

Means-end Relations and a Measure of Efficacy

Jesse Hughes (j.hughes@tm.tue.nl)
Section of Philosophy, Technical University of Eindhoven

Albert Esterline (esterlin@ncat.edu) and Bahram Kimiaghalam
(bahram@ncat.edu)
Department of Computer Science, North Carolina A&T State University

Abstract. Propositional dynamic logic (PDL) provides a natural setting for semantics of means-end relations involving non-determinism, but such models do not include probabilistic features common to much practical reasoning involving means and ends. We alter the semantics for PDL by adding probabilities to the transition systems and interpreting dynamic formulas $\langle\alpha\rangle\varphi$ as fuzzy predicates about the reliability of α as a means to φ . This gives our semantics a measure of efficacy for means-end relations.

Keywords: Means-end relations, fuzzy modal logic, propositional dynamic logic, efficacy

1. Introduction

Practical reasoning is concerned with deriving actions (or intentions to act) from certain propositions. This distinct form of reasoning has been studied at least since Aristotle's time and enjoyed renewed interest recently beginning with von Wright's landmark article (von Wright, 1963). The topic has gained a wider audience in recent years, due to its application in artificial intelligence (in, e.g., (Pollock, 2002)).

We are especially interested in arguments involving means-end relations. A typical argument of this kind involves premises like the following:

- (1) an assertion that an agent A desires some *end* φ ,
- (2) an assertion that (possibly given some precondition ψ) the action α is related to the realization of φ ,
- (3) an assertion of some factual matter, such as that the precondition ψ is true.

Premises of type (2) express causal relations about the world (or, perhaps, *beliefs* about causal relations). Such premises are essential to practical reasoning, since they give the motivational force for the argument. The reason to *do* the action α is that it is related in the right



© 2005 Kluwer Academic Publishers. Printed in the Netherlands.

way to the desired condition φ . Because one wants φ to be realized, he will be motivated to do α . We call such premises (*conditional means-end relations*), since they assert that the action α is a *means* to the end φ . We will focus here on *local means-end relations*, which assert that *in this world*, the action α is related to the realization of φ , independent of any precondition.

In order to evaluate practical syllogisms (like those found in (von Wright, 1963)), one needs a semantics for the premises involved. In particular, means-end syllogisms should be evaluated with respect to the meaning of the means-end relations. This work is part of an ongoing project to understand the natural language semantics of means-end talk, partly in order to evaluate syllogisms like those discussed by von Wright. But our interest also arises from a different source.

Artifactual functions are evidently associated with means-end relations. A functional ascription (“Staplers are for fastening papers.”) entails a related means-end assertion, e.g., that staplers can be used to fasten papers. We aim to investigate the relation between functions and means-end relations, but we need a clear semantics for means-end relations to get the project off the ground. Hence, our interest in means-end semantics extends beyond their roles in practical syllogisms. We want a semantics for “ α is a means to φ ” that is independent of the desirability of φ . The degree to which a nuclear warhead is appropriate for decimating cities is independent of the desirability of such mass killings¹.

As a consequence, our work will appear different than many of the discussions of means-end relations, goal planning and the like found in current AI literature. Our primary aim is to present a formal semantics for means-end relations which (1) approximates natural language means-end talk and (2) can be used in a semantics for artifactual functions². For this reason, we avoid many of the difficult topics of practical reasoning, including the frame and ramification problems. Certainly, one would like the resulting semantics to be applicable in the actual reasoning one does from ends to means, but we postpone these issues for present and use a simple formalism for our initial development.

We use Propositional Dynamic Logic (PDL) as our formal setting for means-end relations. An end is a condition we wish to make true (or keep true). A means must be some way of changing the world so that our end is realized. This suggests possible world semantics and

¹ Perhaps the primary use of nuclear weapons is deterrence, but the reason that such weapons deter is that everyone involved believes they are effective in mass killing.

² See (Hughes, 2005) for preliminary work on the relationship between means-end relations and artifactual functions.

that means correspond to transition structures between worlds, i.e., to actions in a dynamic logic. We realize that this conceptual setting does not accommodate all natural language means-end talk and in particular does not include the common reference to objects as means. We will discuss these kinds of means-end relations in a subsequent paper and focus here on actions as means.

An alternative tradition for practical reasoning involves temporal logic. Recently, Mark Brown (Brown, 2005) suggested a means-end semantics involving such logics with *stit* (see-to-it-that) operators. His logic includes sophisticated temporally defined ends, such as making φ true for a certain period, attainable for some time in the future and so on, and these are useful features lacking in our present account. However, he identifies means to an end as certain formulas expressing ability, which does not seem quite right. Indeed, since his logic has no place for actions as syntactic entities, it's hard to see how it can represent means at all. Rather, it seems closer to a logic of ability (very different from his prior account in (Brown, 1988)). Thus, despite the attractive features of Brown's use of temporal logic, we prefer PDL for means-end semantics, since we are committed to means as actions and, as Meyer says (Meyer, 2000), in PDL actions appear as "first-class citizens".

In choosing PDL for our means-end semantics, we are following a healthy tradition in current AI research, surveyed in (Meyer, 2000), with additional examples in (Castilho et al., 2002; Castilho et al., 1999; Prendinger and Schurz, 1996; Zhang and Foo, 2002). But where this work is concerned with feasibly deriving plans from goals or defeasibly deducing consequences of actions given partial information, we are interested in the conceptual analysis of natural language via formal semantics. Consequently, we allow ourselves the luxury of *local* means-end relations as a first approximation in order to simplify presentation and defer the real and difficult problems involving the reasoning capacities of actual agents.

One of the primary characteristics of means to an end is its efficacy: its propensity to achieve the related end. Many of the plans we form involve uncertainty, whether due to the agents' ignorance (when the agent plays Three Card Monte, say), randomness in the world (playing a game of chance, e.g., craps) or unpredictability of skills-based actions (playing a game of physical skill, e.g., darts). Thus, there are several reasons one may wish to assign probabilities to the outcomes of ones' actions. Our agent comparing options for achieving his end will surely be interested in the likelihood that each action will bring about his end. Efficacy does not determine the agent's choice, but it often influences it.

Moreover, efficacy is essential in a semantics for artifactual functions. Different artifact tokens may be distinguished by their effectiveness in fulfilling their functional roles. A token which performs particularly badly (when compared to normal tokens of its type) is a malfunctioning token. In order to develop these ideas, it is necessary to add a measure of efficacy to our means-end semantics³. But artifactual efficacy does not depend on utility, goal planning or many of the related topics of agent and action research. Hence, we investigate efficacy without introducing a measure of desirability for the related end. In other words, we add probabilities to our semantics without developing a decision theoretic approach common in much of the AI literature.

There is another reason that we avoid a utility-based semantics. The desirability of φ is distinct from whether α is a reliable means to φ . Desirability is relevant in an agent's decision, but it is not relevant in determining the efficacy of means to a given end. Thus, while Qualitative Decision Theory (QDT) (Boutilier, 1994) represents goals as formulas, it nonetheless presupposes a desirability ranking of possible worlds which we find inappropriate for our current analysis. Of course, we could use QDT to measure efficacy, by calculating the expected utility of various means under the assignment of utility one to φ worlds and zero to non- φ worlds. But in this case, it is unclear how to measure efficacy when φ is a complex end involving dynamic operators. See Section 3 for a brief discussion of this issue.

Assigning probabilities to the outcomes of actions corresponds to shifting from non-deterministic transition systems to probabilistic transition systems. Probabilistic transition systems are well-studied in computer science, with *Probabilistic Computation Tree Logic* (pCTL) a common language for reasoning about and characterizing such structures (Hansson and Jonsson, 1994; Bianco and de Alfaro, 1995). pCTL is a rich language, and, like temporal logic, includes temporal operators such as “while” and “until” that could be useful in our setting (and may be considered in later work). However, pCTL introduces probabilities by subscripting operators with values in $[0, 1]$. This mechanism doesn't satisfy our needs in reasoning about compositions of actions. We would like to express the reliability of α in realizing φ even when φ is a complex formula involving further actions. In pCTL, one must first decide what probabilities to attach to the actions occurring in φ , and this does not seem appropriate for our analysis of means-end relations.

Instead, we take our inspiration from the connection between fuzzy logic and probabilities found in (Hájek et al., 1995; Hájek et al., 2000).

³ We present a semantics of artifactual functions involving efficacy and malfunction in (Hughes and Zwart, 2005).

Degrees of truth aren't the same as probabilities, but probabilities do define fuzzy predicates. Hajek, et al., discussed the fuzzy predicates “probably φ ” induced by probabilities assigned to a propositional language. Similarly, we will be interested in fuzzy predicates “ α is a reliable means to φ .” The truth degrees of these predicates is defined by the probabilistic transition system.

This approach is closely related to *Probabilistic PDL*, introduced in (Kozen, 1983). There are two primary differences in the approaches. First, Kozen chooses a different interpretation of the propositional connectives, but this is a fairly superficial difference. A more fundamental difference is the application and interpretation of the logics. Where our logic is built on fuzzy propositions, Kozen's analogous type is arbitrary measurable functions. Our interpretation assigns truth degrees to formula-world pairs, whereas Kozen's interpretation assigns uninterpreted arithmetic values. Finally, our dynamic operators are used to construct certain vague propositions, namely that a given action is a “reliable” means to an end. It is less clear how to interpret Kozen's operators.

We begin with an overview of propositional dynamic logic and our definitions of some basic means-end relations. We then discuss probabilistic transition systems and introduce a fuzzy version of PDL for expressing efficacy of means to an end in this setting. In Section 4, we re-introduce non-determinism to our new PDL models by moving to an analog of minimal models (i.e., neighborhood semantics).

2. Propositional dynamic logic

Propositional Dynamic Logic (PDL) is a logic of actions, typically used to reason about computer program behavior but also in representing agent planning to realize goals (Meyer, 2000). We refer the reader to (Harel, 1984), from which we take much of the following material.

The syntax consists of two disjoint types: the set Π_0 of atomic actions and the set Φ_0 of atomic propositions. From these two sets, we inductively define the sets Π of actions and Φ of formulas as follows

- $\{\top\} \cup \Phi_0 \subseteq \Phi$;
- if $\varphi, \psi \in \Phi$ then $\neg\varphi$ and $\varphi \wedge \psi \in \Phi$;
- if $\alpha \in \Pi$ and $\varphi \in \Phi$ then $[\alpha]\varphi \in \Phi$;
- $\Pi_0 \subseteq \Pi$;
- if $\alpha, \beta \in \Pi$ then $(\alpha; \beta)$, $(\alpha \cup \beta)$ and α^* are in Π ;

- if $\varphi \in \Phi$ then $\varphi? \in \Pi$.

We introduce the value \perp , the connectives \vee , \rightarrow and \leftrightarrow and the weak operator $\langle \alpha \rangle$ as usual. The action constructors are intended thus: the semicolon denotes sequential composition, the union non-deterministic choice and the asterisk finite iterations. A formula $[\alpha]\varphi$ roughly expresses that doing α will necessarily realize φ (with no guarantee that α can be done) and $\langle \alpha \rangle \varphi$ that it is possible to realize φ by doing α (and hence that α can be done).

A *PDL frame* \mathbf{F} for Π_0 consists of a set \mathcal{W} of worlds (or states) and an interpretation

$$\llbracket - \rrbracket^{\mathbf{F}} : \Pi_0 \rightarrow (\mathcal{P}\mathcal{W})^{\mathcal{W}}$$

of actions via non-deterministic transition systems. Here \mathcal{P} denotes the powerset functor and exponentiation A^B denotes the set of functions $B \rightarrow A$.

The interpretation $\Pi_0 \rightarrow (\mathcal{P}\mathcal{W})^{\mathcal{W}}$ assigns to each $m \in \Pi_0$ a function $\llbracket m \rrbracket : \mathcal{W} \rightarrow \mathcal{P}\mathcal{W}$. For $w \in \mathcal{W}$, we interpret $\llbracket m \rrbracket(w) \subseteq \mathcal{W}$ as the set of possible⁴ outcomes of doing m in w and we write $w \xrightarrow{m} w'$ for $w' \in \llbracket m \rrbracket^{\mathbf{F}}(w)$. Clearly, a PDL frame is just the same as a labeled transition system with nodes $w \in \mathcal{W}$ and labels $m \in \Pi_0$.

A *PDL model* \mathbf{M} is a frame \mathbf{F} together with a valuation

$$\llbracket - \rrbracket^{\mathbf{M}} : \Phi_0 \rightarrow \mathcal{P}\mathcal{W}$$

of atomic propositions. We abuse notation by adopting Scott brackets for both the valuation of atomic propositions and the interpretation of atomic actions, but since our sets Π_0 and Φ_0 are disjoint, no confusion should result. We omit the superscripts hereafter.

A valuation assigns to each atomic proposition $P \in \Phi_0$ a set $\llbracket P \rrbracket \subseteq \mathcal{W}$ of worlds. We interpret $\llbracket P \rrbracket$ as the set of worlds in which P is true.

We extend the valuation of atomic propositions to a function

$$\llbracket - \rrbracket : \Phi \rightarrow \mathcal{P}\mathcal{W}$$

and the interpretation of atomic actions to a function

$$\llbracket - \rrbracket : \Pi \rightarrow (\mathcal{P}\mathcal{W})^{\mathcal{W}}$$

recursively as in Table I. We say that w *satisfies* φ or that φ *is true* in w just in case $w \in \llbracket \varphi \rrbracket$. In this case, we write $\mathbf{M}, w \models \varphi$ or just $w \models \varphi$ when \mathbf{M} is understood by context. We write $\mathbf{M} \models \varphi$ and say that φ is *valid in* \mathbf{M} if for every $w \in \mathcal{W}$ we have $w \models \varphi$ and we write $\models \varphi$ if $\mathbf{M} \models \varphi$ for every model \mathbf{M} . In this case, we say that φ is *valid*.

⁴ Or “normal” or “reasonably expected” or ...

Table I. Extension of valuation to Φ and interpretation to Π .

<u>On formulas</u>
$\llbracket \top \rrbracket = \mathcal{W}$
$\llbracket \neg\varphi \rrbracket = \mathcal{W} \setminus \llbracket \varphi \rrbracket$
$\llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket$
$\llbracket [\alpha]\varphi \rrbracket = \{w \in \mathcal{W} \mid \llbracket \alpha \rrbracket(w) \subseteq \llbracket \varphi \rrbracket\}$
<u>On actions</u>
$\llbracket [\alpha; \beta] \rrbracket(w) = \{w' \in \mathcal{W} \mid \exists w'' \in \mathcal{W} . w \xrightarrow{\alpha} w'' \xrightarrow{\beta} w'\}$
$\llbracket [\alpha \cup \beta] \rrbracket(w) = \llbracket \alpha \rrbracket(w) \cup \llbracket \beta \rrbracket(w)$
$\llbracket \alpha^* \rrbracket(w) = \bigcup_{n=0}^{\infty} \underbrace{\llbracket [\alpha; \dots; \alpha] \rrbracket(w)}_{n \text{ times}}$
$\llbracket [\varphi?] \rrbracket(w) = \begin{cases} \{w\} & \text{if } w \in \llbracket \varphi \rrbracket; \\ \emptyset & \text{else.} \end{cases}$

Let us call an action α *prohibited in w* if $\llbracket \alpha \rrbracket(w) = \emptyset$. Intuitively, such actions cannot be performed⁵ in w . If α is prohibited in w , then $w \models [\alpha]\varphi$ for any $\varphi \in \Phi$ (including \perp), but $w \not\models \langle \alpha \rangle \varphi$ for any $\varphi \in \Phi$ (not even \top).

Example 2.1. Consider the example of a footrace about to begin⁶. The starter has a (one-shot) pistol and the race will begin once the pistol discharges a blank, but this requires that the pistol is loaded. We will construct a very simple model for this case consisting of only two atomic predicates:

Started true if the race has started,
Loaded true if the pistol is loaded.

⁵ In computer science, one often interprets such actions as *non-terminating*.

⁶ This example is very similar to the Yale shooting problem (Hanks and McDermott, 1987) but we have a different purpose in mind. Because we are primarily interested in semantics of means-end relations rather than practical reasoning, we postpone the difficulties of the frame problem and the appropriateness nonmonotonic logics for later consideration.

Our language will also include two atomic actions:

load the starter loads the pistol,
 fire the starter pulls the trigger.

Note that the action `fire` does not imply that the pistol discharges a blank, but only that the starter pulls the trigger. Our action name `fire` may be a bit misleading in this respect, but it is more suggestive than `pull` and less awkward than `pulltrigger`.

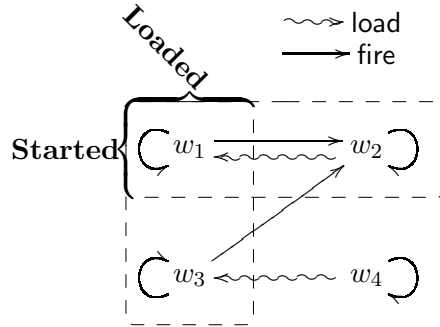


Figure 1. A simple PDL model for a footrace started by firing a pistol.

We consider a model of four worlds, so that each combination of atomic predicates is represented (Figure 1). The transition structure for `load` takes a world in which the pistol is unloaded to one in which it is loaded without altering the truth value for **Started**. The act is prohibited in worlds which already satisfy **Loaded**. One can't load an already loaded pistol.

For worlds in which **Loaded** is false, the action `fire` produces no change. Successfully firing the pistol in the world w_1 in which the race has already begun moves one to w_2 , in which the state of **Started** is unchanged, but **Loaded** is false. Successfully firing the pistol in world w_3 not only unloads the gun, but also starts the race, so again we have a transition to w_2 .

Just to make our model properly non-deterministic, we add the possibility that a blank misfires, say, due to a weak spring in the pistol. In this case, the blank is still capable of firing but our hammer has not struck it with enough force to fire it. Hence, we add reflexive `fire`-transitions in worlds w_1 and w_3 .

The reader may confirm that the equations in Table II are actually satisfied by the model in Figure 1.

To complete our introduction to PDL, we present in Table III the standard axiom system for PDL, taken from (Harel, 1984). This system is sound and complete for our semantics, i.e., $\vdash \varphi$ iff $\models \varphi$.

Table II. Some facts about the model in Figure 1.

$\llbracket \mathbf{Started} \rrbracket$	$= \{w_1, w_2\}$
$\llbracket \mathbf{Loaded} \rrbracket$	$= \{w_1, w_3\}$
$\llbracket [\mathbf{fire}] \mathbf{Started} \rrbracket$	$= \llbracket \mathbf{Started} \rrbracket$
$\llbracket \langle \mathbf{fire} \rangle \mathbf{Started} \rrbracket$	$= \llbracket \mathbf{Started} \rrbracket \cup \llbracket \mathbf{Loaded} \rrbracket$
$\llbracket [\mathbf{load}] \mathbf{Loaded} \rrbracket$	$= \mathcal{W}$
$\llbracket \langle \mathbf{load} \rangle \mathbf{Loaded} \rrbracket$	$= \llbracket \neg \mathbf{Loaded} \rrbracket$
$\llbracket [(\mathbf{load} \cup \mathbf{fire})^*] \mathbf{Started} \rrbracket$	$= \llbracket \mathbf{Started} \rrbracket$
$\llbracket \langle (\mathbf{load} \cup \mathbf{fire})^* \rangle \mathbf{Started} \rrbracket$	$= \mathcal{W}$
$\llbracket [\mathbf{Loaded}?; \mathbf{fire}] \mathbf{Started} \rrbracket$	$= \llbracket \mathbf{Started} \rrbracket \cup \llbracket \neg \mathbf{Loaded} \rrbracket$
$\llbracket \langle \mathbf{Loaded}?; \mathbf{fire} \rangle \mathbf{Started} \rrbracket$	$= \llbracket \mathbf{Loaded} \rrbracket$

Table III. The theory PDL

Axioms	
Tautology	Every propositional tautology
Distributivity	$[\alpha](\varphi \wedge \psi) \leftrightarrow ([\alpha]\varphi \wedge [\alpha]\psi)$
Composition	$[\alpha; \beta]\varphi \leftrightarrow [\alpha][\beta]\varphi$
Choice	$[\alpha \cup \beta]\varphi \leftrightarrow ([\alpha]\varphi \wedge [\beta]\varphi)$
Iteration	$[\alpha^*]\varphi \leftrightarrow (\varphi \wedge [\alpha][\alpha^*]\varphi)$
Test	$[\psi?]\varphi \leftrightarrow (\psi \rightarrow \varphi)$
Induction	$[\alpha^*](\varphi \rightarrow [\alpha]\varphi) \rightarrow (\varphi \rightarrow [\alpha^*]\varphi)$
K	$[\alpha](\varphi \rightarrow \psi) \rightarrow ([\alpha]\varphi \rightarrow [\alpha]\psi)$
Inference rules	
Modus Ponens	$\varphi, \varphi \rightarrow \psi / \psi$
Necessitation	$\varphi / [\alpha]\varphi$

In (Hughes et al., 2005), we presented a semantics for various means-end relations in PDL, including weakly and strongly sufficient and necessary means. We omit the discussion of necessary means here but briefly present the definitions for the two kinds of sufficient means.

DEFINITION 1. *Let $w \in \mathcal{W}$. We say that an action α is a weakly sufficient means to a formula φ if $w \models \langle \alpha \rangle \varphi$, i.e., if there is a w' such that $w \xrightarrow{\alpha} w'$ and $w' \models \varphi$.*

We say that α is a strongly sufficient means to φ if $w \models [\alpha]\varphi \wedge \langle \alpha \rangle \top$, i.e., if for every w' such that $w \xrightarrow{\alpha} w'$ we have $w' \models \varphi$ and furthermore α is not prohibited in w ⁷.

Our use of PDL is similar to its application in agent planning, including (Castilho et al., 2002; Castilho et al., 1999; Meyer, 2000; Prendinger and Schurz, 1996; Zhang and Foo, 2002). Because we are primarily interested in semantics rather than practical reasoning, however, we defer the difficult problems found in such works, including the frame, qualification and ramification problems. Instead, in this first approximation, we focus on “local” means-end relations — the relations that hold at a particular world in a fully defined model. Actual agents would not typically have the luxury of reasoning about such relations, since they would not know which possible world is the actual world. Nonetheless, such local relations are appropriate for our task here, which focuses on introducing a measure of efficacy to means-end relations. In later work, we may consider conditional means-end relations which more closely represent an agent’s knowledge.

With these caveats in mind, let us tentatively sketch some practical consequences of local means-end relations. In order to realize φ , one must do *some* weakly sufficient means. If he does a strongly sufficient means, then he can be sure to realize his end. One may be tempted to say that a strongly sufficient means is more efficacious than a (strictly) weakly sufficient means, but this seems a bit naive. It depends on the source of the non-determinism in our model. If the non-determinism reflects only an agent’s non-deterministic choices in performing the act, then a weakly sufficient means is as sufficient as its strong counterpart. If the non-determinism is due to influences beyond the control of the agent (or reflects an agent’s ignorance), then a strongly sufficient means is more efficacious than a properly weakly sufficient means.

In any case, this is a very crude, all-or-nothing measure of efficacy. Non-determinism allows for a set of possible outcomes, but possibility doesn’t lend itself well to measuring the likelihood of attaining one’s end. It is very common that an act with uncertain outcome still allows one to estimate the relative probability of each outcome and this is surely an important feature of practical reasoning. In order to represent this feature, we shift from non-deterministic transition systems to probabilistic transition systems.

⁷ In (Zhang and Foo, 2002), the authors interpret $[\alpha]\varphi \wedge \langle \alpha \rangle \top$ as “ α must cause φ ”.

3. Adding probabilities

Propositional dynamic logic is a natural choice for means-end semantics but how to adapt PDL to probabilistic transition systems so that we can include a measure of efficacy? There are two traditional approaches to altering PDL for probabilistic transition systems. The first, as found in (Hansson and Jonsson, 1994), involves indexing the dynamic operators by values in $[0, 1]$ and interpreting $[\alpha]_{\geq x}\varphi$ as the probability that φ is realized after doing α is at least x . This approach doesn't lend itself to measuring the reliability of α in realizing φ when φ involves dynamic operators: one is forced to choose reliability values for each operator in φ . Alternatively, probabilistic PDL (as presented in (Kozen, 1983)) assigns a single value to each world-formula pair, but does so in a purely abstract manner: propositional variables are replaced by arbitrary measurable functions and this induces measurable functions for each "formula" syntactically constructed, but these functions are uninterpreted.

We seek a minimal change to our PDL semantics so that means-end relations include a measure of efficacy. This measure of efficacy should be naturally induced by adding probabilities to the possible world semantics of PDL. Our desire for simplicity suggests that the resulting semantics be truth-functional, but the conflicts between truth-functional semantics and probabilities are well-known. Nonetheless, truth-functional semantics *are* arguably appropriate for vague predicates, and the notions of efficacy and reliability are characteristically "fuzzy". Thus, we pursue a truth-functional semantics for efficacy using probabilities to define fuzzy propositions about the reliability of a means to an end. The resulting semantics is structurally similar to probabilistic PDL, but with explicit interpretations of the values assigned to world-value pairs: these values are *truth-degrees* in fuzzy logic.


We begin by adding probabilities⁸ to the PDL frames discussed in Section 2. A PDL frame is a set \mathcal{W} of worlds with a dynamic interpretation $\llbracket - \rrbracket : \Pi_0 \rightarrow (\mathcal{PW})^{\mathcal{W}}$, where $w' \in \llbracket m \rrbracket(w)$ means that w' is a possible outcome of doing m in w . We replace this *set* of possible outcomes with a *discrete sub-distribution* on the set of worlds. In other words, we are interested in interpretations

$$\llbracket - \rrbracket : \Pi_0 \rightarrow (\mathcal{DW})^{\mathcal{W}}$$

⁸ We are purposely agnostic about the interpretation of probabilities. We hope that the resulting formalism works equally well with objective and subjective probabilities and we use examples involving each of these kinds of probabilities.

where $\mathcal{DX} = \{p : X \rightarrow [0, 1] \mid \sum_{x \in X} p(x) \leq 1\}$. Note that the distribution depends only on the action and current world, and not on the history that brought us to this world. In other words, adding probabilities to PDL this way naturally involves the so-called *Markov assumption* (Puterman, 1994) that states determine the relevant probabilities independent of history.

Given $m \in \Pi_0$ and $w, w' \in \mathcal{W}$, we interpret $\llbracket m \rrbracket(w)(w')$ as the probability that doing m in w results in world w' . For example, suppose that `flip` denotes the action of flipping a coin and \mathbf{H} (\mathbf{T} , resp.) denotes the proposition that the coin just flipped came up heads (tails, resp.). The transition structure for a world in which the coin has not been flipped may look something like this.




We have chosen interpretations involving sub-distributions instead of distributions so that some actions have a positive probability (less than one) of being prohibited. The value

$$1 - \sum_{w' \in \mathcal{W}} \llbracket m \rrbracket(w)(w')$$

is the probability that act m is *prohibited* in world w . We denote this by $\llbracket m \rrbracket(w)(\emptyset)$. In most examples we have in mind, this value is either 0 or 1 for atomic actions m . For composite actions α (in particular, those involving the test operator), we cannot rule out the case that the value $\llbracket \alpha \rrbracket(w)(\emptyset)$ is in $(0, 1)$ and so, for uniformity's sake, we choose to allow $\llbracket m \rrbracket(w)(\emptyset) \in (0, 1)$ for atomic actions too.

Remark 3.1. The notion of actions prohibited with positive probability may seem problematic: what does it mean if our agent attempts to perform α when $0 < \llbracket \alpha \rrbracket(w)(\emptyset) < 1$? This is especially puzzling when α is atomic, but even in the case of composite actions, what happens when an agent attempts but fails to perform α ? For instance, suppose that one wants to open a combination lock to a safe and then grab the now-reachable contents. If `unlock` has a chance $0 \leq x \leq 1$ of opening the safe, then what should `unlock; take` mean?

This is not an issue specific to our probabilistic situation. The same question could be asked in the non-deterministic situation shown below.



In world w , it may be the case that doing m results in a world in which n is prohibited, so how should one interpret the action $m;n$ in this context? Either one interprets the composite by assuming that m results in w_2 or one gives an interpretation of doing n in world w_1 (where n is prohibited). The latter is easily done in typical computer science settings: interpret prohibition in terms of non-termination. This option is not suitable for our setting. But we suggest that, in either probabilistic or non-deterministic settings, one shouldn't expect sensible interpretations of prohibited actions. In some situations, actions simply cannot be performed and sometimes (as in the case of `unlock;take`), this is not obvious at the start. Our PDL models simply convey what worlds may result when various actions *are* performed and do not give information about worlds which result when a particular action *cannot* be executed⁹.

We have chosen a fuzzy logic to yield a truth-functional semantics for probabilistic PDL, but one must be careful here. One of the mantras of fuzzy logic is that fuzziness (i.e., vagueness) is not the same as probabilities. But the predicate “ m reliably realizes φ in w ” is a fuzzy predicate and we will define the truth value of this predicate in terms of the probability that doing m in w results in a world w' satisfying φ .

Thus, we introduce for each action α a dynamic operator¹⁰ $\langle\alpha\rangle$ and interpret $\langle\alpha\rangle\varphi$ as the fuzzy predicate “ α reliably realizes φ in w .” Explicitly

$$\llbracket\langle\alpha\rangle\varphi\rrbracket(w) = \sum_{w' \in \mathcal{W}} \llbracket\alpha\rrbracket(w)(w') \cdot \llbracket\varphi\rrbracket(w').$$

⁹ This response seems more plausible when we do not allow atomic actions to have non-trivial prohibition values and we are happy with this solution. We do not restrict our logic so that $\llbracket\alpha\rrbracket(w)(\emptyset) \in \{0, 1\}$, however, since this is fundamentally a philosophical choice that affects our logical properties very little.

¹⁰ It has been suggested that the notation $\langle\alpha\rangle$ is misleading here, confusing the conventional non-deterministic operator and our fuzzy operator. In overloading this notation we are following the conventions of (Kozen, 1983).

We also claim that $\langle\alpha\rangle\varphi$ is an appropriate interpretation of the statement “ α is a *sufficient means* to φ ” in this probabilistic setting. (We do not distinguish weakly and strongly sufficient means in this model.) The truth degree of $\langle\alpha\rangle\varphi$ in w is the *efficacy* of α as a means to φ in w .

The dynamic operator will be the basis for our fuzzy version of PDL, in which the truth degrees of certain fuzzy predicates correspond to probabilities involving the uncertain outcomes of actions. Accordingly, our interpretation of fuzzy PDL will take formulas to *fuzzy* sets of worlds, where the value $\llbracket\varphi\rrbracket(w)$ is the degree to which φ is true in w .

A *fuzzy PDL frame*, then, is a set \mathcal{W} of worlds together with a dynamic interpretation

$$\llbracket-\rrbracket : \Pi_0 \rightarrow (\mathcal{D}\mathcal{W})^{\mathcal{W}}.$$

A *fuzzy PDL model* extends a fuzzy PDL frame by adding a fuzzy valuation $\llbracket-\rrbracket : \Phi_0 \rightarrow \mathcal{F}\mathcal{W}$, where $\mathcal{F}X$ is the set of fuzzy subsets of X , i.e.,

$$\mathcal{F}X = \{p : X \rightarrow [0, 1]\}.$$

We interpret $\llbracket P \rrbracket(w)$ as the truth degree of P in world w .

Call the valuation *crisp* if for each atomic proposition P the set $\llbracket P \rrbracket$ is crisp and call the interpretation of Π_0 *crisp* if each set $\llbracket m \rrbracket(w)$ is crisp (and consequently either a singleton or empty, so that our model is deterministic). There is conceptual interest in focusing on crisp valuations: one may wish to focus on fuzziness that arises from uncertain outcomes and ignore any other sources of fuzziness in means-end relations. Nonetheless, we won’t restrict to crisp valuations since it adds little to the logical properties of our theory and allowing fuzzy atomic propositions more closely reflects the kinds of ends persons adopt (see Example 3.2 below, for instance).

There are many families of fuzzy propositional logic. We must decide on which of these families are best suited for means-end semantics involving probabilities. Our primary aim here is a conceptual analysis of means-end relations and so we will justify our semantics for the boolean connectives in terms of suitability for means-end semantics rather than attractive logical properties.

Let us first decide the appropriate semantics for conjunctions. For this, we are primarily motivated by the analysis of formulas of the form $\langle\alpha\rangle\varphi \wedge \langle\beta\rangle\psi$. What truth degree ought to be assigned to such formulas? The conjunction expresses the reliability of α in realizing φ and of β in realizing ψ . We find it natural to say that, in a given world w ,

$$\llbracket \langle\alpha\rangle\varphi \wedge \langle\beta\rangle\psi \rrbracket(w) = \min\{\llbracket \langle\alpha\rangle\varphi \rrbracket(w), \llbracket \langle\beta\rangle\psi \rrbracket(w)\}.$$

We adopt this equation for arbitrary conjunctions as well, of course. It is here that the technical details of our semantics diverges from that of (Kozen, 1983), due to our interpretation of efficacy in terms of reliably realizing one's end.

Example 3.2. A fickle gift-giver is shopping for a sweater, which he may give as a present to his brother later that day or may keep for himself, as the mood strikes him. If he chooses to give the sweater as a gift, then he wants his brother to appreciate the gift. If he chooses to keep the sweater, then he wants it to be a sweater of his liking. Thus, his end is the conjunction

$$\langle \text{give} \rangle \mathbf{BroHappy} \wedge \langle \text{keep} \rangle \mathbf{MeHappy}.$$

The means under consideration are actions buy_x where x is one of the sweaters at the store. In judging the reliability of each means realizing his end, he must determine the truth degree of the conjunction in the various worlds which result from doing buy_x .

In order to reliably reach his ultimate end, whether it is **BroHappy** or **MeHappy**, he must be sure that the corresponding action is a reliable means to that end. If he is unsure now *which* ultimate end he will pursue, then he had better make sure that **give** is a reliable means to **BroHappy** and **keep** is a reliable means to **MeHappy**. Thus, the minimum of the two is the truth degree relevant for our fickle giver.

Put differently: the truth degree of the formula

$$\langle \text{buy}_x \rangle (\langle \text{give} \rangle \mathbf{BroHappy} \wedge \langle \text{keep} \rangle \mathbf{MeHappy})$$

expresses the reliability of buy_x in realizing the end $\langle \text{give} \rangle \mathbf{BroHappy} \wedge \langle \text{keep} \rangle \mathbf{MeHappy}$. It does not express the desirability of performing buy_x . Efficacy is about reliability in realizing ends and this is different than desirability of the outcome. Insofar as the agent desires a reliable means to the conjunction, he should choose a means which gives each of the conjuncts a high truth degree.

As an anonymous referee noted, our work would be greatly simplified if we did not include complex ends like $\langle \alpha \rangle P \wedge \langle \beta \rangle Q$. In that case, we could restrict the dynamic operators to apply to formulas of the form $\langle \alpha_1 \rangle \dots \langle \alpha_k \rangle \varphi$ where φ is purely propositional and so we would not need the full fuzzy logic interpretations we give here. Indeed, in this case, we could have used qualitative decision theory and calculated efficacy under the assumption that φ worlds have value 1, non- φ worlds value 0 and determining the value of doing the composite $\alpha_1; \dots; \alpha_k$. We have resisted that alternative on the grounds that complex ends are often relative to both practical agents and artifactual functions and in

such cases, it is not clear how to use qualitative decision theory (which worlds are $\langle\alpha\rangle P \wedge \langle\beta\rangle Q$ worlds?). We offer a second example of the former here to partially justify the claim that such ends are relevant and to give further motivation to the interpretation we have selected.

Example 3.3. Consider a situation in which there are two cooperative but non-communicating actors, J and L . Suppose that J knows that L will try either to realize P by doing m_L or to realize Q by doing n_L , but he does not know which alternative L will choose. J wishes to act so that L is likely to succeed whichever alternative she chooses. In other words, he desires a means that is reliable in realizing the conjunction

$$\langle m_L \rangle P \wedge \langle n_L \rangle Q,$$

i.e., so that m_L is a reliable means to P and n_L is a reliable means to Q . It seems plausible to calculate the truth degree of the conjunction as the minimum of the truth degrees of the conjuncts, since this is the degree to which J can expect that L will realize her end.

On the other hand, if J knows that L is ambivalent regarding her options and will choose whichever will more likely succeed, then he is interested in a different end. He will be interested in maximizing $\max\{\langle m_L \rangle P, \langle n_L \rangle Q\}$. This is the natural disjunction of $\langle m_L \rangle P$ and $\langle n_L \rangle Q$, given that we define conjunction as the minimum. Thus, in this case, J desires a reliable means to

$$\langle m_L \rangle P \vee \langle n_L \rangle Q$$

and our semantics assigns the appropriate value to this formula in each world.

We define disjunction and implication in terms of the conjunction, so that

$$\begin{aligned} \llbracket \varphi \vee \psi \rrbracket(w) &= \max\{\llbracket \varphi \rrbracket(w), \llbracket \psi \rrbracket(w)\} \\ \llbracket \varphi \rightarrow \psi \rrbracket(w) &= \begin{cases} 1 & \text{if } \llbracket \varphi \rrbracket(w) \leq \llbracket \psi \rrbracket(w) \\ \llbracket \psi \rrbracket(w) & \text{else} \end{cases} \end{aligned}$$

These are the so-called Gödel interpretations for \wedge , \vee and \rightarrow .

Let us turn our attention to negation, then. A formula $\neg\langle\alpha\rangle\varphi$ expresses that α does not reliably realize φ . Such formulas occur naturally as ends, in particular in games in which one tries to prevent his opponent from winning. In such situations, one's end may be that a particular move by his opponent does *not* reliably result in a winning state for his opponent.

There are two obvious alternatives for negation. The first is the *Gödel negation*, defined by

$$\llbracket \neg\varphi \rrbracket(w) = \begin{cases} 1 & \text{if } \llbracket \varphi \rrbracket(w) = 0 \\ 0 & \text{else} \end{cases}$$

The other alternative is sometimes called *Lukasiewicz negation*, defined by


$$\llbracket \neg\varphi \rrbracket(w) = 1 - \llbracket \varphi \rrbracket(w).$$

Gödel negation yields a better behaved logic, since it satisfies

$$\llbracket \neg\varphi \rrbracket(w) = \llbracket \varphi \rightarrow \perp \rrbracket(w),$$

so it is with reluctance that we abandon the Gödel negation. Nonetheless, we do not think that the Gödel negation is the right negation for our purposes, since it results in extreme truth values for the fuzzy proposition “ m does not reliably realize φ .”

Example 3.4. An eight-ball player a considers placing his ball so that it partially blocks his opponent b 's winning shot. Let us represent his end as $\neg(\text{shoot}_b)\mathbf{Win}$ and the means he is considering as block_a . Let us suppose that the blocking shot is easy to make but still leaves small positive probability x that his opponent sinks the eight ball. The transition system representing this situation is shown below.



If we interpret negation as Gödel negation, then the efficacy of block_a in achieving a 's end is 0, whereas if we use the Lukasiewicz negation, the same efficacy evaluates to $1 - x$. We find the latter more reasonable in evaluating the appropriateness of block_a .

From another perspective, our choice of boolean operators is quite natural. In classical PDL, $\llbracket \varphi \rrbracket$ denotes the crisp set $\{w \in \mathcal{W} \mid w \models \varphi\}$. In our fuzzy PDL, the same construction denotes the corresponding fuzzy set, so that $\llbracket \varphi \rrbracket(w)$ is the degree to which w is an element of the set of worlds satisfying φ (i.e., we can view the relation \models as a fuzzy relation). In these terms, the boolean connectives \wedge , \vee and \neg are

interpreted as

$$\begin{aligned} \llbracket \varphi \wedge \psi \rrbracket &= \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket \\ \llbracket \varphi \vee \psi \rrbracket &= \llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket \\ \llbracket \neg \varphi \rrbracket &= \mathcal{W} \setminus \llbracket \varphi \rrbracket \end{aligned}$$

where \cap , \cup and \setminus are the *standard operations* on fuzzy sets (Klir and Yuan, 1995). The interpretation of \rightarrow is the R-implication for the standard operations. Our fuzzy PDL, then, fits in with the “broad sense” of fuzzy logic, rather than the “narrow sense” developed by Hájek, et al.

Nonetheless, we hope to revisit the interpretations of the boolean connectives at a later date, when we may more closely examine the applicability of the LII logic found in (Hájek et al., 1995). Perhaps a variation of this logic can satisfy both our goals: semantic plausibility and logical reasonableness.

Remark 3.5. Under our interpretation, the conjugate $[\alpha]$ defined by $\neg\langle\alpha\rangle\neg$ of the dynamic operator $\langle\alpha\rangle$ is easy to understand:

$$\llbracket [\alpha]\varphi \rrbracket(w) = \llbracket \langle\alpha\rangle\varphi \rrbracket(w) + \llbracket [\alpha] \rrbracket(w)(\emptyset).$$

That is, the truth degree of the conjugate is the truth degree of $\langle\alpha\rangle\varphi$ plus the probability that α is prohibited (the same relation can be found in probabilistic PDL (Kozen, 1983)). Thus, the conjugate bears some resemblance to the strong dynamic operator from Section 2 (hence our notation), since $[\alpha]\varphi$ is trivially satisfied in a world in which α is prohibited with probability 1. It is unusual in that $\llbracket [\alpha]\varphi \rrbracket(w) \geq \llbracket \langle\alpha\rangle\varphi \rrbracket(w)$, which may be unexpected in terms of the usual strong/weak distinction of modal operators.

Of course, we must also extend the dynamic interpretation

$$\llbracket - \rrbracket : \Pi_0 \rightarrow (\mathcal{D}\mathcal{W})^{\mathcal{W}}$$

to arbitrary actions. For this, we must omit non-deterministic choice and iteration until the following section, when we add non-determinism to our logic. We give the extension for the remaining actions along with a summary of our semantics for formulas in Table IV.

Remark 3.6. We have defined an interpretation $\llbracket - \rrbracket : \Pi \rightarrow (\mathcal{D}\mathcal{W})^{\mathcal{W}}$. Thus, for each α , we have $\llbracket [\alpha] \rrbracket : \mathcal{W} \rightarrow \mathcal{D}\mathcal{W}$. Since \mathcal{D} is a subfunctor of the fuzzy subset functor \mathcal{F} , this means that the relation $\xrightarrow{\alpha}$ is a special kind of fuzzy relation: a subset of $\mathcal{F}(\mathcal{W} \times \mathcal{W})$ such that for all w ,

$$\sum_{w' \in \mathcal{W}} w \xrightarrow{\alpha} w' \leq 1.$$

Table IV. A fuzzy semantics for PDL

<u>On formulas</u>	
$\llbracket \top \rrbracket(w)$	$= 1$
$\llbracket \varphi \wedge \psi \rrbracket(w)$	$= \min\{\llbracket \varphi \rrbracket(w), \llbracket \psi \rrbracket(w)\}$
$\llbracket \varphi \vee \psi \rrbracket(w)$	$= \max\{\llbracket \varphi \rrbracket(w), \llbracket \psi \rrbracket(w)\}$
$\llbracket \neg \varphi \rrbracket(w)$	$= 1 - \llbracket \varphi \rrbracket(w)$
$\llbracket \varphi \rightarrow \psi \rrbracket(w)$	$= \begin{cases} 1 & \text{if } \llbracket \varphi \rrbracket(w) \leq \llbracket \psi \rrbracket(w), \\ \llbracket \psi \rrbracket(w) & \text{else;} \end{cases}$
$\llbracket \langle \alpha \rangle \varphi \rrbracket(w)$	$= \sum_{w' \in \mathcal{W}} \llbracket \alpha \rrbracket(w)(w') \cdot \llbracket \varphi \rrbracket(w')$
<u>On actions</u>	
$\llbracket \alpha; \beta \rrbracket(w)(w')$	$= \sum_{w'' \in \mathcal{W}} \llbracket \alpha \rrbracket(w)(w'') \cdot \llbracket \beta \rrbracket(w'')(w')$
$\llbracket \varphi? \rrbracket(w)(w')$	$= \begin{cases} \llbracket \varphi \rrbracket(w) & \text{if } w = w'; \\ 0 & \text{else.} \end{cases}$

In these terms, we may interpret $\llbracket \langle \alpha \rangle \varphi \rrbracket(w)$ as the fuzzy predicate

$$\bigvee_{w' \in \mathcal{W}} (w \xrightarrow{\alpha} w' \wedge w' \in \llbracket \varphi \rrbracket),$$

where \bigvee and \wedge are interpreted as in product logic (see (Hájek et al., 1995; Hájek, 1998)) rather than the “standard” connectives. This perspective gives more reason to believe that LII is an appropriate alternative setting for our means-end logic.

Example 3.7. We revisit the footrace from Example 2.1, exploiting the extra flexibility of the fuzzy set semantics. Previously, we included transitions $w_1 \rightarrow w_1$ and $w_3 \rightarrow w_3$ to represent the possibility that the gun misfires. But one might expect this occurrence to be less probable as a successful firing of the gun. We’ve assigned the probability of misfire 0.05, the probability of successfully firing 0.95 and every other relevant probability 1. We do not assume that a misfire increases the probability of subsequent misfires.

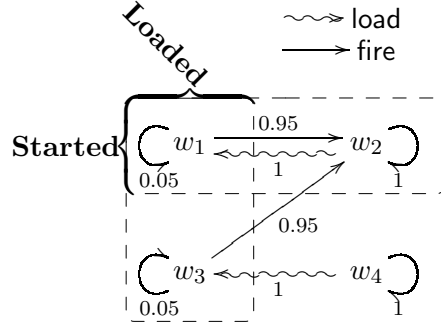


Figure 2. A probabilistic version of the PDL model from Figure 1.

In Figure 2, we present the fuzzy version of our footrace example, together with the interpretations of those formulas from Figure 1 (omitting those involving iteration and choice).

Table V. Truth degrees of various propositions about the model in Figure 2.

	w_1	w_2	w_3	w_4
$\llbracket \text{Started} \rrbracket$	1	1	0	0
$\llbracket \text{Loaded} \rrbracket$	0	1	0	1
$\llbracket [\text{fire}] \text{Started} \rrbracket$	1	1	0.95	0
$\llbracket \langle \text{fire} \rangle \text{Started} \rrbracket$	1	1	0.95	0
$\llbracket [\text{load}] \text{Loaded} \rrbracket$	1	1	1	1
$\llbracket \langle \text{load} \rangle \text{Loaded} \rrbracket$	0	1	0	1
$\llbracket [\text{Loaded?}; \text{fire}] \text{Started} \rrbracket$	1	1	0.95	1
$\llbracket \langle \text{Loaded?}; \text{fire} \rangle \text{Started} \rrbracket$	1	0	0.95	0
$\llbracket [\text{load}; \text{fire}] \text{Started} \rrbracket$	1	1	1	0.95
$\llbracket \langle \text{load}; \text{fire} \rangle \text{Started} \rrbracket$	0	1	0	0.95

We say that a formula φ is *valid* if, for every fuzzy PDL model and world w , we have $\llbracket \varphi \rrbracket(w) = 1$. We say that a rule of inference $\varphi_1, \dots, \varphi_n / \psi$ is *sound* if for every model

$$\min \{ \inf \{ \llbracket \varphi_i \rrbracket(w) \mid w \in \mathcal{W} \} \mid 1 \leq i \leq n \} \leq \inf \{ \llbracket \psi \rrbracket(w) \mid w \in \mathcal{W} \}.$$

Reasoning in terms of expected outcomes is rather different than quantifying over all possible outcomes, as one does in classical PDL. Consequently, we lose some of the classical axioms for PDL. We summarize the logical features of fuzzy PDL in Table VI, which refers to the De Morgan and Implication axioms found in Table VII. We omit the proofs for the positive results, which are straightforward, but we include a counterexample for distributivity below. We also omit completeness considerations for this logic, due to existing negative results


(see (Pelletier, 2004)). Because conceptual analysis is our primary aim, we do not regard the lack of a complete calculus to be a fatal defect for our logic.

Table VI. Summary of the logical features of fuzzy PDL

<u>Positive results</u>	<u>Negative results</u>
Axioms De Morgan and Implication axioms (Table VII) are valid. Composition axiom (Table III) is valid.	Axioms K, Distributivity and Test (Table III) are not valid.
Rules Modus ponens, cut and necessitation (Table VII) are sound.	

It is easy to see that the Test axiom is not valid, but this is not surprising. As long as we interpret the dynamic operators in terms of expected outcomes (i.e., weighted averages), we cannot expect the conditional operator to act as implication.

The fuzzy PDL semantics also does not satisfy the distributivity equivalence $\langle \alpha \rangle (\varphi \vee \psi) \leftrightarrow (\langle \alpha \rangle \varphi \vee \langle \alpha \rangle \psi)$. For instance, consider the following situation.



In this case $\llbracket \langle m \rangle (P \vee Q) \rrbracket (w) = 1$ but $\llbracket \langle m \rangle P \vee \langle m \rangle Q \rrbracket (w) = \frac{1}{2}$.

But the failure to accommodate distributivity is in keeping with our interpretations of $\langle \alpha \rangle$ in terms of reliability. Let us consider the conjunction $\langle \alpha \rangle \varphi \vee \langle \alpha \rangle \psi$. This asserts two local means-end relations: that α reliably realizes φ or that α reliably realizes ψ . But this is different than the degree to which α is a reliable means to $\varphi \vee \psi$, which is our interpretation of $\langle \alpha \rangle (\varphi \vee \psi)$. In other words, the distributive axiom simply doesn't reflect the intended interpretation of our fuzzy set semantics and so its failure is an expected feature rather than a shortcoming of our interpretation.

Table VII. de Morgan algebra and implication axioms

de Morgan axioms	
Involution	$\neg\neg\varphi \leftrightarrow \varphi$
Commutativity	$\varphi \wedge \psi \leftrightarrow \psi \wedge \varphi$ $\varphi \vee \psi \leftrightarrow \psi \vee \varphi$
Associativity	$\varphi \wedge (\psi \wedge \vartheta) \leftrightarrow (\varphi \wedge \psi) \wedge \vartheta$ $\varphi \vee (\psi \vee \vartheta) \leftrightarrow (\varphi \vee \psi) \vee \vartheta$
Distributivity	$\varphi \wedge (\psi \vee \vartheta) \leftrightarrow (\varphi \wedge \psi) \vee (\varphi \wedge \vartheta)$ $\varphi \vee (\psi \wedge \vartheta) \leftrightarrow (\varphi \vee \psi) \wedge (\varphi \vee \vartheta)$
Idempotence	$\varphi \wedge \varphi \leftrightarrow \varphi$ $\varphi \vee \varphi \leftrightarrow \varphi$
Absorption	$\varphi \wedge (\varphi \vee \psi) \leftrightarrow \varphi$ $\varphi \vee (\varphi \wedge \psi) \leftrightarrow \varphi$
Absorption by \perp and \top	$\varphi \wedge \perp \leftrightarrow \perp$ $\varphi \vee \top \leftrightarrow \top$
Identity	$\varphi \wedge \top \leftrightarrow \varphi$ $\varphi \vee \perp \leftrightarrow \varphi$
de Morgan's laws	$\neg(\varphi \wedge \psi) \leftrightarrow \neg\varphi \vee \neg\psi$ $\neg(\varphi \vee \psi) \leftrightarrow \neg\varphi \wedge \neg\psi$

Implication axioms

Reflexivity	$\varphi \rightarrow \varphi$
Weakening	$(\varphi \rightarrow \psi) \rightarrow ((\varphi \wedge \vartheta) \rightarrow \psi)$ $(\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow (\psi \vee \vartheta))$
Currying	$((\varphi \wedge \psi) \rightarrow \vartheta) \leftrightarrow (\varphi \rightarrow (\psi \rightarrow \vartheta))$
Preserves \wedge	$((\varphi \rightarrow \psi) \wedge (\varphi \rightarrow \vartheta)) \leftrightarrow (\varphi \rightarrow (\psi \wedge \vartheta))$
Preserves \top	$\varphi \rightarrow \top$

Inference rules

Modus Ponens	$\varphi, \varphi \rightarrow \psi / \psi$
Cut	$\varphi \rightarrow \psi, \psi \rightarrow \vartheta / \varphi \rightarrow \vartheta$
Necessitation	$\varphi / [\alpha]\varphi$

In the next section, we introduce more complex models that allow for a representation of non-deterministic choice and iteration.

4. Minimal fuzzy models

Non-determinism is a natural feature of PDL and of means-end relations, but it is missing from our fuzzy models from the previous section. Probabilities are useful for introducing efficacy, but by replacing the powerset operator \mathcal{P} by the subdistribution operator \mathcal{D} , we lost some of the expressiveness of PDL. In particular, there is no natural way to interpret non-deterministic choice using the models from Section 3. We don't want to assign probabilities to the outcome of a non-deterministic choice, since this would be a fundamental re-interpretation—it would replace the non-deterministic feature with a related probabilistic choice¹¹. But probabilistic choices aren't non-deterministic in the usual sense.

We fix this problem by turning to a kind of neighborhood semantics. A *minimal fuzzy PDL frame* for Π_0 is a set \mathcal{W} of worlds, together with a dynamic interpretation¹²

$$\llbracket - \rrbracket : \Pi_0 \rightarrow (\mathcal{PDW})^{\mathcal{W}}.$$

This interpretation takes an atomic action m and world w to a set of discrete subdistributions $\llbracket m \rrbracket(w)$ on \mathcal{W} . The elements of $\llbracket m \rrbracket(w)$ represent alternatives which may be selected non-deterministically. A *minimal fuzzy PDL model* adds a valuation $\llbracket - \rrbracket : \Phi_0 \rightarrow \mathcal{FW}$ to a minimal fuzzy PDL frame.

Our dynamic structure mixes fuzzy and non-deterministic features. Thus, we define our semantics by mixing the non-deterministic semantics of Section 2 with the fuzzy set semantics of the previous section.

For technical reasons, we require that our set \mathcal{W} of worlds is finite, so that all of the sums that appear here are finite sums and hence commute with limits. In particular, this means that for every family

¹¹ In (Kozen, 1983), we see exactly this solution: the author replaces the non-deterministic choice construction in favor of taking linear combinations of subformulas.

¹² One could also choose a fuzzier version of the dynamic interpretations, namely

$$\llbracket - \rrbracket : \Pi_0 \rightarrow (\mathcal{FDW})^{\mathcal{W}}.$$

This more complicated semantics would distract from our main aim, discussing the role of probabilities in means-end relations, and is unnecessary for re-introducing the action constructions \cup and $*$.

$(S_w)_{w \in \mathcal{W}}$ of sets, we have

$$\sum_{w \in \mathcal{W}} \inf S_w = \inf \left\{ \sum_{w \in \mathcal{W}} \sigma_w \mid \sigma \in \prod_{w \in \mathcal{W}} S_w \right\}. \quad (1)$$

Here, σ is a \mathcal{W} -indexed vector $\sigma = \langle \sigma_{w_1}, \sigma_{w_2}, \dots, \sigma_{w_n} \rangle$ and σ_w is the w 'th projection of σ . We will be particularly interested in the product $\prod_{w \in \mathcal{W}} \llbracket \alpha \rrbracket(w)$. In this case, elements σ of this product are vectors of subdistributions, so that for each w , $\sigma_w \in \llbracket \alpha \rrbracket(w) \subseteq \mathcal{D}(\mathcal{W})$.

In the sequel, we use p, q, \dots to range over subdistributions on \mathcal{W} and write $p(\emptyset)$ for $1 - \sum_{w \in \mathcal{W}} p(w)$. At a given world w , the interpretation $\llbracket \alpha \rrbracket(w)$ of α in w is a set of subdistributions. We interpret $p \in \llbracket \alpha \rrbracket(w)$ as follows: $p(w')$ as the probability that doing α in w yields w' *provided* that p is the subdistribution non-deterministically chosen. This serves as the motivation for our minimal fuzzy semantics presented in Table VIII.

The formula $\langle \alpha \rangle \varphi$ is interpreted as “one *can* reliably realize φ by doing α ,” where “can” here is the can of ability described in (Brown, 1988). Put differently, $\langle \alpha \rangle \varphi$ means, “There is a manner of doing α to reliably realize φ .” The conjugate formula $[\alpha] \varphi$ represents the vague proposition, “However α is done, doing α will reliably realize φ ,” while still allowing that α may be prohibited. These two operators can be used to define weakly and strongly sufficient means as in Section 2. The truth degree of $\langle \alpha \rangle \varphi$ ($[\alpha] \varphi \wedge \langle \alpha \rangle \top$, resp.) in w is the *efficacy* of α as a weakly sufficient (strongly sufficient, resp.) means to φ in w .

We use the standard fuzzy infinitary disjunction (conjunction, resp.) to assign the least upper bound (greatest lower bound, resp.) of the relevant expected values in interpreting $\langle \alpha \rangle \varphi$ ($[\alpha] \psi$, resp.):

$$\begin{aligned} \llbracket \langle \alpha \rangle \varphi \rrbracket(w) &= \sup \left\{ \sum_{w' \in \mathcal{W}} p(w') \cdot \llbracket \varphi \rrbracket(w') \mid p \in \llbracket \alpha \rrbracket(w) \right\}, \\ \llbracket [\alpha] \varphi \rrbracket(w) &= \inf \left\{ p(\emptyset) + \sum_{w' \in \mathcal{W}} p(w') \cdot \llbracket \varphi \rrbracket(w') \mid p \in \llbracket \alpha \rrbracket(w) \right\} \end{aligned}$$

Composition is a bit complicated, since we must compose each pair of sub-distributions found in the interpretations of α and β , like so:

$$\llbracket \alpha; \beta \rrbracket(w) = \left\{ \lambda w'. \sum_{w'' \in \mathcal{W}} p(w'') \cdot \sigma_{w''}(w') \mid p \in \llbracket \alpha \rrbracket(w), \sigma \in \prod_{v \in \mathcal{W}} \llbracket \beta \rrbracket(v) \right\}$$

Non-deterministic choice and iteration now can be defined exactly as in the standard PDL semantics, and the conditional operator is interpreted by a singleton consisting of its interpretation from Section 3. For iteration, we apply the convention that $\llbracket \alpha^0 \rrbracket(w) = \{w\}$, so that $\{w\} \in \llbracket \alpha^* \rrbracket(w)$.

Table VIII. A minimal fuzzy semantics for PDL

<u>On formulas</u>	
$\llbracket \top \rrbracket(w)$	$= 1$
$\llbracket \varphi \wedge \psi \rrbracket(w)$	$= \min\{\llbracket \varphi \rrbracket(w), \llbracket \psi \rrbracket(w)\}$
$\llbracket \varphi \vee \psi \rrbracket(w)$	$= \max\{\llbracket \varphi \rrbracket(w), \llbracket \psi \rrbracket(w)\}$
$\llbracket \neg \varphi \rrbracket(w)$	$= 1 - \llbracket \varphi \rrbracket(w)$
$\llbracket \varphi \rightarrow \psi \rrbracket(w)$	$= \begin{cases} 1 & \text{if } \llbracket \varphi \rrbracket(w) \leq \llbracket \psi \rrbracket(w), \\ \llbracket \psi \rrbracket(w) & \text{else} \end{cases}$
$\llbracket \langle \alpha \rangle \varphi \rrbracket(w)$	$= \sup \left\{ \sum_{w' \in \mathcal{W}} p(w') \cdot \llbracket \varphi \rrbracket(w') \mid p \in \llbracket \alpha \rrbracket(w) \right\}$
<u>On actions</u>	
$\llbracket \alpha; \beta \rrbracket(w)$	$= (\text{see text.})$
$\llbracket \alpha \cup \beta \rrbracket(w)$	$= (\llbracket \alpha \rrbracket \cup \llbracket \beta \rrbracket)(w)$
$\llbracket \alpha^* \rrbracket(w)$	$= \bigcup_{n=0}^{\infty} \underbrace{\llbracket \alpha; \dots; \alpha \rrbracket(w)}_{n \text{ times}}$
$\llbracket \varphi? \rrbracket(w)$	$= \{p_{\varphi}^w\}$ where $p_{\varphi}^w(w') = \begin{cases} \llbracket \varphi \rrbracket(w) & \text{if } w = w'; \\ 0 & \text{else.} \end{cases}$

Example 4.1. We illustrate the minimal fuzzy model semantics by returning to the footrace example from the previous section. Previously, we had to defer on examining formulas involving iteration and choice, since the normal fuzzy set semantics has no natural definition for these constructions. Minimal fuzzy models correct this deficiency and allow us to complete the table from Figure 2, as shown in Table IX.

Table IX. The completion of the table in Figure 2.

	w_1	w_2	w_3	w_4
$\llbracket [(\text{load} \cup \text{fire})^*] \text{Started} \rrbracket$	1	1	0	0
$\llbracket \langle (\text{load} \cup \text{fire})^* \rangle \text{Started} \rrbracket$	1	1	1	1

We summarize the primary logical features of minimal fuzzy PDL in Table X. We omit the proofs of the positive results, which are straightforward but tedious. Since the minimal fuzzy semantics reduces (in special cases) to the fuzzy semantics from the previous section, the negative results discussed there apply here too. Namely, our minimal fuzzy semantics does not satisfy the axioms Test, Distributivity or K.

Table X. Summary of the logical features of minimal fuzzy PDL

<u>Positive results</u>	<u>Negative results</u>
<p>Axioms De Morgan, Implication and Composition axioms are valid. Choice and iteration (Table III) are also valid.</p> <p>Rules Modus ponens, cut and necessitation are sound.</p>	<p>Axioms K, Distributivity and Test are not valid. Induction (Table III) is not valid.</p>

Of course, our motivation in this section isn't to fix the problems of Section 3. We introduced the more complicated semantics so that we could represent non-deterministic choice and iteration. Their defining axioms are valid for this class of models, but the Induction axiom is not valid. This is not too surprising, given that K is not a valid axiom. The model in Figure 3 is a counterexample to Induction.




Figure 3. A counterexample to Induction

The region interior to the dotted curve marks the area where $\llbracket P \rrbracket$ evaluates to at least $\frac{2}{3}$ (in particular, $\llbracket P \rrbracket(w) = \frac{2}{3}$), while in the interior of the solid curve $\llbracket P \rrbracket$ evaluates to 1 and $\llbracket P \rrbracket(w_3) = 0$. We are interested

in calculating

$$\llbracket [m^*](P \rightarrow [m]P) \rightarrow (P \rightarrow [m^*]P) \rrbracket(w).$$

This takes some work, but it is not particularly difficult to verify

$$\begin{aligned} \llbracket [m^*]P \rrbracket(w) &= \frac{1}{3} \\ \llbracket [m^*](P \rightarrow [m]P) \rrbracket(w) &= \frac{2}{3}. \end{aligned}$$

Consequently, the induction axiom evaluates to $\frac{1}{3}$ in world w .

5. Concluding remarks

Probabilities, whether interpreted subjectively or objectively, are an essential characteristic of means-end relations. An agent choosing between two different plans to realize a common end is likely to consider the propensity of each plan to successfully bring about the end. This may not be the deciding factor for an agent (he may choose a less efficacious means in order to avoid undesirable side effects, for instance) but it is surely an important factor. Additionally, a semantics for artifactual functions in terms of means-end relations is much richer if the relations include a measure of efficacy, measuring the degree to which an artifact fulfills its function.

Traditional propositional dynamic logic represents non-determinism well, but has no probabilistic features at all. Languages like pCTL, on the other hand, are designed for describing probabilistic transition systems, but do so by introducing indexed operators that explicitly include the probabilities. Probabilistic PDL comes closer to our needs, but replaces propositional atoms with uninterpreted measurable functions. Decision-theoretic approaches naturally include measures of efficacy, but do so by presupposing desirability rankings that mix efficacy and desirability together. While this is appropriate for generating practical plans of action, it seems difficult to separate the resulting notion of efficacy from desirability and it is not altogether clear how to calculate efficacy of means to complex ends like $\langle \alpha \rangle P \wedge \langle \beta \rangle Q$.

Because we are interested in a truth-functional language, we chose to introduce a fuzzy logic interpretation of PDL, in which certain probabilistic formulas are translated to fuzzy predicates (as in (Hájek et al., 1995)). Thus, we use the probabilities in order to create vagueness and vagueness is exactly the phenomenon that fuzzy logics explore.

At present, we have included an interpretation of the boolean connectives that stresses intuitively plausible semantics at the expense of

nice logical properties. This is not satisfying, but we are hopeful that later work will allow a better behaved logic. We view the interpretation of the boolean connectives as a parameter to our model in any case and stress instead the new work involving the interpretation of dynamic operators as fuzzy predicates. This is the original contribution to probabilities and logic that we offer.

Aside from the boolean connectives, there are many ways to extend this cursory presentation of efficacy in means-end relations. In particular, we employ our definition of efficacy to functional ascriptions in (Hughes and Zwart, 2005) in order to compare the appropriateness of various artifacts in realizing their functions and in defining one important kind of malfunction.

Acknowledgements

We would like to thank the anonymous referees for their helpful comments and references to existing literature. Additionally, our development and presentation has been aided by discussions with Ling Cheung, Kit Fine, Sven Ove Hansson, Lamber Royakkers, Yde Venema and Sjoerd Zwart.

References

- Bianco, A. and L. de Alfaro: 1995, ‘Model Checking of Probabilistic and Nondeterministic Systems’. *FSTTCS: Foundations of Software Technology and Theoretical Computer Science* **15**.
- Boutillier, C.: 1994, ‘Toward a logic for qualitative decision theory’. In: *Proceedings of the 4th International Conference on Knowledge Representation and Reasoning, KR’94*. pp. 75–86.
- Brown, M. A.: 1988, ‘On the Logic of Ability’. *Journal of Philosophical Logic* **17**, 1–26.
- Brown, M. A.: 2005, ‘Means and Ends in Branching Time’. Presented at the Norms, Reasoning and Knowledge in Technology workshop.
- Castilho, M., A. Herzig, and I. Varzinczak: 2002, ‘It depends on the context! A decidable logic of actions and plans based on a ternary dependence relation’. In: *9th Intl. Workshop on Non-Monotonic Reasoning NMR’2002*.
- Castilho, M. A., O. Gasquet, and A. Herzig: 1999, ‘Formalizing action and change in modal logic I: the frame problem’. *Journal of Logic and Computation* **9**(5).
- Hájek, P.: 1998, *Metamathematics of Fuzzy Logic*, Vol. 4 of *Trends in Logic*. Kluwer Academic Publishers.
- Hájek, P., L. Godo, and F. Esteva: 1995, ‘Fuzzy Logic and Probability’. In: *Proceedings of the Eleventh Annual Conference on Uncertainty in Artificial Intelligence (UAI-95)*. Montreal, Quebec, Canada.

- Hájek, P., L. Godo, and F. Esteva: 2000, 'Reasoning about Probability using Fuzzy Logic'. *Neural Network World* **10**, 811–824.
- Hanks, S. and D. McDermott: 1987, 'Default Reasoning, Nonmonotonic Logics, and the Frame Problem'. In: M. L. Ginsberg (ed.): *Readings in Nonmonotonic Reasoning*. Los Altos, CA: Kaufmann, pp. 390–395.
- Hansson, H. and B. Jonsson: 1994, 'A Logic for Reasoning About Time and Reliability'. *Formal Aspects of Computing* **6**(5), 512–535.
- Harel, D.: 1984, 'Dynamic Logic'. In: D. Gabbay and F. Guenther (eds.): *Handbook of Philosophical Logic*, Vol. II. D. Reidel Publishing Company, pp. 497–604.
- Hughes, J.: 2005, 'Means-end relations and Artifacts Functions: a Sketch'. Presented at the Norms, Reasoning and Knowledge in Technology workshop, available at <http://phiwumbda.org/~jesse/papers/index.html>.
- Hughes, J., P. Kroes, and S. Zwart: 2005, 'A Semantics for Means-End Relations'. Presented at SEP 2005, available at <http://phiwumbda.org/~jesse/papers/index.html>.
- Hughes, J. and S. Zwart: 2005, 'A Semantics for Functional Efficacy'. Presented at SPT 2005, available at <http://phiwumbda.org/~jesse/papers/index.html>.
- Klir, G. J. and B. Yuan: 1995, *Fuzzy Sets and Fuzzy Logic: Theory and Applications*. Prentice Hall PTR.
- Kozen, D.: 1983, 'A Probabilistic PDL'. In: *STOC '83: Proceedings of the fifteenth annual ACM symposium on Theory of computing*. New York, NY, USA, pp. 291–297.
- Meyer, J.-J. C.: 2000, 'Dynamic logic for reasoning about actions and agents'. In: *Logic-based artificial intelligence*. Norwell, MA, USA: Kluwer Academic Publishers, pp. 281–311.
- Pelletier, F. J.: 2004, 'On Some Alleged Misconceptions About Fuzzy Logic'. *Artificial Intelligence Review* **22**(1), 71–82.
- Pollock, J. L.: 2002, 'The Logical Foundations of Means-End Reasoning'. In: R. Elio (ed.): *Common Sense, Reasoning and Rationality*. Oxford University Press.
- Prendinger, H. and G. Schurz: 1996, 'Reasoning about Action and Change. A Dynamic Logic Approach.'. *Journal of Logic, Language and Information* **5**(2), 209–245.
- Puterman, M. L.: 1994, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc.
- von Wright, G. H.: 1963, 'Practical Inference'. *The Philosophical Review* **72**(2), 159–179.
- Zhang, D. and N. Foo: 2002, 'Dealing with the ramification problem in the extended propositional dynamic logic'. In: F. Wolter, H. Wansing, M. de Rijke, and M. Zakharyashev (eds.): *Advances in Modal Logic*, Vol. 3. World Scientific, pp. 173–191.

