1994], [Boutilier 1994b] and [Boutilier 1994a]. This work is less relevant to the focus of the initial project that is presented below (which, as we explain below, concentrates on causal aspects of preference), than it is to the subsequent part of the project that takes into account nonmonotonicity.

3. Motivating the treatment of action and time

We start with the idea, which is now current in many other approaches to action, that time branches. A moment m is a node in branching time. A history is a maximal linearly ordered set of moments.

There is a set $\mathcal{S}$ of states. We associate a state $state(m) \in \mathcal{S}$ with each moment m. We have in mind a language of "fluents" that keeps track of the changing phenomena that bear on action and decision. Each formula of this language is assigned a truth value in each state, relative to a model. For planning and decision purposes, moments that are associated with the same state are equivalent; but we do not rule out the possibility that different moments may have the same associated state. Such moments may differ in ways that don't bear directly on planning and action; for instance, the histories that lead to these moments may differ.[5]

There is a set *Agents* of agents. Agents are unanalyzed primitives; we do not need to inquire into the nature of agents. Unlike some other theories that deal with branching time and action, we treat actions as primitive: there is a set *Actions* of (fully specified) actions. The nature of actions is vital to our project. Starting with the notion of action that is incorporated in the STRIPS model of planning, we will generalize the model to take concurrency and indeterminism into account, while attempting to preserve some of the desirable causal properties that STRIPS actions enjoy.

Intuitively, the actions in *Actions* correspond to the result of appropriately binding all variables of the action types of some reasoning domain. These actions differ if and only if their action types differ, or their action types are the same and they have different variable bindings (considered as tuples of individuals). We assume in this paper that the relevant variable bindings affect argument positions for the agent of the actions, and for other objects that may be involved in the action. Actions have no "localizing" argument positions like time and place. Thus, we should think of the members of *Actions* as action types that are relatively specific. They are not individual occurrences, but types of occurrences that can be multiply instantiated; for instance, the action of unstacking b_1 by $Agent_1$ may be performed many times.

Consider a blocks microworld $World_1$, where as usual there is only a single agent. (So we can suppress reference to the agent.) In this world there might be only three action types, *Pickup*, *PutOnBlock*, and *PutOnTable*. Instantiating the variables associated with these action types yields a *fully specified action*; when we refer without qualification to an action, a fully specified action is intended. In $World_1$, $Pickup(b)$ and $PutOnTable(b)$ will be actions for each block b in the domain. And for every pair of blocks b_1 and b_2, $PutOnBlock(b_1, b_2)$ will be an action. Then $Actions_{World_1}$ would be the set in which all of these actions are collected together.

In most planning formalisms (but see, for instance, [Grosz & Kraus 1993]) there is only one agent. We now generalize the theory to take multiple agents into account.

In models with multiple agents, each action type has an argument position that has to be instantiated with an agent in obtaining fully specified actions. In each model M there will be a function $agent_M$ that takes actions of the model to their agents in the model.[6] When

[5] We assume that aspects of histories that bear on planning and action are recorded in the states.

[6] Though actions can have groups as agents, we have not yet begun to reckon with group agency; in this paper, we assume that all actions have exactly one agent.

reference to a model is clear, we sometimes drop subscripts: e.g., we may write '$agent(a)$' when '$agent_M(a)$' is clearly intended.

In the STRIPS model of action, an agent can only perform one action at a moment, and the repertoire of feasible actions is determined by the state associated with the moment and the preconditions of the available fully specified actions.

We will relax the assumptions of STRIPS in two ways: (1) we will allow many actions to occur at the same moment,[7] and (2) we will allow multiple possible results even when the same actions are performed at a moment. However, in doing this we wish to maintain an important insight of STRIPS: that the global outcome when actions $a_1, \ldots, a_n$ are concurrently performed at m—that is, the state associated with the moment that results when the actions are performed concurrently in m—is determined by the local outcomes of the separate actions. As we will see, this decomposition of global causality into the local results of actions can be useful in formulating decision theory.

To provide the indeterminacy required by (2), we associate a set $A\text{-}Outcomes_M(a, s)$ of possible outcomes with each action a and state s in a model M. These "action-outcomes" should be thought of as imposing constraints on global states; more vividly, they could be thought of as parts of global states.[8] To keep the terminology clear, we refer to local outcomes (outcomes of of actions) as "a-outcomes", and to global outcomes (the resulting outcome states) as "s-outcomes". The fact that an action a is performed in state s (and, let's say, no other actions) no longer suffices to determine a unique global state, or s-outcome. But we assume that, given the a-outcome of a, a unique s-outcome will be forthcoming. In general, when a set A of actions is performed concurrently, we need to take *action-outcome patterns* for the action A into account. Such a pattern, for an initial state s, is a set of pairs of the form $AO = \{\langle a, o\rangle \ : \ a \in A\}$, where A is a jointly feasible set of actions at s, and $o \in A\text{-}Outcomes(a, s)$.

In effect, we are preserving the STRIPS assumption that the only changes that occur are induced by actions. In STRIPS and the action formalisms derived from it, this assumption is enforced by "frame axioms" (or "inertial axioms"), which can either be formalized using some nonmonotonic logic (as in, for instance, [Shanahan 1995]), or by monotonic conditions on actions (as in, for instance, [Schubert 1990]). Here, we do not care how this is done; we simply assume that there is a function *s-outcome* that takes action-outcome patterns into a global state.[9] And in our examples we assume that this function does not violate the usual frame constraints.[10]

We will illustrate the nondeterministic case with a domain in which the actions are coin tosses. Here, the a-outcome of an action is a coin's position after it is tossed; so each action

[7] In the general case, some of these actions may be performed by the same agent; but we also allow concurrent actions by different agents.

[8] The idea of a-outcomes is similar in some ways to the "hidden variables" technique that is used in [Lin 1996] to make indeterministic actions deterministic. Though Lin's approach is deterministic while ours is nondeterministic, we conjeture that the two are equivalent with respect to a language that does not mention the hidden variables or outcomes.

[9] Our treatment of multi-agent domains is similar to the accounts that have appeared in the planning literature; for instance, [Georgeff 1987] and [Lansky 1987]. The treatment of concurrency is especially similar to that of [Reiter 1996]. We had not seen Reiter's paper until a near-final draft of this paper had been written.

[10] This principle of "outcome determinism" is related to the notion of "epistemic completeness" of of [Lin & Shoham 1992]. But (1) it is model-theoretic, not proof-theoretic, and (2) it provides for indeterminism.

has two a-outcomes. The s-outcome of performing an action, or set of concurrent actions (in a world in which the positions of coins are the only fluents) is the configuration of all the coin positions after the action or actions have been performed.

Let $World_2$ be a microworld in which there is only one agent, and the objects consist of two coins, $coin1$ and $coin2$. Each coin can have either of two polarities: heads up or tails up. Here, the possible states, or s-outcomes, could be represented as sets of the form

(1) $\{x, y\}$, where $x \in \{\mathrm{H}_1, \mathrm{T}_1\}$ and $y \in \{\mathrm{H}_2, \mathrm{T}_2\}$.[11]

Let's assume that there are two actions in $World_2$, $toss(coin1)$ and $toss(coin2)$. The a-outcomes of these actions are as follows.

(2) a-outcomes of $toss(coin1)$: $\{\mathrm{H}_1, \mathrm{T}_1\}$
a-outcomes of $toss(coin2)$: $\{\mathrm{H}_2, \mathrm{T}_2\}$

This rendition of a-outcomes and s-outcomes has the advantage of representating s-outcomes as supersets of a-outcomes, so that a-outcomes are, in a sense, literally parts of states. Here, the a-outcomes determine the s-outcomes by invoking the frame constraint that coins that are not tossed remain in their previous positions, while tossed coins receive their a-outcome positions.

Note that a branching-time model can easily be recovered from a model like $World_2$, if we retain the STRIPS assumption that the only transitions between moments are determined by the performance of actions, and that a frame principle of minimal change applies to these transitions. For instance, if a moment is in the state

$$\{\mathrm{H}_1, \mathrm{T}_2\},$$

then two s-outcomes are possible if $toss(coin1)$ is performed:

- One s-outcome of $toss(coin1)$: $\{\mathrm{H}_1, \mathrm{T}_2\}$
- Another s-outcome of $toss(coin1)$: $\{\mathrm{T}_1, \mathrm{T}_2\}$

Obviously, all four states of $World_2$ should be possible if both coin-tossing actions are performed concurrently; but we can't accommodate this possibility until we provide for concurrent actions. We are now in a position to remedy this limitation. At the same time, we will take another feature of STRIPS into account—the assumption that actions have preconditions. We do this by associating with each model M a function that determines a set $joint\text{-}acts_M(s)$ of sets of actions; each set of actions in $joint\text{-}acts_M(s)$ represents a combination of actions that could be performed in s. In STRIPS, $joint\text{-}acts_M(s)$ would consist only of unit sets (unless doing nothing counts as a transition, in which case the empty set would also be allowed). These would be the unit sets $\{a\}$ such that the preconditions of a are satisfied in s. When concurrent actions are allowed, it may happen that there is never interference between different actions: all sets of actions actions are *jointly feasible*. In this case,

$$joint\text{-}acts_M(s) = \{A : A \subseteq Actions_M \text{ and the preconditions of } a \text{ are satisfied in } s \text{ for all } a \in A.\}$$

[11] Think of (1) as a scheme that generates sets of objects, which we call *constraints*, rather than as a set of formulas denoting truth values.

But in general, we can expect more complicated patterns of concurrent actions and their outcomes. This is illustated by the following example.

$World_3$ is a blocks microworld. Associated with this world are:

- Two robot arms arm_1 and arm_2;
- Four blocks b_1, b_2, b_3, and b_4;
- Actions $PickUpAndPutOn(x, y, z)$, where $x, y \in \{b_1, b_2, b_3, b_4\}$ and $z \in \{arm_1, arm_2\}$;
- The fifty-four states consisting of all possible stacking configurations of the four blocks. Of these states, let s_1 be the one in which b_1, b_2, b_3, and b_4 are OnTable, and let s_2 be the one in which b_1 is on b_2, and b_3 and b_4 are OnTable.

The preconditions of $PickUpAndPutOn(x, y, z)$ are that x is clear and y is clear. The unit concurrent action sets are determined as follows.

$\{PickUpAndPutOn(x, y, z)\} \in joint\text{-}acts_{World_3}(s)$ if and only if s satisfies $Clear(x)$ and $Clear(y)$ in $World_3$, where $x \neq y$.

Thus, for instance,

$$\{PickUpAndPutOn(b_3, b_1, arm_1)\} \in joint\text{-}acts_{World_3}(s_1),$$

but

$$\{PickUpAndPutOn(b_2, b_3, arm_1)\} \notin joint\text{-}acts_{World_3}(s_1).$$

Since there are only two arms, if two actions are performed concurrently in $World_3$ then they must be performed by different arms. Moreover, there is a possibility of interference when an attempt is made to perform two actions; this could occur if arm_1 attempts to put b_1 on b_2 in s_2 while at the same time arm_2 attempts to do the same thing.

Notice that this limitation cannot be expressed as a precondition on the actions. We are not speaking here of a case in which one arm moves before the other, but are supposing that the actions are genuinely concurrent. A condition about what the other arm is doing simultaneously is not a *pre*-condition.

It is very natural to speak in such cases, as we have just done, in terms of "attempting" to perform an action. But in using such language, we are speaking in a way that can't be modeled by STRIPS, or even by the extension that we are developing. In STRIPS, to attempt an action is to achieve its results. Even when action is rendered nondeterministic, as in STIT,[12] the consequences associated with the action are those that occur in every alternative in which the action is performed. As we ordinarily speak of action, the goals or postconditions are expressed as defaults which may not always be achieved. Crossing the street is an action, which if initiated may sometimes fail to achieve the goal. We believe that a nonmonotonic extension of the present theory will capture this aspect of action, and provide a bridge to the theories of agency in the literature of linguistic semantics, as in [Dowty 1979]. But this is a task for a later paper. The alternative treatment, which we will adopt here, is to only allow

[12] "STIT" stands for "Seeing to it that"; see [Nuel D. Belnap & Perloff 1988]. In a later paper we will provide a more detailed comparison of the theory we develop here and the STIT approach to action.

feasibility sets of noninterfering concurrent actions. The formalization is imperfect; because if the control systems for the arms are independent, it may well happen that interferences occur, with unpredictable results. But, as we say, we cannot model this in the present framework without cutting the connection between actions and their postconditions.

The two-member a-outcome sets in $World_3$ are then determined as follows.

$\{PickUpAndPutOn(x_1, y_1, z_1), PickUpAndPutOn(x_2, y_2, z_2)\} \in joint\text{-}acts_{World_3}(s)$ if and only if $z_1 \neq z_2$ and $\{z_1 \neq z_2\} \subseteq \{arm_1, arm_2\}$ and s satisfies $clear(x_1)$, $clear(y_1)$, $clear(x_2)$, and $clear(y_2)$ in $World_3$, where $x_1 \neq y_1$, $x_2 \neq y_2$, $x_1 \neq x_2$, $x_1 \neq y_2$, $y_1 \neq x_2$, and $y_1 \neq y_2$.

In $World_3$, this condition has the consequence that a pair of actions can only occur concurrently if they are performed by different arms on disjoint pairs of blocks in s_1. As we have modeled it, $World_3$ is a deterministic world; there is only one a-outcome per action, as in STRIPS.

Putting together together concurrency and indeterminism, imagine a coin microworld $World_4$ with two agents, *Fred* and *Jane*, two coins, *coin*1 and *coin*2, and three action types: *TurnUp* (this action results in a coin being heads up), *TurnDown* (this action results in a coin being tails up), and *Toss*. There are then twelve actions:

$TurnUp(coin1, Fred)$, $TurnUp(coin2, Fred)$, $TurnUp(coin1, Jane)$, $TurnUp(coin2, Jane)$, $TurnDown(coin1, Fred)$, $TurnDown(coin2, Fred)$, $TurnDown(coin1, Jane)$, $TurnDown(coin2, Jane)$, $toss(coin1, Fred)$, $toss(coin2, Fred)$, $toss(coin1, Jane)$, and $toss(coin2, Jane)$.

We will assume that *Fred* has control over *coin*1 and *Jane* has control over *coin*2; the preconditions of six of these actions are then never met. The remaining *Toss* actions have no preconditions. The remaining *TurnUp* actions presuppose that the coin is tails up, and the remaining *TurnDown* actions presuppose that the coin is heads up. Any set A of actions is jointly feasible in s as long as long as the preconditions of each action are satisfied in s, and A does not contain more than one action by the any one agent. The a-outcomes of the actions are as expected.

a-outcomes of $TurnUp(coin1, Fred)$: $\{\mathrm{H}_1\}$

a-outcomes of $TurnUp(coin2, Jane)$: $\{\mathrm{H}_2\}$

a-outcomes of $TurnDown(coin1, Fred)$: $\{\mathrm{T}_1\}$

a-outcomes of $TurnDown(coin2, Jane)$: $\{\mathrm{T}_2\}$

a-outcomes of $toss(coin1, Fred)$: $\{\mathrm{H}_1, \mathrm{T}_1\}$

a-outcomes of $toss(coin2, Jane)$: $\{\mathrm{H}_2, \mathrm{T}_2\}$

It should be clear how a world like $World_4$ generates a branching time model. There are, of course, four possible states:

$$s_0 = \{\mathrm{H}_1, \mathrm{H}_2\}, \qquad s_1 = \{\mathrm{H}_1, \mathrm{T}_2\}, \qquad s_2 = \{\mathrm{T}_1, \mathrm{H}_2\}, \qquad s_3 = \{\mathrm{T}_1, \mathrm{T}_2\}.$$

Suppose that we start at a moment m_0 with associated state s_0. If we allow any jointly feasible combination of concurrent actions, and count the empty set as such a set, there are sixteen action-outcome patterns for s_0. These patterns and their associated s-outcomes for s_0 are as follows. (Remember, in constructing these models, we are applying frame constraints.)

$AO_1 = \emptyset$, $\textit{s-outcome}(AO_1, s_0) = s_0$;

$AO_2 = \{\langle \textit{TurnDown}(\textit{coin}1, \textit{Fred}), \mathrm{T}_1\rangle\}$, $\textit{s-outcome}(AO_2, s_0) = s_3$;

$AO_3 = \{\langle \textit{TurnDown}(\textit{coin}2, \textit{Jane}), \mathrm{T}_2\rangle\}$, $\textit{s-outcome}(AO_3, s_0) = s_1$;

$AO_4 = \{\langle \textit{toss}(\textit{coin}1, \textit{Fred}), \mathrm{H}_1\rangle\}$, $\textit{s-outcome}(AO_4, s_0) = s_1$;

$AO_5 = \{\langle \textit{toss}(\textit{coin}1, \textit{Fred}), \mathrm{T}_1\rangle\}$, $\textit{s-outcome}(AO_5, s_0) = s_2$;

$AO_6 = \{\langle \textit{toss}(\textit{coin}2, \textit{Jane}), \mathrm{H}_2\rangle\}$, $\textit{s-outcome}(AO_6, s_0) = s_0$;

$AO_7 = \{\langle \textit{toss}(\textit{coin}2, \textit{Jane}), \mathrm{T}_2\rangle\}$, $\textit{s-outcome}(AO_7, s_0) = s_1$;

$AO_8 = \{\langle \textit{TurnDown}(\textit{coin}1, \textit{Fred}), \mathrm{T}_1\rangle\langle \textit{TurnDown}(\textit{coin}2, \textit{Jane}), \mathrm{T}_2\rangle\}$,
$\textit{s-outcome}(AO_8, s_0) = s_3$;

$AO_9 = \{\langle \textit{TurnDown}(\textit{coin}1, \textit{Fred}), \mathrm{T}_1\rangle, \langle \textit{toss}(\textit{coin}2, \textit{Jane}), \mathrm{H}_2\rangle\}$,
$\textit{s-outcome}(AO_9, s_0) = s_2$;

$AO_{10} = \{\langle \textit{TurnDown}(\textit{coin}1, \textit{Fred}), \mathrm{T}_1\rangle, \langle \textit{toss}(\textit{coin}2, \textit{Jane}), \mathrm{T}_2\rangle\}$,
$\textit{s-outcome}(AO_{10}, s_0) = s_3$;

$AO_{11} = \{\langle \textit{toss}(\textit{coin}1, \textit{Fred}), \mathrm{H}_1\rangle, \langle \textit{TurnDown}(\textit{coin}2, \textit{Jane}), \mathrm{T}_2\rangle\}$,
$\textit{s-outcome}(AO_{11}, s_0) = s_1$

$AO_{12} = \{\langle \textit{toss}(\textit{coin}1, \textit{Fred}), \mathrm{T}_1\rangle, \langle \textit{TurnDown}(\textit{coin}2, \textit{Jane}), \mathrm{T}_2\rangle\}$,
$\textit{s-outcome}(AO_{12}, s_0) = s_3$

$AO_{13} = \{\langle \textit{toss}(\textit{coin}1, \textit{Fred}), \mathrm{H}_1\rangle, \langle \textit{toss}(\textit{coin}2, \textit{Jane}), \mathrm{H}_2\rangle\}$,
$\textit{s-outcome}(AO_{13}, s_0) = s_0$.

$AO_{14} = \{\langle \textit{toss}(\textit{coin}1, \textit{Fred}), \mathrm{H}_1\rangle, \langle \textit{toss}(\textit{coin}2, \textit{Jane}), \mathrm{T}_2\rangle\}$,
$\textit{s-outcome}(AO_{14}, s_0) = s_1$.

$AO_{15} = \{\langle \textit{toss}(\textit{coin}1, \textit{Fred}), \mathrm{T}_1\rangle, \langle \textit{toss}(\textit{coin}2, \textit{Jane}), \mathrm{H}_2\rangle\}$,
$\textit{s-outcome}(AO_{15}, s_0) = s_2$.

$AO_{16} = \{\langle \textit{toss}(\textit{coin}1, \textit{Fred}), \mathrm{T}_1\rangle, \langle \textit{toss}(\textit{coin}2, \textit{Jane}), \mathrm{T}_2\rangle\}$,
$\textit{s-outcome}(AO_{16}, s_0) = s_3$.

We then generate the immediate successors of m_0 by creating a new moment for each action-outcome pattern AO, and assigning this new moment the s-outcome for AO. In diagramming these models, it is natural to display the tree as a graph, and to label each edge with the action-outcome pattern that produces it.

Applying this process to m_0, we obtain sixteen successor moments. Part of this first level of the construction will look like this.

4. Models

Here we summarize in more formal terms the account of models that has emerged from the previous discussion.

A model M consists of the following components. (The relatively large number of components is due to the treatment of actions, states, moments, and outcomes as primitives.)

1. A nonempty set $\mathcal{D}_M$. *Explanation:* the domain of individuals.

2. A set $\mathcal{M}_M$. *Explanation:* the moments of the model. (Moments aren't clock times; they are nodes in a branching time tree, to which states are assigned.)
3. A member m_0^M of $\mathcal{M}_M$. *Explanation:* the root moment for the temporal structure of the model.
4. A function $successors_M$ from $\mathcal{M}_M$ to subsets of $\mathcal{M}_M$. *Explanation.* This function returns the immediate successors of a moment.
5. A set $\mathcal{S}_M$. *Explanation:* the states of the model.
6. A function $state_M$ from $\mathcal{M}_M$ to $\mathcal{S}_M$. *Explanation:* the state assignment function of the model.
7. A set $Actions_M$. *Explanation:* the actions of the model. See Section 3 for explanation of what is meant intuitively by an action.
8. A nonempty subset $Agents_M$ of $\mathcal{D}_M$. *Explanation:* the agents of the model.
9. A nonempty set $Outcomes_M$. *Explanation:* the possible a-outcomes of the model.
10. A function $agent_M$ from $Actions_M$ to $Agents_M$. *Explanation:* This function returns the agent associated with each action.
11. A function $A\text{-}Outcomes_M$ from $Actions_M \times \mathcal{S}_M$ to the power set of $Outcomes_M$. *Explanation:* This function returns the set of outcomes associated with an action in a state.
12. A function $joint\text{-}acts_M$ from $\mathcal{S}_M$ to the power set of $Actions_M$. *Explanation:* This function returns the sets of actions that can be concurrently performed in a state.

Definition 1. (Action-outcome patterns.)
An *action-outcome pattern* for M is a set AO of pairs $\langle a, o\rangle$, where $a \in Actions_M$ and $o \in Outcomes_M$.
$Actions(AO) = \{a : \langle a, o\rangle \in AO\}$. $A\text{-}Outcomes(AO) = \{o : \langle a, o\rangle \in AO\}$.
$Actions_{M,p}(AO) = \{a \in Actions_M : agent_m(a) = p\}$.
A pattern AO is *feasible in* s, where $s \in \mathcal{S}_M$, if (1)
$outcome(\langle a, o\rangle) \in A\text{-}Outcomes_M(action(\langle a, o\rangle), s)$ whenever $\langle a, o\rangle \in AO$, and (2)
$Actions(AO) \in joint\text{-}acts_M(s)$.
$AO\text{-}Patterns_M(s)$ is the set of all feasible outcome patterns for M in state s.
$AO\text{-}Patterns_M = \bigcup\{AO\text{-}Patterns_M(s) : s \in \mathcal{S}_M\}$.

Explanation: An action-outcome pattern is an association of outcomes with a set of actions. The pattern is feasible in s if the outcomes are appropriate for the actions in s, and the set of actions is jointly feasible in s.

13. A function $s\text{-}outcome_M$ from states and action-outcome patterns appropriate for the state to states. I.e., if AO is an appropriate action-outcome pattern for s, then $s\text{-}outcome_M(s, AO) \in \mathcal{S}_M$.
    Explanation: This function returns the outcome state that results when action-outcome pattern occurs in an initial state.
14. A valuation V_M, providing appropriate values $V_M(X, s)$ for individual constants and predicates X, for each state s in $\mathcal{S}_M$.

Explanation: This is a familiar first-order interpretation, parameterized for states.

We impose two requirements on the models of the theory. The purpose of these conditions is to guarantee that the successor relation on instants is uniquely determined by the s-outcomes of jointly feasible actions, that histories in the model are constructed by applying the successor relation iteratively to the root moment m_0, and that the depth of branches is uniform.

Condition 1. For all $m \in \mathcal{M}_M$, either (1) $successors_M(m) = \emptyset$, or (2) there is a one-to-one function F^M_m from $AO\text{-}Patterns_M(s)$ to $successors_M(m)$, such that for all $m' \in successors_M(m)$, and $AO \in AO\text{-}Patterns_M$, $state_M(F^M_{m'}(AO)) = s\text{-}outcome_M(state_M(m), AO)$.

Definition 2. (The *result* function.)
$result_M(AO, m) = F^M_m(AO)$.

Definition 3. (Histories.)
A *history* on M is a maximal chain over the tree with root m^M_0 and successor function $successors_M$. $Histories_M$ is the set of histories of M.

Definition 4. (Depth of moments.)
$depth_M$ is a function from $\mathcal{M}_M$ to ω, such that: $depth_M(m_0) = 0$ and for all $m' \in successors_M(m)$ we have $depth_M(m') = depth_M(m) + 1$.

Definition 5. (Depth of histories.)
If h is a finite history of M, let $end\text{-}moment(h)$ be the last moment in h, and let $depth_M(h) = depth_M(end\text{-}moment(h))$. If h is an infinite history, let $depth_M(h) = \omega$.

Condition 2. The histories of a model M are of uniform depth. I.e., there is an $\nu_M \leq \omega$ such that for all histories h of the model, $depth(h) = \nu$. In that case, we say that $depth(M) = \nu_M$.

Definition 6. (Action-outcome sequences.)
Let $AO\text{-}Sequences^*_M$ consist of all sequences $\langle AO_i \rangle_{i<depth(M)}$, where $AO_i \in AO\text{-}Patterns_M$. Where $\alpha \in AO\text{-}Sequences^*_M$, let $state_M(\alpha, 0) = state_M(m^M_0)$ and $state_M(\alpha, i+1) = s\text{-}outcome_M(state_M(\alpha, i), \alpha_i)$, where $i + 1 < depth(M)$ and $s\text{-}outcome_M(state_M(\alpha, i), \alpha_i)$ is defined. Let $AO\text{-}Sequences_M$ be the set of all members α of $AO\text{-}Sequences^*_M$ such that $state(\alpha, i)$ is defined for all $i < depth(M)$, and $\alpha_i \in AO\text{-}Patterns_M(state_M(\alpha, i))$.

Conditions 1 and 2 ensure that any model can be generated from a root moment and its associated state, by iterating the process of generating successor moments to some fixed, uniform depth (which may be infinite). Moreover, each successor moment is uniquely determined by a feasible action-outcome pattern. They also ensure that the histories of a model are in one-to-one corresponence with $AO\text{-}Sequences_M$. We use this fact to establish an action-based notation for histories.

Definition 7. (Notation for histories.)
Where $\alpha \in \textit{AO-Sequences}_M$, let $\textit{history}_M(\alpha)$ be the unique history of M corresponding to α. And let $\textit{ao-sequence}_M(h)$ be the unique AO-sequence in M that corresponds to h, where $h \in \textit{Histories}_M$.

Definition 8. (Strategies.)
A strategy for $p \in \textit{Agents}_M$ is any sequence $\langle a_i \rangle_{i<depth(M)}$ of actions, where $\textit{agent}_M(a_i) = p$ for all $i < \textit{depth}(M)$. $\textit{Strategies}_M$ is the set of all strategies of M. $\textit{agent}(\sigma) = p$ iff σ is a strategy for p. Where σ is finite, with length n, $\textit{last}(\sigma) = \sigma_n$. $\textit{Strategies}_M(p)$ is the set of strategies M for agent p.

Definition 9. (Set of outcome histories for a strategy.)
Where $\sigma \in \textit{Strategies}_M$ and $\textit{agent}(\sigma) = p$, let $\textit{Histories}_M(\sigma) = \{\textit{history}_M\sigma : \sigma_i = a\}$, where $\langle a, o\rangle \in \alpha_i$ and $\textit{agent}(a) = p$.

Definition 10. (Restriction of M to n)
Where M is a model with $\textit{depth}(M) = \nu$ and $n < \nu$, the restriction of M to n is the model whose histories are obtained by truncating the histories of M to depth n.

Definition 11. (Notation for action-outcome sequences.)
Where $\alpha \in \textit{AO-Sequences}_M$, and $\textit{depth}_M = n$, let $\alpha ^\frown \langle AO\rangle$ be the sequence α' such that $\alpha'_i = \alpha_i$ for $i < n$ and $\alpha'_n = AO$.

To simplify our account of the conditional, we will impose the further condition that at each moment, each agent performs one and only one action. In models in which the action types include a null action, and are closed under conjunction, this assumption does not impose any real restrictions.

Condition 3. Let $AO \in \textit{AO-Patterns}_M$. For each $p \in \textit{Agents}_M$, there is a unique $AO(p) \in \textit{Actions}_M$ such that $\textit{agent}_M(a) = p$, where $AO(p) = \langle a, o\rangle$ for some o.

5. Conditionals and causality

The general semantic formalisms for conditionals[13] imposed abstract constraints on conditional selection functions, or on closeness relations among worlds. The difficulty these formalisms have had in showing how selection functions can be constructed in realistic cases has been a chronic source of philosophical criticism of these formalisms.[14] And in the AI literature, the task of specifying a selection has also proved to be highly intractable in domains with interacting joint constraints; see [Ginsburg 1985].

It is useful to look at conditional constructions as presenting a problem in nonmonotonic reasoning that is similar to the frame problem in temporal reasoning. By default, $[A \wedge C] \;\Box\!\!\rightarrow B$ holds if $A \;\Box\!\!\rightarrow B$ holds. Exceptions to this default are provided by specific "counterfactual causal" rules. The problem is how to provide these rules.[15]

[13]See [Stalnaker & Thomason 1970] and [Lewis 1973].

[14]These complaints have often been relatively unfocused. But the contrast between the relative lack of constraint on general selection functions and the complexity and detail of conditional reasoning has disturbed even the developers of these logics.

[15]Such an approach has been suggested in various places. See, for instance, [Horty & Thomason 1991] and [Asher & Morreau 1991].

In most current action formalisms, a causal predicate is used to relate actions to their direct effects. But there is no agreement on the extent to which a general theory of causality is needed in such formalisms. The need for causal information is much greater when conditionals are at stake.

To see this, notice that the $World_4$ of Section 3 we lack information about conditionals, because we lack causal information. Suppose, for instance, that in the first turn, Jane tosses her coin and it comes heads up; Fred tosses his coin and it comes up tails. In the second turn, Jane tosses her coin and it comes up heads; Fred turns his coin tails up. Now we ask: *would* Fred have turned his coin tails up in the second turn *even if* Jane's coin had come up heads in the first turn? We don't really know. If Fred and Jane are acting without reference to each other (they are in different places, with no communication), Fred would have turned his coin tails up even if Jane's coin had come up differently. If they are able to observe each other, and Fred is imitating the results of Jane's previous toss, then Fred would have turned his coin heads up if Jane's coin had come up heads. If Fred is influenced in his planning by Jane's previous toss, then Fred may have turned his coin heads up if Jane's coin had come up heads.

For this reason, we will add causal information to our causal model. This information takes the form of a set of relations having any of the following three forms:

$$Action_1 \quad\quad Action_2 \qquad\qquad a\text{-}outcome \quad\quad Action \qquad\qquad a\text{-}outcome \quad\quad _o\, Action$$

we could think of the causal model then, as a directed graph whose nodes are either actions or outcomes, and with three sorts of edges.

The relation $Action_1$ $Action_2$ means that whether $Action_2$ is performed can depend on $Action_1$'s having just been performed. The relation *a-outcome* *Action* means that whether *Action* is performed can depend on *aoutcome*'s having just come about. The relation *a-outcome* $_o$ *Action* means that the a-outcomes of *Action* can depend on *aoutcome*'s having just come about.

This representation of causal information incorporates several simplifying assumptions. For instance, it involves a qualitative form of the Markov condition: there can be no delayed influences. But we believe it provides a basis for formulating many decision problems.

We can illustrate the role of causal relations by returning to the coin world. Suppose that we think *Fred*'s action depends in some complex, perhaps poorly understood way, on *Jane's* previous action. Then the model should contain the following causal relations. (Call the world with these causal relations $World_{4,1}$.)

$$\begin{aligned}&\mathbf{a}(coin1, Jane) \quad\quad \mathbf{b}(coin2, Fred),\\ &\quad \text{where } \mathbf{a} \in \{TurnUp(coin2, Jane), TurnDown(coin2, Jane), toss(coin2, Jane)\}\\ &\quad \text{and } \mathbf{b} \in \{toss(coin1, Fred), TurnUp(coin1, Fred), TurnDown(coin1, Fred)\}.\end{aligned}$$

We want $TurnUp(coin2, Jane)$ $TurnUp(coin1, Fred)$, for instance, among the causal relations to indicate how certain conditionals will be affected by the hypothesis that *Fred* takes *Jane*'s actions into account in deciding what to do. For example, suppose that *Jane* turned up *coin*2 and then *Fred* turned up *coini*. The dependence between *Jane*'s actions and *Fred*'s subsequent actions means that if *Jane* had flipped *coin*2, then *Fred* might not then have turned *coin*1 down. (Contrast this with the case in which the actions are independent; they take place in different places, with no transfer of information. In this case, *Fred* would still have turned *coin*1 down, even if *Jane* had flipped *coin*2.)

Suppose now that *Jane* takes the position of *coin*1 into account in deciding what action to take. (In fact, she has some formula for deciding what action to take, though we may not know what this formula is.) Then we want the following additional causal relations.

$$\begin{array}{llll} \mathrm{H}_1 & \mathit{TurnUp}(\mathit{coin}2, \mathit{Jane}) & \mathrm{H}_1 & \mathit{TurnDown}(\mathit{coin}2, \mathit{Jane}) \\ \mathrm{H}_1 & \mathit{toss}(\mathit{coin}2, \mathit{Jane}) & \mathrm{T}_1 & \mathit{TurnUp}(\mathit{coin}2, \mathit{Jane}) \\ \mathrm{T}_1 & \mathit{TurnDown}(\mathit{coin}2, \mathit{Jane}) & \mathrm{T}_1 & \mathit{toss}(\mathit{coin}2, \mathit{Jane}) \end{array}$$

$\mathrm{H}_1 \quad \mathit{TurnUp}(\mathit{coin}2, \mathit{Jane})$, for instance, is among the causal relations because if *Jane* turned *coin*2 up after *coin*1 came up heads, then she might have done something else if *coin*1 had come up tails.

Finally, suppose that when coins are tossed, they are put on a coin flipping device in their current postion (e.g., they are put on the device heads up if they landed heads up on the previous toss). Then the outcome of a toss of a coin depends on its initial position; this is represented by the following causal relations.

$$\begin{array}{llll} \mathrm{H}_1 & {}_o\, \mathit{toss}(\mathit{coin}1, \mathit{Fred}) & \mathrm{T}_1 & {}_o\, \mathit{toss}(\mathit{coin}1, \mathit{Fred}) \\ \mathrm{H}_2 & {}_o\, \mathit{toss}(\mathit{coin}2, \mathit{Fred}) & \mathrm{T}_2 & {}_o\, \mathit{toss}(\mathit{coin}2, \mathit{Fred}) \end{array}$$

In this case, $\mathrm{H}_1 \quad {}_o\, \mathit{toss}(\mathit{coin}1, \mathit{Fred})$ is among the causal relations because under our causal hypothesis the fact that *coin*1 came up heads can influence the result of the next toss. (The sort of counterfactual influence that we have in mind here is perfectly compatible with the hypothesis that *coin*1 is statistically fair.)

Definition 12. ($\mathit{Causal\text{-}Rels}_M$.)
$\mathit{Causal\text{-}Rels}_M$ is the set of causal relations of the model M.

6. Causal independence and conditional selection

We now show how a conditional selection function can be defined on action models. Like the conditional function of [Thomason & Gupta 1980], this function respects the structure of branching time; details of the two constructions differ.

The following definitions are relativized to a fixed model M, which we assume meets the conditions of Section 4. Also, we will not try to define a conditional for arbitrary antecedents. Instead, we will confine ourselves to antecedents concerning alternative actions that might have been performed; such antecedents are adequate for the decision theoretic applications that we envisage.

We begin by defining the set $\mathit{closest}_M(p, h^*, \sigma)$ for models M of depth 1, where $p \in \mathit{Agents}_M$, $h^* \in \mathit{Histories}_M$, and $\sigma \in \mathit{Strategies}_M(p)$.

Definition 13. ($AO_1 \sim_{M,p} AO_2$.)
Where $AO_1, AO_2 \in \mathit{AO\text{-}Patterns}$ and $p \in \mathit{Agents}$, let $AO_1 \sim_{M,p} AO_2$ if and only if for all $a \in \mathit{Actions}_M$, if $\mathit{agent}_M(a) \neq p$ then $\langle a, o\rangle \in AO_1$ iff $\langle a, o\rangle \in AO_2$

Explanation: Two AO-patterns are similar with respect to p if they differ (if at all) only with respect to p's action and its outcomes.

Definition 14. ($\mathit{closest}_M(p, h^*, \sigma)$: Basis Case.)

Given a model M with $depth(M) = 1$, as characterized in Section 4, let $s_0 = state(m_0)$, $\sigma = \langle a \rangle$, where $agent(a) = p$, $AO^* \in AO\text{-}Patterns_M(s_0)$, $h^* = history(\langle AO^* \rangle)$, $p \in Agents$, and $\langle a \rangle \in Strategies_M(p)$. Then:

(1) If $a \in Actions(AO^*)$ then $closest_M(p, h^*, \langle a \rangle) = \{h^*\}$; and

(2) Otherwise,
$closest_M(p, h^*, \langle a \rangle) = \{history_M(\langle AO \rangle) : a \in Actions(AO) \text{ and } AO \sim_p AO^*\}$.

Explanation: The closest histories to h^* with respect to strategy α for agent p are the histories that involve minimal changes to the actions that led to h^*: the actions of agents other than p are unchanged, and the action recommended by α is substituted for p's action in h^*. All outcomes of p's alternative actions are allowed. Clause (1) ensures that in case α actually yields h^* then h^* is the unique closest history.[16]

We now show how to define $closest_M(p, h^*, \sigma)$ for models M of depth $n + 1$.

Definition 15. ($AO_1 \mathit{I\text{-}Precedes}_{M,s,p,AO,AO',AO''} AO_2$.)
$AO_1 \mathit{I\text{-}Precedes}_{M,s,p,AO,AO',AO''} AO_2$ if and only if for some $q \in Agents_M$, $q \neq p$, we have $AO_1 \sim_p AO_2$ and, where $\langle a, o \rangle \in AO$ and $agent(a) = q$, either:

(1) (i) There is no relation $a' \quad a$ in $Causal\text{-}Rels_M$, where $a' \in Actions(AO')$ and $a' \notin Actions(AO'')$, and (ii) there is no relation $o' \quad a$ in $Causal\text{-}Rels_M$, where $o' \in A\text{-}Outcomes(AO')$ and $o' \notin A\text{-}Outcomes(AO'')$, and (iii) $a \in Actions(AO_1)$ but $a \notin Actions(AO_2)$, or

(2) (i) There is no relation $a' \quad a$ in $Causal\text{-}Rels_M$, where $a' \in Actions(AO')$ and $a' \notin Actions(AO'')$, and (ii) there is no relation $o' \quad a$ in $Causal\text{-}Rels_M$, where $o' \in A\text{-}Outcomes(AO')$ and $o' \notin A\text{-}Outcomes(AO'')$, and (iii) there is no relation $o' \quad_o a$ in $Causal\text{-}Rels_M$, where $o' \in A\text{-}Outcomes(AO')$ and $o' \notin A\text{-}Outcomes(AO'')$, and (iv) $a \in Actions(AO_1)$ and $a \in Actions(AO_2)$, but (v) $\langle a, o \rangle \in AO_1$ and $\langle a, o \rangle \notin AO_2$.

Definition 16. ($AO_1 <_{M,s,p,AO,AO',AO''} AO_2$.)
$<_{M,s,p,AO,AO',AO}$ is the transitive closure of $\mathit{I\text{-}Precedes}_{M,s,p,AO,AO',AO''}$.

Explanation: Imagine that AO-pattern AO'' produces a state s, where both AO_1 and AO_2 are feasible, and that AO' produces a moment at which AO is then applied. Let p be an agent. Then AO_1 is closer to AO than AO_2, with respect to s, p, AO, AO', and AO'' if AO_1 can be produced from AO by fewer changes of action-outcome pairs, where these changes respect respect the applicable causal relations for the preceding actions of outcomes.

Definition 17. ($closest_M(p, h^*, \sigma)$: Inductive Case.)

[16]This clause, and a similar clause in the inductive step, will ensure a general property of conditionals known as "centering"; see [Lewis 1973].

Given a model M_{n+1} with $depth(M_{n+1}) = n+1$, where $n > 0$, let M_n be M_{n+1} restricted to n. Let $h^*_{n+1} = history_{M_{n+1}}(\alpha^*_n)$, where $\alpha^*_n = \alpha^*_{n-2}{}^\frown\langle AO^*_{n-1}\rangle{}^\frown\langle AO^*_n\rangle$, and let $h^*_n = history(\alpha^*_{n-1})$. Let $p \in Agents$ and $\sigma_n = \sigma_{n-1}{}^\frown\langle\alpha\rangle \in Strategies_M(p)$. Let $strategy(\alpha^*_n)$ be the strategy τ for p such that $\tau_i = c$, where $Actions_p(\alpha^*_i) = \{c\}$ whenever α^*_i is defined, and such that τ_i is undefined whenever α^*_i is undefined.
Then either (1) $strategy_p(\alpha^*_n) = \sigma_n$ and $closest_{M_{n+1}}(p, h^*_{n+1}, \sigma_n) = \{h^*_{n+1}\}$, or
(2) $strategy_p(\alpha^*_n) \neq \sigma_n$ and where $\sigma_n = \sigma_{n-1}{}^\frown\langle a\rangle$, we have:
$closest_{M_{n+1}}(p, h^*_{n+1}, \sigma_n) = \{history_{M_{n+1}}(ao\text{-}sequence_{M_{n+1}}(h_{n-1})^\frown\langle AO_{n-1}\rangle\langle AO_n\rangle) :$
$h_n = history_{M_n}(h_{n-1}{}^\frown\langle AO_{n-1}\rangle) \in closest_{M_n}(p, h^*_n, \sigma_{n-1})$ and, where
$\sigma = state(end\text{-}moment(h_n))$, we have (2.1) $AO_n \in AO\text{-}Patterns_M(s)$ and (2.2)
there is no $AO \in AO\text{-}Patterns_M(s)$ such that $AO <_{M,s,p,AO^*,AO^*_{n-1},AO_{n-1}} AO_n\}$.

Explanation: The closest histories to h^* with respect to strategy α for agent p are obtained by recursively finding closest histories to the subhistories of h^*, at each stage making changes that are forced by α and retaining any actions and outcomes that are not influenced by the causal conditions.

The utility of the formalism that has been presented so far consists in its applicability to a variety of examples of conditional reasoning about action. The theory is to be tested according to how successful it is in formalizing and explaining a variety of examples in this domain.

Also, an important part of understanding the theory is seeing how it applies to examples. The intricacy of the definitions that we have just rehearsed was induced by a process of reviewing examples, and modifying versions of the theory when they were unable to deal with these. Without a systematic review of these case studies, it may be hard to see why the theory was constructed as it was, although anyone who is familiar with the counterfactual reasoning in any of its guises is likely to realize that the problem is intrinsically complex.

For these reasons, a reasonably self-contained presentation of the theory should contain a long section dealing with examples. But space limitations make it impossible to do that in this version. Here, we will only sketch a few cases.[17]

Examples illustrating the theory can be constructed by seeing how conditionals would be treated in the microworlds that were presented in Section 3. (Of course, these models need to be supplemented with appropriate sets of causal information.) The following cases illustrate how this can be done.

We will confine ourselves to the $World_4$ model of depth 2. (See Section 3.) We will need the following action-outcome patterns, histories, and strategies.

$AO^*_1 = \{\langle TurnUp(coin1, Fred), \mathrm{H}_1\rangle, \langle toss(coin2, Jane), \mathrm{H}_2\rangle\}$.

$AO^*_2 = \{\langle TurnUp(coin1, Fred), \mathrm{H}_1\rangle, \langle toss(coin2, Jane), \mathrm{H}_2\rangle\}$.

$AO'_2 = \{\langle TurnDown(coin1, Fred), \mathrm{T}_1\rangle, \langle toss(coin2, Jane), \mathrm{H}_2\rangle\}$.

$s'_1 = result(AO_1, m_0)$.

$AO''_2 = \{\langle TurnDown(coin1, Fred), \mathrm{T}_1\rangle, \langle \mathbf{a}, \mathbf{o}\rangle\{ :$
where $\mathbf{a} \in \{TurnUp(coin2, Jane), TurnDown(coin2, Jane), toss(coin2, Jane)\}$
and $\mathbf{o} \in A\text{-}Outcomes(s'_1, \mathbf{a})\}$.

[17]We hope to develop and maintain a version of the paper that will contain appropriate case studies. Consult http://www.pitt.edu/ thomason/thomason.html for progress on this project.

$h^* = history AO_1, AO_2$.

$\sigma_1 = \langle TurnUp(coin1, Fred), TurnDown(coin1, Fred)\rangle$.

$\sigma_2 = \langle TurnDown(coin1, Fred), TurnDown(coin1, Fred)\rangle$.

Example 1. Whatever causal relations are in the model,

$$closest(Fred, h^*, \sigma_1) = \{history(AO_1^*, AO_2')\}.$$

That is, if *Fred* had turned his coin down on the second turn, *Jane* would still have tossed her coin, and it would still have come up heads.

Example 2. Assume that there are no causal relations in the model. In this case,

$$closest(Fred, h^*, \sigma_2) = \{history(AO_1^*, AO_2')\}.$$

That is, if *Fred* had turned his coin down on both turns, *Jane* would still have tossed her coin in both turns, and it would still have come up heads in both turns.

Example 3. Assume the causal relations of $World_{4,1}$; here, *Jane*'s actions can depend on *Fred*'s. Here,

$$closest(Fred, h^*, \sigma_1) = \{history(AO_1^*, AO_2')\}.$$

As before, if *Fred* had turned his coin down on the second turn, *Jane* would still have tossed her coin on that turn, and it would still have come up heads. But

$$closest(Fred, h^*, \sigma_2) = \{history(AO_1^*, AO_2')\}.$$

That is, if *Fred* had turned his coin down on both turns, *Jane* might have done anything on the second turn. (Opinions may differ on whether the outcomes of $toss(coin2, Jane)$ should be resricted in this case. Our definition of closeness is conservative, in allowing all possibilities.)

7. Utilities of histories and dominance

7.1. Two ways to introduce utilities

Classical decision theory assigns a numerical value, or utility, to each fully specified result of a course of action. In our branching-time models, this corresponds to assigning a numerical utility to each history. If we wish to complete the decision-theoretic picture, we can also assign a probability to each history. A qualitative assignment, as we will treat it, is much more modest—there are no probabilities at all, and there is only a "no worse" relation $\succeq$ that compares the relative utility of any two histories. The most natural extension of our concurrent action models to quantitative utilities turns out to yield an account of expected utility that resembles that of [Gibbard & Harper 1978].

We will define a relation of dominance between strategies on each of these approaches, and will state and prove a theorem showing that qualitiative dominance is sound with respect to quantitative dominance.

7.1.1. Quantitative dominance

We obtain a quantitative utility model by adding the following two items to the model components described in Section 4.

1. A function ut from the set $Histories_M$ of M to the reals.
2. A probability measure pr over $Histories_M$.

Gibbard and Harper begin ([Gibbard & Harper 1978, p. 153]) with the idea that the value of an act is the sum over the set of outcomes of the act of the product of the utility of the outcome with the probability of the conditional stating that this outcome would occur if the act were performed. This conditional is interpreted (as in Stalnaker's general logic of conditionals) by a function that takes certain antecedent conditions (conditions of the form 'I do action a') and outcomes into outcomes.

In opting for a function that returns a single outcome, rather than a set of closest outcomes, Gibbard and Harper seem to take a stand on the controversial issue of "conditional excluded middle" in conditional logic. With respect to this issue, they say[18] in effect that they find the single-valued assumption implausible, but assume it for the sake of simplicity in the initial statement of causal decision theory.

We will work with this single-valued version of causal decision theory.[19] We also assume a version of causal decision theory in which the outcomes are histories, and the conditionals have the form

$$\mathrm{Do}(\sigma, h_1, p) \;\Box\!\!\rightarrow h_2.$$

Intuitively, this means that assuming that h_1 will occur, if agent p were to adopt plan σ then h_2 would ensue. So the conditional selection quantitative decision model M will take a plan σ, an agent p, and a history h_1 into a history $condit_M(\sigma, p, h_1)$.

The formula for the quantitative causal utility of σ then sums over histories h the product of the utility of h with the probability of the conditional that says if σ were adopted then h would ensue.

Definition 18. ($ut_{GH}(\sigma, p)$.)
$ut_{GH}(\sigma, p) = \sum_{h \in Histories_M} ut(h) pr(\{h' \,/\, condit(\sigma, p, h) = h'\})$, where $p = agent(\sigma)$.

Quantitative dominance is obtained by comparing utilities.

Definition 19. (GH-Dominance.)
$\sigma \geq \tau$ iff $ut_{GH}(\sigma, Histories) \geq ut_{GH}(\tau, Histories)$.

This is a weak, or "no-worse" notion of dominance. Strong dominance can be defined as follows.

Definition 20. (Strong GH-Dominance.)
$\sigma > \tau$ iff $ut_{GH}(\sigma, Histories) > ut_{GH}(\tau, Histories)$.

[18] In the discussion of their Axiom 2, [Gibbard & Harper 1978, p. 156]

[19] In fact, we do not know if there has been a generalization of the theory to conditionals that do not obey conditional excluded middle. We suspect that such a generalization might not only have to be more complex, but *ad hoc* in certain ways.

7.1.2. Qualitative dominance

A qualitative model, as well as the components listed in Section 4, has the following additional ingredient.

1. A reflexive, transitive ordering $\succeq$.

The relation $h \succeq h'$ means that h is no worse than h'. We do not exclude the possibility that the relation $\succeq$ is inferred through a process of reasoning, as in [Boutilier 1994b] and [Tan & Pearl 1994]. But we do not discuss such reasoning processes here. We let $h \succ h'$ iff $h \succeq h'$ and $h' \not\succeq h$.

Intuitively, a strategy σ qualitatively dominates another strategy τ if given any history h, any history h'' that would ensue if σ were adopted provided that otherwise h occur is no worse than any history that would ensue if τ were adopted. This leads to the following definition.

Definition 21. (Qualitative Dominance for Plans.)
$\sigma \succeq_p \tau$ iff for all $h \in \mathit{Histories}$, all $h_1 \in \mathit{closest}(p, h, \sigma)$, and all $h_2 \in \mathit{closest}(p, h, \tau)$, we have $h_1 \geq h_2$.

Again, we can define a notion of strong dominance.

Definition 22. (Strong Qualitative Dominance for Plans.)
$\sigma \succ_p \tau$ iff $\sigma \succeq_p \tau$ and for some $h \in \mathit{Histories}(\sigma)$, we have $h_1 \succ h_2$ for some $h_1 \in \mathit{closest}(p, h, \sigma)$ and $h_2 \in \mathit{closest}(p, h, \tau)$.

7.2. Soundness

Three things need to be added to a qualitative model to extend it to a quantitative one: a probability measure, an assignment Ut of numerical utilities, and a choice function that selects a member of $\mathit{closest}(p, h, \sigma)$. (The third ingredient is needed to secure a single-valued conditional selection function.) We will say that the extension is *consistent* if Ut is consistent with the qualitative utility ordering given in M; i.e., if $Ut(h) > Ut(h')$ if $h >_M h'$.

Soundness amounts to this: when a qualitative model is consistently extended to a quantitative one in this way, its dominance relation is consistent with the corresponding qualitative dominance relation. More precisely:

Theorem 1. Let M be a qualitative utility model. Let pr be a probability measure over $\mathit{Histories}_M$, and let $s(\sigma, p, h) \in \mathit{closest}(p, h, \sigma)$. Let $M(pr, s, Ut)$ be the quantitative model that is obtained by adding pr to the ingredients in M, letting $\mathit{condit}_{M(pr, s, Ut)} = s$, and by letting $Ut_{M(pr, s, Ut)}$ be consistent with Ut. Then: if $\sigma \succeq_M \tau$ then $\sigma \geq_{M(pr, s, Ut)} \tau$. Similarly, if $\sigma \succ_M \tau$ then $\sigma >_{M(pr, s, Ut)} \tau$.

Proof. Suppose that $\sigma \succeq \tau$. This means that for all $h, h_1, h_2 \in \mathit{Histories}$, If $h_1 \in \mathit{closest}(p, h, \sigma)$ and $h_2 \in \mathit{closest}(p, h, \tau)$ then $h_1 >_M h_2$.

Now,

$ut_{GH}(\sigma, p) = \sum_{h \in \mathit{Histories}_M} ut(h) pr(\{h' \,/\, \mathit{condit}(\sigma, p, h) = h'\})$ (by definition)
$= \sum_{h \in \mathit{Histories}_M} \sum_{h' \in \{h'' \,/\, \mathit{condit}(\sigma, p, h'') = h\}} pr(\{h'\}) ut(h)$ (pr is a probability measure)
$= \sum_{h \in \mathit{Histories}_M} pr(\{h\}) ut(\mathit{condit}(\sigma, p, h))$ (condit is a function)
$> \sum_{h \in \mathit{Histories}_M} pr(\{h\}) ut(\mathit{condit}(\tau, p, h))$ (from qualitative dominance)
$= \sum_{h \in \mathit{Histories}_M} \sum_{h' \in \{h'' \,/\, \mathit{condit}(\tau, p, h'') = h\}} pr(\{h'\}) ut(h)$ (condit is a function)

$= \sum_{h \in Histories_M} ut(h) pr(\{h' \ / \ condit(\tau, p, h) = h'\})$ (pr is a probability measure)
$= ut_{GH}(\tau, p)$ (by definition)

Gibbard and Harper show[20] that under certain conditions, their account of expected utility gives the same results as that of classical decision theory. Their two conditions presuppose a language with an explicit conditional, and so we cannot verify them directly. However, the crucial condition (their "Condition 1") would be validated if we adopted an object language with an explicit conditional, and the other condition is one that one can be expected to hold widely in decision-making examples.

Thus, our soundness result shows (1) that whenever a dominance relation holds in the qualitative theory it will correspond to the recommendations of a generalization of one quantitative account, and (2) that in many cases it will correspond to the recommendations of classical decision theory.

7.3. Examples

If there is space, we need to put in some examples.

8. Some lines of development

Much work remains in developing the technical part of the theory and relating it to the formalization of common sense practical reasoning. Here are some considerations. (1) In the current version, we have not considered what would happen if an explicit conditional were added to the language. This would facilitate a comparison to the general conditional formalisms, as well to Gibbard and Harper's causal decision theory. (2) The definitions of this paper were simplified in a number of ways; for instance, the only allowable conditional antecedents were simple strategies. Thus, they need to be generalized in a number of ways. (3) Unfortunately, the simplifications did not prevent the definitions from becoming unpleasantly complex. It would be good to have a more modular and readily intelligible presentation of the theory, along with some diagramming conventions. (4) The properties of interesting special cases of the theory need to be developed; for instance, the *deterministic case* has many interesting features. (5) It is clear that a causal theory of some sort is required by our conditional, but it is less clear what form this theory should take. The version that we have provided here is rather crude, and certainly does not permit the expression of causal dependencies that are at all complex. The utility of the causal theory needs to be tested by formalizing some representing domains. (6) More generally, the entire theory needs to be tested in this way. We have in mind not only domains borrowed from decision theory, but cases from "cognitive robotics," such as those described in [Reiter 1996]; and [Lansky 1987] also provides complex, realistic examples. (7) The branching time formalism here needs to be compared with the matrix formalisms of decision theory, and with the "STIT" formalisms deriving from [Nuel D. Belnap & Perloff 1988]. A detailed comparison with the "utilities lifting" literature mentioned in Section 2 would also be useful. (8) The current theory does not provide even for qualitative differences in the likelihood of branches. As we intimated in Section 2, we have in mind an account of these differences using ideas from nonmonotonic logic.[21] (9) We know (for instance, from Wellman's work) that dominance

[20]See [Gibbard & Harper 1978, p. 157].

[21]In particular, we have in mind a normalcy relation over histories.

reasoning is implementable. But the extent to which the theory that we have presented here can guide implementations is unclear; this question raises an entire research program, which we have not even begun to think about. (10) Finally, the addition of an epistemic dimension is obviously desirable, and would help us to make contact with many of the most exciting issues in contemporary decision theory.

Bibliography

Allen, J. F. 1990. Two views of intention: Comments on Bratman and on Cohen and Levesque. In Cohen, P. R.; Morgan, J.; and Pollack, M., eds., *Intentions in Communication*. Cambridge, Massachusetts: The MIT Press. 71–75.

Asher, N., and Morreau, M. 1991. Commonsense entailment: a modal theory of non-monotonic reasoning. In Mylopoulos, J., and Reiter, R., eds., *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, 387–392. Los Altos, California: Morgan Kaufmann.

Boutilier, C. 1994a. Defeasible preferences and goal derivation. Unpublished manuscript; Department of Computer Science, University of British Columbia, Vancouver, British Columbia, Canada V6T 1Z4.

Boutilier, C. 1994b. Toward a logic for qualitative decision theory. In Doyle, J.; Sandewall, E.; and Torasso, P., eds., *Principles of Knowledge Representation and Reasoning*, 75–86.

Dowty, D. 1979. *Word Meaning in Montague Grammar*. Dordrecht, Holland: D. Reidel.

Doyle, J., and Wellman, M. 1991. Preferential semantics for goals. In Dean, T., and McKeown, K., eds., *Proceedings of the Ninth National Conference on Artificial Intelligence*, 698–703. Menlo Park, California: American Association for Artificial Intelligence.

Doyle, J., and Wellman, M. 1994. Representing preferences as *ceteris paribus* comparatives. In *Working Notes of the AAAI Spring Symposium on Decision-Theoretic Planning*, 69–75. Menlo Park, California: American Association for Artificial Intelligence.

Doyle, J.; Shoham, Y.; and Wellman, M. P. 1991. A logic of relative desire (preliminary report). In Ras, Z., ed., *Proceedings of the Sixth International Symposium on Methodologies for Intelligent Systems*, Lecture Notes in Computer Science. Berlin: Springer-Verlag.

Fikes, R., and Nilsson, N. 1971. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence* 2:189–208.

Friedman, M., and Savage, L. 1948. The utility analysis of choices involving risk. *Journal of Political Economy* 56:279–304.

Georgeff, M. 1987. Many agents are better than one. In Brown, F. M., ed., *The Frame Problem in Artificial Intelligence: Proceedings of the 1987 Workshop*, 59–75. Los Altos, California: Morgan Kaufmann.

Gibbard, A., and Harper, W. 1978. Counterfactuals and two kinds of expected utility. In Hooker, C.; Leach, J.; and McClennen, E., eds., *Foundations and Applications of Decision Theory*. Dordrecht: D. Reidel. 125–162.

Ginsburg, M. 1985. Counterfactuals. In Joshi, A., ed., *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 80–86. Los Altos, California: Morgan Kaufmann. Attempts to define a selection function for a circuits domain.

Goldszmidt, M., and Pearl, J. forthcoming. Qualitative probabilities for default reasoning, belief revision, and causal modeling. *Artificial Intelligence*.

Grosz, B., and Kraus, S. 1993. Collaborative plans for group activities. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 367–373. Los Altos, California: Morgan Kaufmann.

Horty, J. F., and Thomason, R. 1991. Conditionals and artificial intelligence. *Fundamenta Informaticae* 15:301–324.

Horty, J. F. 1994. Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic* 23(1):35–65.

Jeffrey, R. 1983. *The Logic of Decision*. Chicago: University of Chicago Press, 2 edition.

Jennings, R. E. 1974. Utilitarian semantics for deontic logic. *Journal of Philosophical Logic* 3(4):445–456.

Jennings, R. E. 1985. Can there be a natural deontic logic. *Synthèse* 65:257–274.

Lansky, A. 1987. A representation of parallel activity based on events, structure, and causality. In *Reasoning About Actions and Plans, Proceedings of the 1986 Workshop at Timberline, Oregon*. San Mateo, California: Morgan Kaufmann. 123–160.

Lewis, D. K. 1973. *Counterfactuals*. Cambridge, Massachusetts: Harvard University Press.

Lifschitz, V. 1990. On the semantics of STRIPS. In Allen, J. F.; Hendler, J.; and Tate, A., eds., *Readings in Planning*. San Mateo, California: Morgan Kaufmann. 523–530.

Lin, F., and Shoham, Y. 1992. A logic of knowledge and justified assumptions. *Artificial Intelligence* 57:271–289.

Lin, F. 1996. Abstract operators, indeterminate effects, and the magic predicate. In Buvač, S., and Costello, T., eds., *Working Papers: Common Sense '96*, 96–103. Stanford University: Computer Science Department, Stanford University. Consult http://www-formal.Stanford.edu/tjc/96FCS.

Luce, R. D., and Raiffa, H. 1957. *Games and Decisions*. New York: John Wiley and Sons.

McCarthy, J. 1995. Situation calculus with concurrent events and narrative. Available by anonymous ftp at sail.stanford.edu. This paper is labled "Non Citable Draft." It should not be quoted.

McClennen, E. 1990. *Rationality and Dynamic Choice*. Cambridge, England: Cambridge University Press.

Nuel D. Belnap, J., and Perloff, M. 1988. Seeing to it that: a canonical form for agentives. *Theoria* 54:175–199.

Reiter, R. 1996. Natural actions, concurrency and continuous time in the situation calculus. In Aiello, L. C.; Doyle, J.; and Shapiro, S., eds., *KR'96: Principles of Knowledge Representation and Reasoning*. San Francisco, CA: Morgan Kaufmann. 2–13.

Schubert, L. 1990. Monotonic solution of the frame problem in the situation calculus; an efficient method for worlds with fully specified actions. In Kyburg, H.; Loui, R.; and Carlson, G., eds., *Knowledge Representation and Defeasible Reasoning*. Dordrecht: Kluwer. 23–67.

Shanahan, M. 1995. A circumscriptive calculus of events. *Artificial Intelligence* 77(2):251–284.

Stalnaker, R., and Thomason, R. 1970. A semantic analysis of conditional logic. *Theoria* 36:23–42.

Tan, S., and Pearl, J. 1994. Qualitative decision theory. In Hayes-Roth, B., and Korf, R., eds., *Proceedings of the Twelfth National Conference on Artificial Intelligence*, 928–933. Menlo Park, California: American Association for Artificial Intelligence.

Thomason, R., and Gupta, A. 1980. A theory of conditionals in the context of branching time. *The Philosophical Review* 80:65–90. Reprinted in W. Harper, R. Stalnaker, and G. Pearce, eds., *Ifs,* D. Reidel, Dordrecht, 1981, pp. 299-322.

van Fraassen, B. C. 1972. The logic of conditional obligation. *Journal of Philosophical Logic* 1:417–438.

Wellman, M. P. 1988. Formulation of tradeoffs in planning under uncertainty. Technical Report TR–427, Massachusetts Institute of Technology Laboratory for Computer Science, 545 Technology Square, Cambridge Massachusetts, 02139.

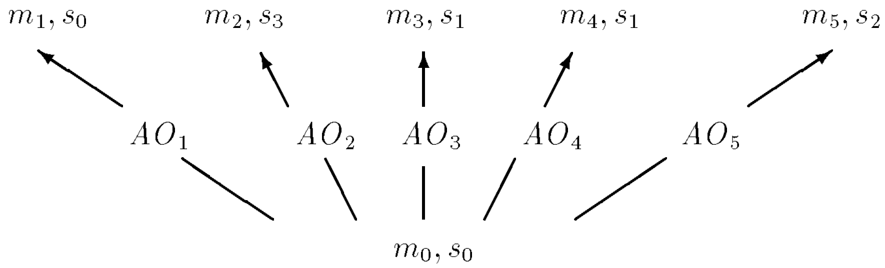

Figure 1: Part of the $World_4$ Construction