DYNAMIC LOGIC WITH LIKELIHOODS

JESSE HUGHES

Abstract. We interpret formulas of dynamic logic in terms of means-end ascriptions. We add a notion of uncertainty by using discrete distributions in the transition structure. We interpret this uncertainty as a function taking pairs of worlds w and formulas φ to $[0,1]$, yielding the *likelihood* that φ is true in w. The definition of the corresponding function is the chief interest here. We offer two different likelihood functions yielding different semantics for our modal logic. Unsurprisingly, neither of these semantics are well-behaved with respect to propositional logic, but the more easily motivated *simple* semantics appears to be less well-behaved than the more complicated and less intuitive *distributive* semantics.

1. Introduction

We are interested in representing means-end ascriptions using a dynamic logic. We regard statements of the form $[m]\varphi$ as asserting that m is a *means* to the *end* φ and that, for our present purposes[1], a means-end ascription expresses nothing more than $[m]\varphi$. Thus, we are committed to interpreting means as actions and ends as formulas. This commitment ought to be examined and defended, but we will not do so here.

We work with propositional logic, for simplicity's sake. Again, a full treatment of means-end ascriptions would probably require quantification, but we postpone that for later consideration. We fix a set **prop** of atomic propositions and a set **act** of actions. For each action m, we introduce an operator $[m]$ and interpret $[m]\varphi$ as asserting that, after doing m, the property φ attains (or likely attains). Thus our language **ME** is the least set containing **prop** and the true formula $\top$ and closed under the construction of negations, conjunctions and application of dynamic operators. We take $\bot$, $\vee$ and $\rightarrow$ as defined.

Our semantics will involve a set $\mathcal{W}$ of alternative worlds. We will define a relation $\models \subseteq \mathcal{W} \times \mathbf{ME}$, written in infix form. If $w \models [m]\varphi$, then we say in w, the act m is a means to φ.

Typically, a particular action may have several possible outcomes. Often one models this situation is by interpreting actions m by maps $[\![m]\!] : \mathcal{W} \rightarrow \mathcal{P}\mathcal{W}$, where $\mathcal{P}$ denotes the power set functor. The set $[\![m]\!](w)$ denotes the set of alternative outcomes (or the set of "normal" alternative outcomes) to doing m in w. Then we say $w \models [m]\varphi$ iff for every $w' \in [\![m]\!](w)$, we have $w' \models \varphi$ (where $w \models \varphi \wedge \psi$, $w \models \neg\psi$ and $w \models P$ are defined in the usual way).

Key words and phrases. dynamic logic, means-end reasoning, probabilistic transition systems.

[1]The account described here is clearly inadequate for a full analysis of means-end ascriptions. At the least, one wants a conditional operator in order to express the pre-conditions which may be explicit or implicit in a means-end ascription. A fully developed logic for means-end ascriptions is a topic for later research.

There are a number of deficiencies in adopting this interpretation to our means-end analysis. Our primary concern is that this semantics fails to include a notion of *efficacy*. By "efficacy", we mean the likelihood that a means will actually yield its end. This is an important characteristic of means and a failure to accommodate this feature is a real deficiency in our model. Efficacy is one of the primary methods of comparing different means to the same end and a common motivation for choosing one means over another. We aim to represent the notion of efficacy by adding probabilistic features to our transition structure. There are many other features of means-end reasoning which we will leave unaddressed for now so that we may fully explore dynamic logic in the presence of likelihoods.

It is perhaps not surprising that adding probabilities to a propositional setting yields some undesirable results. In fact, we offer two different semantics here, each of which has its own advantages and deficiencies, and neither of which is wholly satisfying. Neither of them are well-behaved with respect to propositional logic; they are not even "systems of modal logic" in the sense of Chellas [3]. The first semantics (the "simple" semantics) fails to satisfy $w \models \varphi$ for some tautologies, while the second (the "distributive" semantics) is not closed under modus ponens. Since we are presently interested in a semantic analysis, we don't consider these problems fatal to our project, but they are undeniably disappointing (and simultaneously interesting).

Our approach is distinguished from previous discussions of the ability to bring about an end in a number of ways, and not merely by virtue of explicit probabilities. In both [2] and [7], the authors explore the logic of bringing about an end, but they both omit the means the agent employs. Of course, how an agent brings about an end may be irrelevant for certain analyses, but if we want to discuss means, then we must include the actions. Pollock's work on the logical foundations of means-end reasoning [9] is similar to what we do here, but his goal is different. He is interested in planning algorithms for intelligence, and so it is an exercise in practical reasoning. We are concerned instead with a conceptual analysis of means-end ascriptions through formalization, and so we stress semantic issues while ignoring formal deduction (in the presence of probabilities, it may be unlikely that traditional formal deduction will do us much good anyway).

Alternatively, our models are similar to the probabilistic transition systems in computer science. One can reason about and characterize these structures using the formal language "Probabilistic Computation Tree Logic" or PCTL [6, 1]. PCTL is a rich language with some constructions ("unless" and "until" operators, for instance) that ought to be incorporated in our setting, although we do not do so here. However, PCTL treats probabilities quite differently than we do, by introducing modal operators $[P]_{\geq p}$ representing that f has at least probability p of being true. These operators don't compose in the way that we'd like: there's not a single formula representing that, if one does m and then n, then P has probability at least p of being true. We can specify $[[P]^n_{p_1}]^m_{p_2}$ for each pair p_1 and p_2 such that $p_1 p_2 = p$, but this doesn't capture our intended meaning.

We begin by discussing the features common to both of our semantic approaches. In Section 3 we introduce the simple semantics and in Section 4 we briefly describe the distributive semantics, but unfortunately must omit technical details for reasons of space. In the conclusion, we compare the two but we are presently unable to provide a final judgment on whether the simple or the distributive semantics are appropriate for our setting.

2. The common core

In each of the two semantics we consider, our model consists of a set $\mathcal{W}$ of worlds, a function $[\![-]\!] : \mathbf{prop} \to \mathcal{P}(\mathcal{W})$ mapping atomic propositions to the set of worlds in which they are true and a function $[\![m]\!] : \mathcal{W} \to \mathcal{D}(\mathcal{W})$ for each $m \in \mathbf{act}$, where $\mathcal{D}(X)$ is the set of discrete distributions on X. That is,

$$\mathcal{D}(X) = \{d : X \to [0,1] \mid \sum_{x \in X} d(x) = 1\}.$$

For w, $w' \in \mathcal{W}$ and $m \in \mathbf{act}$, we interpret $[\![m]\!](w)(w')$ as the probability that the result of doing m in world w is world w'. We prefer to be a bit agnostic regarding our interpretation of probabilities. We don't want to commit ourselves whether we mean a subjective or propensity interpretation yet, but hope that what we develop here works reasonably well regardless of which interpretation we adopt.

We will use these distributions to define a likelihood[2] function

$$l : \mathcal{W} \times \mathbf{ME} \to [0,1],$$

assigning to each pair $\langle w, \varphi\rangle$ a value $0 \leq l(w,\varphi) \leq 1$ representing the likelihood that φ is satisfied in world w. For purely propositional formulas, this likelihood is always 0 or 1; we choose to interpret these formulas as either true or false, with no uncertainty involved[3]. For formulas involving dynamic operators, the uncertainty reflects that the outcomes of certain actions are unpredictable.

The likelihood function is intended to provide a sense of efficacy for our means. For purely propositional formulas φ, the efficacy (in w) of an act m in bringing about φ is given by

$$l(w, [m]\varphi) = \sum_{w' \models \varphi} [\![m]\!](w)(w').$$

In other words, the efficacy of m in w for bringing about φ is just the probability that the world which results from doing m in w is a world satisfying φ.

How we calculate the efficacy for $[m]\varphi$ when φ is not purely propositional is the essential difference between the two semantics we discuss here, which we tentatively dub the *simple* likelihood semantics and the *distributive* likelihood semantics.

Regardless of which approach we take, we define the satisfaction relation $\models$ by fixing an $\alpha \in [0,1]$ and saying $w \models \varphi$ iff $l(w,\varphi) \geq \alpha$. Thus, α is a measure of our willingness to accept means to ends. The greater α, the fewer true ascriptions "in w, m is a means to φ." If we fix a high α, then we require a greater certainty that m will actually yield φ before we say that m is a means to φ. It is probably sensible to require that $\alpha > \frac{1}{2}$, since otherwise we may have $w \models \varphi$ and $w \models \neg\varphi$ for some φ.

The simple and distributive semantics treat negations identically: for both the simple and distributive semantics, $l(w, \neg\varphi) = 1 - l(w,\varphi)$, although the treatment of negations is a source of many difficulties for the simple semantics. As consequence,

---

[2]We use the non-technical term *likelihood* rather than the technical term *probability* here, just because our likelihood function does not define a distribution on any set.

[3]We assume that in each world each atomic proposition is either true or false, but this is probably not essential and may be relaxed when appropriate.

negations commute with dynamic operators. This is why we have introduced only one operator $[m]$. Its conjugate $\neg[m]\neg$ is equivalent to $[m]$. There is no "weak" dynamic operator in our development of likelihood semantics.

Finally, we turn to the distinctive feature of our approach to likelihood semantics. Our analysis of conjunctions, especially conjunctions of the form $[m]\varphi \wedge [n]\psi$ where m and n are distinct, is central to the subsequent semantics (both simple and distributive) and it is fair to say that our development stands or falls dependent on this analysis. Recall that we interpret the conjuncts $[m]\varphi$ and $[n]\psi$ as asserting that m is a means to φ and n is a means to ψ, respectively. The conjunction, then, should be true in w just in case it is true in w that m is a means to φ and n to ψ, i.e., it is true just in case $l(w, [m]\varphi) \geq \alpha$ and $l(w, [n]\psi) \geq \alpha$. This analysis suggests the following identity.

$$l(w, [m]\varphi \wedge [n]\psi) = \min\{l(w, [m]\varphi), l(w, [n]\psi)\}. \tag{1}$$

When $m \neq n$, both the simple and distributive semantics agree on this point. Nonetheless, since the identity is uncommon in dynamic logics even with probabilities involved, it is worthwhile to give a few more words in its favor.

For now, let us take a agent-centric view of our means and ends and consider what $w \models [m]\varphi \wedge [n]\psi$ means to the agent. It does not mean that the agent can choose to do *both* m and n in order to yield both φ and ψ. In our current setting, actions are interpreted one at a time and sequentially[4] and so, as soon as he does m (or n), he is (perhaps) no longer in world w. Instead, the conjunction represents a *choice* for our agent. If he decides to pursue φ, then he can likely bring it about by doing m and if he decides to pursue ψ, then he can likely bring it about by doing n. In the special case when φ and ψ are the same formula, then he can likely bring about this end by doing either m or n.

Given that the conjunction $[m]\varphi \wedge [n]\psi$ represents a choice between two actions to yield two (possibly distinct) ends, then we should not require that it is likely that both doing m yields φ and also doing n yields ψ. On the contrary, we are only interested in assuring that (a) it is likely that doing m yields φ and (b) it is also likely that doing n yields ψ. But this is just what Equation (1) says.

Let us summarize the similar features of our two approaches to likelihood semantics. For each, we assume that each action m comes with a an evaluation for atomic propositions $[\![-]\!] : \mathbf{prop} \to \mathcal{P}(\mathcal{W})$ and a transition structure $[\![m]\!] : \mathcal{W} \to \mathcal{D}(\mathcal{W})$. In each case, we use these structures to define a likelihood function $l : \mathcal{W} \times \mathbf{ME} \to [0,1]$ and write $w \models \varphi$ iff $l(w, \varphi) \geq \alpha$ for fixed α. These functions coincide on the purely propositional formulas. In fact, they coincide on all formulas of the form $[m_1] \ldots [m_k]\varphi$ where φ is propositional. They treat all negations and certain conjunctions similarly.

The fundamental difference between the two can be characterized thus: in the simple semantics, the equation (1) holds when $m = n$ as well, but in the distributive semantics, this is not so. More specifically, for the distributive semantics, each dynamic operator $[m]$ distributes over conjunction, i.e.,

$$l(w, [m]\varphi \wedge [m]\psi) = l(w, [m](\varphi \wedge \psi)). \tag{2}$$

The remainder of the paper is devoted to how we effect this difference and the consequences it entails.

[4]This is a fairly strong assumption and may be reconsidered in subsequent research.

3. The simple likelihood semantics

The simple likelihood semantics defines a likelihood function $l : \mathcal{W} \times \mathbf{ME} \to [0,1]$ by induction on the structure of the formula. Explicitly, l is defined as follows.

(A) For atomic P, we define $l(w, P) = 1$ iff $w \in [\![P]\!]$, and $l(w, P) = 0$ otherwise.
(B) $l(w, \neg\varphi) = 1 - l(w, \varphi)$.
(C) $l(w, \varphi \wedge \psi) = \min\{l(w, \varphi), l(w, \psi)\}$.
(D) $l(w, [m]\varphi) = \sum_{w' \in \mathcal{W}} [\![m]\!](w)(w') \cdot l(w', \varphi)$.

The main issue before us is whether the clause for conjunction is appropriate. We have given an analysis in support of this clause in the case in which m and n are distinct and indeed our distributive semantics agrees on this point. The question then is: Should the likelihood of $[m]\varphi \wedge [m]\psi$ really be calculated as the minimum of the two likelihoods?

The short answer is that the previous argument for Equation (1) applies just as easily here. The conjunction $[m]\varphi \wedge [m]\psi$ means just that doing m likely yields φ and also doing m likely yields ψ. It does not assert that doing m likely yields *both* φ and ψ simultaneously. So our previous argument is also an argument for (C).

Unfortunately, this semantics has some particularly unattractive features. It assigns likelihood less than 1 to some tautologies (like $[m]\varphi \vee \neg[m]\varphi$[5], which has the same likelihood as $[m]\varphi \vee [m]\neg\varphi$) and consequently it assigns likelihood greater than 0 to some contradictions. However, for other tautologies (for instance, $P \vee \neg P$ for atomic P), the likelihood in each world is 1. Thus, we see that there are tautologically equivalent formulas (i.e., formulas provably equivalent in propositional logic) φ and ψ such that $l(w, \varphi) \neq l(w, \psi)$. Our simple semantics are badly behaved with respect to tautological equivalence: there are worlds in which $w \models \varphi$ but $w \not\models \psi$ even though $\varphi \equiv \psi$ is provable in propositional logic.

So how does the problem arise? It is not hard to see that, for any formulas φ and ψ and any world w, if $w \models \varphi \to \psi$ and $w \models \varphi$, then $w \models \psi$ (here, we rely on the assumption that $\alpha > \frac{1}{2}$). Hence, given any Hilbert-style axiomatization of propositional calculus, there are axioms which our simple semantics fail to satisfy. But we can be more specific than that.

The set $[0,1]$ is a distributive lattice (totally ordered and complete, in fact). Our definition of l above satisfies that

$$l(w, \varphi \wedge \psi) = l(w, \varphi) \wedge l(w, \psi),$$
$$l(w, \varphi \vee \psi) = l(w, \varphi) \vee l(w, \psi).$$

Hence, in every model, the axioms of a distributive lattice will be satisfied. For instance, in every world w, we have

$$l(w, \varphi \wedge (\psi \vee \theta)) = l(w, (\varphi \wedge \psi) \vee (\varphi \wedge \theta)),$$

and so a world satisfies the formula on the left hand side iff it satisfies the corresponding formula on the right.

But the problem is that our negation is not inherited from the structure on $[0,1]$ and so it does not interact well with conjunction and disjunction (and, consequently, implication behaves badly too). The ordered set $[0,1]$ is a Heyting algebra, but we

[5]Although our example here involves excluded middle, the problems with our simple semantics are not really about intuitionistic vs. classical logic.

didn't define our semantics so that

$$l(w, \varphi \to \psi) = l(w, \varphi) \to l(w, \psi), \tag{3}$$

$$l(w, \neg\varphi) = \neg l(w, \varphi). \tag{4}$$

We could, of course, amend our definition to do just this, but unfortunately the negation in $[0,1]$ is trivial and unsuitable for our purposes: in any Heyting algebra with total order, $\neg x = \bot$ if $x \neq \bot$ and $\neg\bot = \top$. Hence, if we follow this approach, then whenever $l(w, \varphi) > 0$, we have $l(w, \neg\varphi) = 0$. This does not support our intuitions about likelihoods. (Note that the implication on $[0,1]$ is similarly unsuitable.)

Alternatively, we could take our codomain to be some Heyting algebra other than $[0,1]$, one with a proper partial order (like in [4, 5]). Unfortunately, this conflicts with our desire to satisfy the principle (1) from Section 2. We chose a total order just so we could take minima of likelihoods. If we use a proper partial order, then in most models we will have situations in which $l(w, [m]\varphi \wedge [n]\psi)$ strictly less than $l(w, [m]\varphi)$ and $l(w, [n]\psi)$. We do not wish to abandon the argument against this situation, so we are loathe to fix things by moving to a partial order.

4. The distributive likelihood semantics

The distributive semantics rejects the claim that $l(w, [m]\varphi \wedge [m]\psi)$ ought to be calculated by taking the minimum of the likelihoods $l(w, [m]\varphi)$ and $l(w, [m]\psi)$, just the same as for distinct actions. When we argued for the principle (1) for distinct actions in Section 2, part of our reasoning relied on interpreting $w \models [m]\varphi \wedge [n]\psi$ as a choice for the agent: If he chooses to pursue φ, then he can do m and if he chooses to pursue ψ, then he may do n.

But that a world satisfies $[m]\varphi \wedge [m]\psi$ does not present a similar choice for our agent. If he chooses to pursue φ or ψ, then the fact that $w \models [m]\varphi \wedge [m]\psi$ suggests just that he can do m. Of course, in w there might be other means to these ends, but the point is that the formula $[m]\varphi \wedge [m]\psi$ does not present a choice for the agent. It is only a statement about the sufficiency of m itself.

In the simple semantics, we calculated the likelihood of a conjunction by the likelihood of the conjuncts. Here, we want instead to calculate the likelihood that the transition structure that actually attains when we do m satisfies the formula $[m]\varphi \wedge [m]\psi$. When one does m, only one world is the actual result and so there is "after the fact" just a single m-transition out of w. We don't want to assent to $[m]\varphi \wedge [m]\psi$ if it is unlikely that this m-transition satisfies $[m]\varphi \wedge [m]\psi$ in the usual Kripke-style sense of satisfaction.

Of course, we must be careful to retain the feature that the likelihood of $[m]\varphi \wedge [n]\psi$ is calculated as a minimum of the likelihoods of the conjuncts when $m \neq n$. We do this by imagining a tree in which each arrow is labeled with an action. We compute the likelihood that each of the immediate subtrees (the trees one label down) satisfies what it ought and take the minimum of these likelihoods. Because we require that each node has at most one arrow labeled by a given action, our conjunction $[m]\varphi \wedge [m]\psi$ is not satisfied unless it is likely that the m-transition which actually attains satisfies the formula.

Space precludes us from a full, technical presentation of the distributive semantics. Instead, we will try to sketch the main ideas behind the semantics, largely by following an example. Consider the formula $[m]\varphi \wedge ([n]\psi \vee [m]\vartheta)$, which we will abbreviate α. We want a measure of the likelihood that (in w) the transition structures which actually attain for m and n satisfy α, subject to the proviso that $l(w, \alpha)$ is no less than $l(w, [m]\varphi)$ and $l(w, [n]\psi)$ in keeping with our discussion of (1) in Section 2.

Now, any transition structure that satisfies α must also satisfy $[m]\varphi$ and at least one of $[n]\psi$ or $[m]\vartheta$. Thus, there are essentially two ways that a transition structure satisfies α, namely,

(a) the structure satisfies $[m]\varphi$ and $[n]\psi$;
(b) the structure satisfies $[m]\varphi$ and $[m]\vartheta$.

Think of (a) and (b) as *strategies* for satisfying α. Then $l(w, \alpha)$ is calculated by computing the likelihood of the success of each of these strategies and taking the maximum[6].

A little more explicitly: We define the dynamic tree of a formula χ to be a tree with edges labeled by actions (at most one edge labeled with each action) so that the structure of the tree reflects the nesting of dynamic operators in the formula. The dynamic tree of the formula α above is given by

$$\begin{tikzcd} & * \arrow[dl, "m"'] \arrow[dr, "n"] & \\ \mathsf{t}(\varphi) \cup \mathsf{t}(\vartheta) & & \mathsf{t}(\psi) \end{tikzcd}$$

where $\mathsf{t}(\chi)$ represents the dynamic tree for χ and $\cup$ joins two trees by identifying their roots and any identically labeled edges with common source. We define an *interpretation* h of such a tree to be a mapping of the nodes in the tree to worlds in $\mathcal{W}$, i.e., $h : \mathsf{t}(\chi) \to \mathcal{W}$ and define what it means for such an interpretation h to satisfy the formula χ.

A *strategy* is a set $\mathcal{S}$ of interpretations with common root w closed under *substitution*, where substitution is the act of replacing one branch of an interpretation $h \in \mathcal{S}$ with a same-labeled branch from another interpretation $g \in \mathcal{S}$. We next define the likelihood $v(\mathcal{S})$ of the transition which attains falling inside a given strategy $\mathcal{S}$. This is the step in which we take a minimum of $v(\mathcal{S}_m)$ for subtrees of $\mathcal{S}$ beginning with m. Finally we define

$$l(w, \chi) = \max \{v(\mathcal{S}) \mid \mathcal{S} \in I_w(\chi)\},$$

where $I_w(\chi)$ is the set of all strategies $\mathcal{S}$ for $\mathsf{t}(\chi)$ rooted at w.

These semantics have some interesting features. As is common with probabilistic approaches, satisfaction is not closed under conjunction. It is not always the case that if $w \models \vartheta$ and $w \models \chi$ then $w \models \vartheta \wedge \chi$.

However, unlike the semantics in Section 3, for the distributive semantics we have that $l(w, \varphi) = 1$ for every tautology φ (and similarly the likelihood of contradictions is 0). Furthermore, if φ and ψ are tautologically equivalent, then $l(w, \varphi) = l(w, \psi)$.

But if distributive semantics satisfies tautologies and is not closed under conjunction, then it must be the case that our distributive semantics is not closed under the inference rule modus ponens. This is easy to see: For every world w, we

[6]In this example, the same result could be obtained by using the distributive law to rewrite α and then applying a naive calculation, but algebraic transformations do not play an explicit role in our semantics.

have

$$l(w, [m]\varphi \to ([m]\psi \to ([m]\varphi \wedge [m]\psi))) = 1, \tag{5}$$

but there are models in which w satisfies each conjunct $[m]\varphi$ and $[m]\psi$, but not the conjunction.

A similar situation occurs in Kyburg's "Conjunctivitis" [8], prompting his move to a *Weak Deduction Principle*. This principle requires that, if S is a body of reasonably accepted statements, $s_1 \in S$ and $\vdash s_1 \to s_2$, then $s_2 \in S$. Thus, S is not closed under all applications of modus ponens, but only those in which the conditional is a theorem of our logic. If we weaken our logic so that the formula (5) is not a theorem (obviously, we cannot take full propositional logic), then we perhaps satisfy a weak deduction principle as well. We will not investigate this further here, however.

5. Conclusion

We are interested in likelihoods in a dynamic logic in which $[m]\varphi$ represents the statement, "m is a means to φ." We have introduced probabilities into our models by means of functions $[\![m]\!] : \mathcal{W} \to \mathcal{D}(\mathcal{W})$, where $[\![m]\!](w)(w')$ represents the probability that world w' is the result of doing m in w. We then introduced a likelihood function $l : \mathcal{W} \times \mathbf{ME} \to [0, 1]$ where $l(w, \varphi)$ represents the likelihood that φ is satisfied in w and, fixing α, we write $w \models \varphi$ iff $l(w, \varphi) \geq \alpha$. If we take this basic approach for granted, then the main question is clearly: how ought our likelihood function be defined?

We presented two alternatives, dubbed the simple and distributive likelihood semantics, respectively. Both agreed that when $m \neq n$, then

$$l(w, [m]\varphi \wedge [n]\psi) = \min\{l(w, [m]\varphi), l(w, [n]\psi)\},$$

and also that $l(w, \neg\varphi) = 1 - l(w, \varphi)$ (although this is the source of some problems for the simple semantics). The simple semantics adopts the above identity for conjunctions even when $m = n$, whereas the distributive semantics instead imposes the requirement that a dynamic operator distributes over conjunction.

Both approaches behave badly with respect to propositional logic, which is perhaps not surprising: probabilities often seem to conflict with purely deductive reasoning. The simple semantics is closed under modus ponens and assigns likelihood 1 to axioms of distributive lattices, but it assigns likelihood less than 1 to some propositional tautologies and consequently fails to assign the same likelihood to formulas which are tautologically equivalent. The distributive semantics, on the other hand, assigns likelihood 1 to every tautology and also satisfies $l(w, \varphi) = l(w, \psi)$ whenever φ and ψ are tautologically equivalent. However, it is not closed under modus ponens: there are situations in which $w \models \varphi \to \psi$ and $w \models \varphi$ but $w \not\models \psi$. All in all, though, it appears that the structural features of distributive semantics are preferable to the simple semantics. The fact that the simple semantics assign different likelihoods to tautologically equivalent formulas seems a serious drawback, more significant than failure to be closed under modus ponens.

On the other hand, the motivating argument for the simple semantics seems at least as persuasive as the corresponding argument for the distributive semantics. The distributive semantics is vulnerable to charges that it avoids the problems of the simple semantics by effectively redefining conjunction. In the distributive semantics, we consider the likelihood that the transition structure which actually attains satisfies the conjunction $[m]\varphi \wedge [m]\psi$, whereas the simple semantics is interested only in the likelihood of each conjunct. When it comes to the persuasiveness of the motivating arguments for each semantics, it appears that the simple semantics has an edge over the distributive semantics.

It is this tradeoff that we find most interesting here. Our initial goal seems simple and natural enough: how does one explicitly add a notion of likelihood to dynamic logic in order to express that means differ in their efficacy? The result we have sketched involves two distinct approaches, the first easily motivated but less well-behaved than the second. As well, another tradeoff even appears in the structural features of the semantics: the simple semantics is well-behaved with respect to modus ponens but not with respect to tautologies and equivalences, while the opposite is true with the distributive semantics. Which of these approaches is more suitable for means-end ascriptions (or what alternatives do a better job than either) is clearly an open question, ripe for further consideration.

References

[1] Bianco and de Alfaro. Model checking of probabilistic and nondeterministic systems. *FSTTCS: Foundations of Software Technology and Theoretical Computer Science*, 15, 1995.

[2] Mark Brown. On the logic of ability. *Journal of Philosophical Logic*, 17:1–26, 1988.

[3] Brian F. Chellas. *Modal Logic: An Introduction*. Cambridge University Press, 1980.

[4] Melvin Fitting. Many-valued modal logics. *Fundamenta Informaticae*, 15:235–254, 1991.

[5] Melvin Fitting. Many-valued modal logics, II. *Fundamenta Informaticae*, 17:55–73, 1992.

[6] Hans Hansson and Bengt Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.

[7] John F. Horty and Nuel Belnap. The deliberative stit: A study of action, omission, ability and obligation. *Journal of Philosophical Logic*, 24:583–644, 1995.

[8] Henry E. Kyburg. Conjunctivitis. In Marshall Swain, editor, *Induction, Acceptance and Rational Belief*, pages 55–82. Reidel, 1970.

[9] John L. Pollock. The logical foundations of means-end reasoning. In Renée Elio, editor, *Common Sense, Reasoning and Rationality*. Oxford University Press, 2002.

Section of Philosophy and Ethics of Technology, Technical University of Eindhoven
E-mail address: J.Hughes@tm.tue.nl