A Semantics for Means-End Relations

Jesse Hughes   Peter Kroes   Sjoerd Zwart

January 15, 2005

Abstract

There has been relatively much work on practical reasoning in philosophy and artificial intelligence. Typically, such reasoning includes premises regarding means-end relations, but the semantics of such relations is not clear. In this paper, we provide a formal semantics for means-end relations, in particular for necessary and sufficient means-end relations. Our semantics includes a non-monotonic conditional operator, so that related practical reasoning is naturally defeasible. This semantics lays the groundwork for an evaluation of existing theories of practical reasoning as well as a semantic foundation for new work in practical reasoning.

"They were in conversation without speaking. They didn't need to speak. They just changed reality so that they had spoken." Terry Pratchett, *Reaper Man*

1 Introduction

Use of means-end relations and means-end reasoning are integral aspects of linguistic practices in general and engineering linguistic practice in particular. The development of a formal semantics for means-end relations and a formal logic for means-end reasoning may not only contribute to a better understanding of these aspects of linguistic practices, but also to attempts to formally represent functional properties of (technical) objects in engineering data representation systems (such as CAD-systems) [5]. The attribution of functional properties to (technical) objects is intimately related to the use of these objects as means to reach certain ends (cfr. the use-plan approach of [8]).

There is much literature on practical reasoning, including von Wright's important work in [13] and Pollock's means-end reasoning in [12], however there is relatively little work on the *meaning* of the premises in an argument in practical reasoning. We believe that one should be clear on the semantics of means-end relations before presenting a deductive theory for arguments involving those relations (and ending in yet another kind of proposition, such as one describing an intention to act). It is not that we expect our semantic analysis to refute current theories of practical reasoning, but rather that we are uncertain how to evaluate those theories without first being clear on the semantics of the premises that occur in arguments in practical reasoning. In this paper we will present a proposal for how a formal semantics for means-end relations might be developed, so that thereafter we may more confidently judge existing theories of practical reasoning or develop our own such theory. Because of the semantic nature of our work, it is more closely related to Brown's logic of ability [1] or Horty and Belnap's deliberative stit [7] than it is to traditional practical reasoning.

The basic idea underlying our formal semantics is that means are actions to bring about desired states of affairs in the world. In other words, means transform a given world into another one that the agent desires or prefers to the original one. From an epistemological point of view, one of the important elements involved in means-end ascriptions is knowledge of causal relations (deterministic and/or probabilistic). If event a causes event b, then under certain conditions the occurrence of a may be considered to be a (sufficient) means to achieve b; one of these conditions is that it must be (technically) possible to bring about a [13]. What other elements are necessary to justify a means-end ascription, if any, will not be further discussed here. In our formal semantics we simply assume that knowledge of the transformation that a means brings about in a world is given. We are interested in a clarification of the meaning of means-end relations with the help of a formal semantics, not in justification of means-end ascriptions.

Means-end talk in natural language exhibits many different features and of course our aim is to represent these features as faithfully as possible in our formal system. Let us list a number of features of the phenomenology of means-end talk (without pretending to be complete):

(1) the distinction between means and ends

(2) the fact that what is an end in one context, may be a means in another

(3) the fact that in certain cases the distinction between means and ends collapses (ends in themselves)

(4) the distinction between necessary and sufficient means

(5) the fact that means-end conditionals are non-monotonic (leading to the so-called frame problem of practical reasoning)

(6) the fact that entities of different types may constitute means (objects, actions)

(7) the distinction between effective and efficient means

(8) the distinction between good and bad means

These features, however, are to be taken as prima facie features of means-end talk. Our formal reconstruction may show that some of these features have to be reconsidered due to the vagueness and sloppiness of natural language. For instance, from a formal point of view there is clearly a tension between features (1) and (3): if we represent means and ends as objects of different types, then it will be difficult to account for feature (3). In other words, we are involved in a critical formal reconstruction of means-end talk in natural language.

We do not pretend that our system can account for all the above features. We will discuss how our formal means-end semantics accounts for at least the first five features (whether or not in a reformulated form), but some of our discussion will be unfortunately brief. In Section 6, we sketch how to deal with items (3), (6) and (7). Our treatment of items (1), (2), (4) and (5) are an integral part of our proposed semantics and are described in some detail throughout our presentation. We leave the account for the other features for future work.

Finally, our proposal is primarily one for a semantics of means-end relations, not for means-end reasoning. Apart from showing its adequacy in terms of the phenomenological features listed above, it will also have to show its viability by being a fruitful basis for developing a formal logic of means-end reasoning. We will leave the development of a formal logic of means-end reasoning for the future, but there is one aspect of means-end reasoning that requires our attention here, namely its non-monotonicity, because it is related to our analysis of means-end relations.

2 Propositional Dynamic Logic

Means-end reasoning is about the adjustment of the actual world to realize a sought-after situation that may fail to be the case in our (current) actual world. Consequently, it concerns bringing about some change in the present state of affairs such that some sentences φ describing this favorable end which are false in the actual world, will be true after a successful application of some means. Thus, we are naturally drawn to a semantics in which means-applications correspond to transitions between possible worlds. The apparently dynamic nature of means suggests *Propositional Dynamic Logic* (PDL) for this task[1].

PDL is a logic of actions, typically used to reason about computer program behavior. It is a multi modal propositional language where each atomic action corresponds with an accessibility relation. The strong modal operator $[\,]alpha\varphi$ expresses that if one does α, φ will be realized; the weak modal operator $\langle\alpha\rangle\varphi$, is defined as usual and means that α can be done, and after doing α, φ may be realized. We refer the reader to [6], from which we take much of the following material. We simplify our presentation by omitting the iteration operation α^*. For our introduction to means-end semantics, iteration is probably more distracting than necessary.

The syntax of PDL is based on 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 $\varphi \wedge \psi$ are in Φ;
- if $\alpha \in \Pi$ and $\varphi \in \Phi$ then $[\alpha]\varphi \in \Phi$;
- $\Pi_0 \subseteq \Pi$;
- if $\alpha, \beta \in \Pi$ then $\alpha;\beta$ and $\alpha \cup \beta$ are in Π.
- if $\varphi \in \Phi$ then $\varphi? \in \Pi$.

We introduce the propositional constant $\bot$, the connectives $\neg$, $\vee$ and $\rightarrow$ and the weak operator $\langle\alpha\rangle$ as usual. The modal operator $[\alpha]\varphi$ expresses that, if one does α, then φ will be realized. The action constructors are intended thus: the semicolon denotes sequential composition and the union $\alpha \cup \beta$ of actions represents non-deterministic choice between α and β.

A *PDL frame* $\mathbf{F}$ for Π_0 consists of a set $\mathcal{W}$ of worlds (or states) and a *dynamic interpretation* $[\![-]\!]^{\mathbf{F}} : \Pi_0 \rightarrow (\mathcal{PW})^{\mathcal{W}}$ 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 $\Pi_0 \rightarrow (\mathcal{PW})^{\mathcal{W}}$ assigns to each $m \in \Pi_0$ a function $[\![m]\!] : \mathcal{W} \rightarrow \mathcal{PW}$. For $w \in \mathcal{W}$, we interpret $[\![m]\!](w)$ as the set of possible[2] outcomes of doing m in w. Clearly, a PDL frame is just the same as a labeled transition system with nodes $w \in \mathcal{W}$ and labels $m \in \Pi_0$. We sometimes write $w \xrightarrow{m} w'$ for $w' \in [\![m]\!](w)$.

A *PDL model* is a frame $\mathbf{F}$ together with a *valuation* $[\![-]\!]^{\mathcal{M}} : \Phi_0 \rightarrow \mathcal{PW}$ 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 often omit the superscripts $\mathbf{F}$ and $\mathcal{M}$. A valuation assigns to each atomic

[1] The modal μ-calculus is another strong candidate, providing more expressive power than PDL. The models are essentially the same as PDL models, namely labeled transition systems, but the μ-calculus adds the strength of fixed point operators. This allows action constructions such as: do α until φ holds. It seems likely that such constructions play a role in means-end relations, but we prefer to use the simpler PDL for now and postpone a discussion of means in the μ-calculus for later investigations.

[2] Or "normal" or "reasonably expected" or ....

proposition $P \in \Phi_0$ a set $[\![P]\!] \subseteq \mathcal{W}$ of worlds. We interpret $[\![P]\!]$ as the set of worlds in which P is true.

We extend the valuation of atomic propositions to a function $[\![-]\!] : \Phi \to \mathcal{P}\mathcal{W}$ and the interpretation of atomic actions to a function $[\![-]\!] : \Pi \to (\mathcal{P}\mathcal{W})^{\mathcal{W}}$ recursively as shown in Table 1.

On formulas

$$
\begin{aligned}
[\![\top]\!] &= \mathcal{W} \\
[\![\neg\varphi]\!] &= \mathcal{W} \setminus [\![\varphi]\!] \\
[\![\varphi \wedge \psi]\!] &= [\![\varphi]\!] \cap [\![\psi]\!] \\
[\![[\alpha]\varphi]\!] &= \{ w \in \mathcal{W} \mid [\![\alpha]\!](w) \subseteq [\![\varphi]\!] \}
\end{aligned}
$$

On actions

$$
\begin{aligned}
[\![\alpha;\beta]\!](w) &= \{ w' \in \mathcal{W} \mid \exists w'' \in \mathcal{W} \,.\, w'' \in [\![\alpha]\!](w) \text{ and } w' \in [\![\beta]\!](w'') \} \\
[\![\alpha \cup \beta]\!](w) &= [\![\alpha]\!](w) \cup [\![\beta]\!](w) \\
[\![\varphi?]\!](w) &= \begin{cases} \{w\} & \text{if } w \in [\![\varphi]\!]; \\ \emptyset & \text{else.} \end{cases}
\end{aligned}
$$

Table 1: Extension of valuation to Φ and interpretation to Π.

We say that w *satisfies* φ or that φ *is true* in w just in case $w \in [\![\varphi]\!]$. In this case, we write $\mathcal{M}, w \models \varphi$ or just $w \models \varphi$ when $\mathcal{M}$ is understood by context. We write $\mathcal{M} \models \varphi$ if for every $w \in \mathcal{W}$ we have $w \models \varphi$ and we write $\models \varphi$ if $\mathcal{M} \models \varphi$ for every model $\mathcal{M}$. In this case, we say that φ is valid.

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

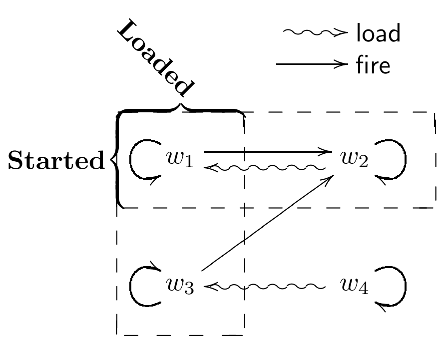

Sample facts about our model		
$[\![\textbf{Started}]\!]$	$=$	$\{w_1, w_2\}$
$[\![\textbf{Loaded}]\!]$	$=$	$\{w_1, w_3\}$
$[\![[\mathsf{fire}]\textbf{Started}]\!]$	$=$	$[\![\textbf{Started}]\!]$
$[\![\langle\mathsf{fire}\rangle\textbf{Started}]\!]$	$=$	$[\![\textbf{Started}]\!] \cup [\![\textbf{Loaded}]\!]$
$[\![[\mathsf{load}]\textbf{Loaded}]\!]$	$=$	$\mathcal{W}$
$[\![\langle\mathsf{load}\rangle\textbf{Loaded}]\!]$	$=$	$[\![\neg\textbf{Loaded}]\!]$
$[\![[\textbf{Loaded}?;\mathsf{fire}]\textbf{Started}]\!]$	$=$	$[\![\textbf{Started}]\!] \cup [\![\neg\textbf{Loaded}]\!]$
$[\![\langle\textbf{Loaded}?;\mathsf{fire}\rangle\textbf{Started}]\!]$	$=$	$[\![\textbf{Loaded}]\!]$

Figure 1: A sample PDL model

Example 2.1. Consider the example of a footrace about to begin. The starter has a (one-shot) pistol and the race will begin as soon as the pistol discharges a blank. 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.

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.

We consider a model of four worlds, so that each combination of atomic predicates is represented. See Figure 1, in which an arrow $w \rightsquigarrow w'$ denotes that $w' \in [\![\mathsf{load}]\!](w)$ and $w \longrightarrow w'$ that $w' \in [\![\mathsf{fire}]\!](w)$. We assume that one cannot load an already loaded gun. Further, just to make the model more interesting, we assume that our starter pistol may misfire. When a loaded pistol misfires, nothing in the world changes, so that fire has reflexive transitions in w_1 and w_3 in addition to the transitions representing successful discharge of a blank.

Let us consider the last but one equation in Figure 1. In w_1, checking whether the gun is loaded results in w_1 and then fire (successful discharge or misfire) results in a world where the race is started. Thus, w_1 satisfies $[\mathbf{Loaded}?;\mathsf{fire}]\mathbf{Started}$, as shown. In worlds w_2 and w_4 the action **Loaded**? returns the empty set and the empty set satisfies $[\mathsf{fire}]\mathbf{Started}$ trivially. Consequently, w_2 and w_4 also satisfy $[\mathbf{Loaded}?;\mathsf{fire}]\mathbf{Started}$. Finally, w_3 fails to satisfy the formula. The result of applying **Loaded**? in w_3 is w_3 again, but $w_3 \not\models [\mathsf{fire}]\mathbf{Started}$ since the gun may misfire, resulting in world w_3 (again) where the race has not been started. The reader may confirm that the model also satisfies the other equations in Figure 1.

To complete our introduction to PDL, we present in Table 2 the standard axiom system for PDL, taken from [6]. For rules of inference, we write φ/ψ to mean: From φ infer ψ. We omit the proof that this system is sound and complete for our semantics, i.e. that $\vdash \varphi$ iff $\models \varphi$.

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	$[\alpha \cup \beta]\varphi \leftrightarrow ([\alpha]\varphi \wedge [\beta]\varphi)$
Test	$[\psi?]\varphi \leftrightarrow (\psi \rightarrow \varphi)$
K	$[\alpha](\varphi \rightarrow \psi) \rightarrow ([\alpha]\varphi \rightarrow [\alpha]\psi)$

Inference rules

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

Table 2: The theory PDL

3 Means-end relations in PDL

When we say that an action α is a sufficient means for the end φ in w, we mean that, if one does α in w, then φ will be realized. However, we must be careful to avoid trivial ascriptions, as when the action α is prohibited in w. If one cannot do α, then surely it is not a means to any end at all. Thus, α is a sufficient means to φ in w if (1) doing α in w ensures that φ and (2) one *can* do α. This yields the following definition.

Definition 3.1. An action α is a *(strongly) sufficient means* to φ in w iff

$$w \models [\alpha]\varphi \wedge \langle\alpha\rangle\top.$$

We say that α is a *weakly sufficient means* to φ in w iff

$$w \models \langle\alpha\rangle\varphi.$$

Note that, because actions and formulas are disjoint, we see that means and ends are distinct, satisfying (1) in the introduction.

Semantically, α is a sufficient means to φ in w iff $\emptyset \subsetneq [\![\alpha]\!](w) \subseteq [\![\varphi]\!]$ and is weakly sufficient iff $[\![\alpha]\!](w) \cap [\![\varphi]\!] \neq \emptyset$. In case that one wants to realize φ, then he may be sure to do so by performing any sufficient means (but there may be reasons that he chooses not to perform any sufficient means, of course). He *cannot* realize φ *without* performing a weakly sufficient means.

The practical consequences of sufficient means are difficult to analyze. One cannot say that an agent should (on pain of practical irrationality, say) either give up his end or perform a given sufficient means. An agent may give up the certainty of realizing his end in order to avoid undesired consequences from sufficient means. (Some might argue that he should give up his end or perform *some* weakly sufficient means, however.)

Thus, many treatments of practical reasoning (including von Wright's important contribution in [13]) spend considerable time on analyzing necessary means rather than sufficient means. Necessary means yield relatively clear practical conclusions. According to von Wright, for instance, if one wants φ, then one must be willing to do what is necessary to realize φ. Indeed, he writes (emphasis in original):

> "Instead of saying 'he will act' I could also have said 'he will necessarily act.' This, moreover, is *logical* necessity. For, if action does not follow, we should have to describe the subject's case by saying either that he did not in fact *want* his professed object of desire or did not, after all, *think it necessary* to do the act in order to get the wanted thing." [13]

Regardless of whether one agrees with von Wright's strong claim, it supports the view that necessary means come with relatively clear practical consequences and that these consequences are simpler than the practical consequences of sufficient means.

However, it appears that the *semantics* of necessary means is considerably subtler than the semantics of sufficient means. Sufficiency is relatively straightforward: α is sufficient just in case doing α is sure to realize one's end. Necessary means, as they appear in the literature, are more complicated. Roughly, if one wants to achieve his end, he must execute any necessary means, *but possibly as part of a complex action.* In von Wright's example of the hut, for instance, he assumes that one must heat the hut in order to make it habitable, but he does not assume that this action by itself might actually make the hut habitable. Perhaps other acts are necessary, too.

Thus, if α is a necessary means to φ, it should not be possible to realize φ without doing α, but one must be precise on the meaning of "without doing α". We will discuss this issue later, but let us take the meaning of "without doing α" as given for now. We write $\beta \preccurlyeq_w \alpha$ to denote that, in w, one cannot do β without also doing α, i.e. that β *involves* α.

Definition 3.2. An action α is a *necessary means* to φ in w if the following hold.

1. There is an action β such that $w \models \langle\beta\rangle\varphi$ and $\beta \preccurlyeq_w \alpha$.
2. For every action β, if $w \models \langle\beta\rangle\varphi$ then $\beta \preccurlyeq_w \alpha$.

Definitions 3.1 and 3.2 are summarized (with other definitions) in Table 5 on page 11.

If one wants to represent necessary means in a single formula, we may add a two new action constructions $\neg\alpha$ and $\widehat{\alpha}$, where

$$[\![\neg\alpha]\!](w) = \bigcup \{ [\![\beta]\!](w) \mid \beta \not\preccurlyeq_w \alpha \},$$
$$[\![\widehat{\alpha}]\!](w) = \bigcup \{ [\![\beta]\!](w) \mid \beta \preccurlyeq_w \alpha \}.$$

Then α is a necessary means to φ just in case $w \models \langle\widehat{\alpha}\rangle\varphi \wedge \neg\langle\neg\alpha\rangle\varphi$. However, nothing in the sequel really depends on these constructions or the ability to represent the necessary means-end relation as a single formula.

Let us turn our attention, then, to the relation $\preccurlyeq_w$. Naively, one might define $\beta \preccurlyeq_w \alpha$ iff $\beta = \alpha$, but this is clearly too strong. If α can realize φ, then so can $\top?;\alpha$ and so there would be *no* necessary means according to this definition. Alternatively, one might define $\beta \preccurlyeq_w \alpha$ iff $[\![\beta]\!](w) \subseteq [\![\alpha]\!](w)$, but this solution is not satisfactory either. Flipping a coin produces the same set of possible outcomes as choosing to lay the coin on the table face up or face down, but there is no sense in which the latter "involves[3]" the former (or vice versa).

With these simple suggestions out of the way, let us consider what properties *are* appropriate to the relation $\preccurlyeq_w$. As suggested by the notation, we expect it to be a pre-order: reflexive and transitive. We also expect that $\alpha \preccurlyeq_w (\alpha \cup \beta)$: one cannot do α without doing $\alpha \cup \beta$. Put differently, the claim $w \models \langle\alpha\rangle\varphi$ should not refute the claim that $\alpha \cup \beta$ is a necessary means to φ.

Somewhat more controversially, we think that $w \models \langle\alpha;\beta\rangle\varphi$ should not refute the claim that β is a necessary means for φ (nor the claim that α is necessary). Consider slightly complicating our footrace example by adding a safety toggle switch to our pistol, together with an action toggle. It seems to us that in w_3, the action fire is a necessary means to **Started** and that this is true even though $\mathsf{toggle};\mathsf{toggle};\mathsf{fire}$ might also realize **Started**. The sequence $\mathsf{toggle};\mathsf{toggle};\mathsf{fire}$ includes the action fire and so the practical importance of necessary means is still present. The complex action does not refute the claim that, in order to realize **Started**, one must do fire (but he may do it as part of a complex action).

These and a few other considerations lead to the following definition of the $\preccurlyeq$ family of pre-orders. We introduce a family $\{ \preccurlyeq_S \mid S \subseteq \mathcal{W} \}$ of pre-orders defined by the deductive system[4] in Table 3, and define $\beta \preccurlyeq_w \alpha$ iff $\vdash \beta \preccurlyeq_{\{w\}} \alpha$. In the table, we use the abbreviation

$$[\![\beta]\!](S) = \{ w' \in \mathcal{W} \mid \exists w \in S \,.\, w' \in [\![\beta]\!](w) \}$$

and we write $\alpha \approx_S \beta$ for $\alpha \preccurlyeq_S \beta \ \wedge \ \beta \preccurlyeq_S \alpha$.

[3]We find a notion of involvement in [4] which the authors indicate can be easily extended to include sequential composition, but the obvious extension does not satisfy our requirements for composition. Also, they do not include the test operator, which leads to some of the complexity in our notation, specifically the need to subscript $\preccurlyeq$ with sets of worlds. A full comparison of our involvement relation with that found in ibid may be the subject of further research.

[4]Note: we define the relation $\preccurlyeq_w$ via a deductive system, but that doesn't mean that the relation itself is inherently syntactic. It depends strongly on our model, as can be seen in the axiom for $\varphi? \preccurlyeq_S \psi?$ and the rule of inference concluding with $\beta;\alpha \preccurlyeq_S \beta;\gamma$. The deductive system is simply a convenient means of defining our relation.

Axioms

$$\alpha \preccurlyeq_S \alpha \qquad \varphi? \preccurlyeq_S \psi? \text{ if } S \models \varphi \rightarrow \psi \qquad \bot?; \alpha \preccurlyeq_S \bot? \qquad \top?; \alpha \preccurlyeq_S \top?$$

$$\bot? \preccurlyeq_S \alpha \qquad \alpha \preccurlyeq_S \top? \qquad \alpha; \bot? \preccurlyeq_S \bot? \qquad \alpha; \top? \preccurlyeq_S \top?$$

$$\alpha \preccurlyeq_S \alpha \cup \beta \qquad \alpha \cup \beta \preccurlyeq_S \beta \cup \alpha$$

$$\alpha \cup (\beta \cup \gamma) \approx_S (\alpha \cup \beta) \cup \gamma \qquad \alpha \cup \bot? \preccurlyeq_S \alpha$$

Rules

$$\frac{\alpha \preccurlyeq_S \beta \quad \beta \preccurlyeq_S \gamma}{\alpha \preccurlyeq_S \gamma} \qquad \frac{\alpha \preccurlyeq_S \gamma}{\alpha; \beta \preccurlyeq_S \gamma; \beta}$$

$$\frac{\alpha \preccurlyeq_S \gamma \quad \beta \preccurlyeq_S \gamma}{\alpha \cup \beta \preccurlyeq_S \gamma} \qquad \frac{\alpha \preccurlyeq_{[\![\beta]\!](S)} \gamma}{\beta; \alpha \preccurlyeq_S \beta; \gamma}$$

$$\frac{\alpha \preccurlyeq_T \beta}{\alpha \preccurlyeq_S \beta} \text{ given } S \subseteq T \qquad \frac{\alpha \preccurlyeq_S \beta \quad \alpha \preccurlyeq_T \beta}{\alpha \preccurlyeq_{S \cup T} \beta}$$

Table 3: The deductive system for $\preccurlyeq_S$.

For some systems, one might also include some additional axioms $m \preccurlyeq_{\mathcal{W}} n$ representing relations among the atomic actions. We make no assumption about whether atomic actions might be related in this way.

Example 3.3. We return to the footrace from Example 2.1 and investigate some sufficient and necessary means for starting the race.

For w_1 and w_2, the action $\top?$ is both a necessary and sufficient means for **Started**. Indeed, *any* non-prohibited action is a sufficient means for **Started** in these worlds. In w_2, the action load is a sufficient means for **Started**, but not in w_1, since it is prohibited there.

For w_3 and w_4, there is *no* sufficient means for starting the race, since the possibility of misfire precludes any guarantee that **Started** will be realized. In w_3, fire is weakly sufficient and in w_4, the composite $\mathsf{load}; \mathsf{fire}$ is weakly sufficient.

In w_3 and w_4, the action fire is a necessary means to **Started**. In w_4, the action load is also a necessary means, as is the composite $\mathsf{load}; \mathsf{fire}$.

In every world, the action

$$\alpha \stackrel{\text{def}}{=} \textbf{Started}? \cup \big(\neg\textbf{Started}?; (\textbf{Loaded}? \cup (\neg\textbf{Loaded}?; \mathsf{load})); \mathsf{fire}\big)$$

is a necessary and weakly sufficient means to **Started**. Moreover, it is maximally effective, in the sense that for all w, and for all actions β, if β is a (strongly) sufficient means to **Started** in w, then so is α.

4 Non-monotonic conditionals

Very often, achieving an end requires a sequence of actions. In our Example 2.1, for instance, if the actual world is w_4 (pistol unloaded and race not started), then it is necessary to load the pistol and fire it to start the race. Although we emphasize semantics over reasoning in the present paper, we want a semantics that supports the reasoning that yields such complex plans.

We are not yet prepared to present a system of practical reasoning, but we will nonetheless sketch some arguments that we think approximates some part of the reasoning involved in forming plans. Our footrace judge in world w_4 may decide to execute $\mathsf{load};\mathsf{fire}$ on the grounds that (1) $\mathsf{load};\mathsf{fire}$ is a necessary means to **Started**, (2) $\mathsf{load};\mathsf{fire}$ is a weakly sufficient means to **Started** and (3) no other means is (strongly) sufficient for **Started**, so that $\mathsf{load};\mathsf{fire}$ is as good a plan as any. Let us focus on (2) and ask how he might be persuaded that it is true, i.e. that $w \models \langle \mathsf{load};\mathsf{fire}\rangle \mathbf{Started}$, which surely involves the simplest reasoning of the three relevant facts.

Even here, we hesitate to present a formal deduction of this theoretical claim, but we suppose that the reasoning may go something like this. If the gun is loaded, then firing the gun may start the race and in w_4, the action load makes the gun loaded. Hence, we believe that the reasoning involves conditionals and can be represented so:

$$\frac{\begin{array}{l}\mathbf{Loaded} \Rightarrow \langle \mathsf{fire}\rangle \mathbf{Started}\\ \langle \mathsf{load}\rangle \mathbf{Loaded}\end{array}}{\langle \mathsf{load};\mathsf{fire}\rangle \mathbf{Started}} \tag{4.1}$$

Such conditionals are, we think, essential to means-end semantics. They justify the formation of complex plans by making the role of intermediate ends explicit: an *intermediate end* is an end which is selected because it is a precondition for a means-end relation involving another desired (possibly intermediate) end.

Such means-end reasoning is typically *defeasible*. This is better illustrated by another example, which we call The Shortsighted Suitor. The reasoning goes thus.

$$\frac{\begin{array}{l}\text{If I had money then she might agree to my proposal for marriage.}\\ \text{Robbing her is a means to having money.}\end{array}}{\text{If I robbed her then she might agree to my proposal for marriage.}}$$

This argument can be represented thus:

$$\frac{\begin{array}{l}\mathbf{Money} \Rightarrow \langle \mathsf{ask}\rangle \mathbf{Married}\\ \langle \mathsf{rob}\rangle \mathbf{Money}\end{array}}{\langle \mathsf{rob};\mathsf{ask}\rangle \mathbf{Married}} \tag{4.2}$$

The argument form is thus very similar to that found in the previous example. It is nonetheless obvious that the conclusion is unwarranted: if I rob my sweetheart, it is unlikely that she will marry me (we assume that such a clueless suitor forgets to mask his identity).

The problem here is a familiar issue in practical reasoning, often called the frame problem [3]. One may try to avoid it by denying that the conditional $\mathbf{Money} \Rightarrow \langle \mathsf{ask}\rangle \mathbf{Married}$ is true: after all, we have found a situation in which it is false, haven't we? But this solution merely sidesteps the issue by pretending that practical reasoning is easy. Reasoning about ends frequently involves conditionals such as this (see, e.g. , [12]) and our means-end semantics ought to reflect this.

Instead, we analyze the situation thus: While it is true that $\mathbf{Money} \Rightarrow \langle \mathsf{ask}\rangle \mathbf{Married}$, it is false that $(\mathbf{Money} \wedge \mathbf{Hate}) \Rightarrow \langle \mathsf{ask}\rangle \mathbf{Married}$, where **Hate** is the proposition that she hates our suitor. Since $[\mathsf{rob}]\mathbf{Hate}$ is also true, we see where the above argument goes wrong. Our conditional operator is *non-monotonic*.

The literature on conditional operators is broad, but we hope that a few simple definitions will satisfy our purposes. At present, we value flexibility over logical commitments, pending further reflection. We propose the following (tentative) semantics for our conditional operator. We add to our PDL frames a function $r : \mathcal{PW} \to (\mathcal{PW})^{\mathcal{W}}$ satisfying that[5] for every world w and

[5]One probably wants some non-trivial relations to hold between the conditional operator and the dynamic operators, such as the axiom $[\alpha](\varphi \Rightarrow \psi) \to ([\alpha]\varphi \Rightarrow [\alpha]\psi)$. Such features can be introduced by adding appropriate restrictions to r, but we will not investigate them here.

set $S \subseteq \mathcal{W}$,

$$r(S)(w) \subseteq S.$$

We interpret $r(S)(w)$ to be the set of S-worlds that are reasonably close to w. The idea is similar to the minimal-change or small-change conditionals discussed in [9], but one important difference is that we do not require that $w \in r(S)(w)$ if $w \in S$. Our conditionals are intended to capture a sense of normality: *normally*, given φ, ψ is true. There's no requirement that the actual world is "normal".

We extend the semantics of Section 2 to include

$$[\![\varphi \Rightarrow \psi]\!] = \{\, w \in \mathcal{W} \mid r([\![\varphi]\!])(w) \subseteq [\![\psi]\!] \,\}\,.$$

In other words, $\varphi \Rightarrow \psi$ evaluates to true at w iff each of the φ-satisfying worlds relevant to w also satisfy ψ. Because $r(w, S) \subseteq S$, every world w satisfies $\varphi \Rightarrow \varphi$.

Our models satisfy the following axioms and inference rules, taken from [9] and [10]. (This list is not minimal: axioms CC and CM, for instance, are derivable from the remainder.)

Axioms

ID:	$\varphi \Rightarrow \varphi$
CC:	$((\varphi \Rightarrow \psi) \wedge (\varphi \Rightarrow \chi)) \rightarrow (\varphi \Rightarrow (\psi \wedge \chi))$
CM:	$(\varphi \Rightarrow (\psi \wedge \chi)) \rightarrow ((\varphi \Rightarrow \psi) \wedge (\varphi \Rightarrow \chi))$

Inference rules

RCEC:	$\varphi \leftrightarrow \psi$	/	$(\chi \Rightarrow \varphi) \leftrightarrow (\chi \Rightarrow \psi)$	
RCK:	$(\varphi_1 \wedge \ldots \wedge \varphi_n) \rightarrow \psi$	/	$((\chi \Rightarrow \varphi_1) \wedge \ldots \wedge (\chi \Rightarrow \varphi_n)) \rightarrow (\chi \Rightarrow \psi)$	$(n \geq 0)$
RCEA:	$\varphi \leftrightarrow \psi$	/	$(\varphi \Rightarrow \chi) \leftrightarrow (\psi \Rightarrow \chi)$	
RCE:	$\varphi \rightarrow \psi$	/	$\varphi \Rightarrow \psi$	

Table 4: Logical properties of $\Rightarrow$.

Clearly, one would like a fuller discussion of our conditional semantics and its appropriateness for means-end reasoning. We consider the semantics presented here as fairly minimal in its commitments, so that later revisions may provide further commitments rather than retract existing commitments. This is in keeping with our bias for flexibility.

5 Sufficient and necessary pre-conditions

We read a sentence $\psi \Rightarrow [\alpha]\varphi$ as, "Given ψ, doing α will realize φ," but this reading should be weakened in the presence of non-monotonic conditionals. It is better to read it as, "*Normally*, ψ implies that doing α will realize φ." We use such conditional formulas to define conditional means-end relations.

Definition 5.1. In a world w, an action α *is a (weakly, resp.) sufficient means to* φ *given* ψ if

$$w \models \psi \Rightarrow ([\alpha]\varphi \wedge \langle\alpha\rangle\top),$$

($w \models \psi \Rightarrow \langle\alpha\rangle\varphi$, resp.). Similarly, α *is a necessary means to* φ *given* ψ if

1. there is some $w' \in r([\![\psi]\!])(w)$ and $\beta \preccurlyeq_{w'} \alpha$ such that $w' \models \langle\beta\rangle\varphi$;

2. For every action β and world $w' \in r(\llbracket\psi\rrbracket)(w)$, if $w' \models \langle\beta\rangle\varphi$ then $\beta \preccurlyeq_w \alpha$.

Such conditional relations can be used for creating tentative plans. If α is a sufficient means to φ given ψ, then *normally*, one has the option of doing α to realize φ when ψ holds. However, there may be certain circumstances in which ψ holds and doing α fails to realize φ. (Again, our focus is not on defeasible practical reasoning at present, but our means-end semantics should provide some explanation as to *why* such reasoning is naturally defeasible.)

The formula ψ in Definition 5.1 is a *sufficient precondition* for the (sufficient or necessary) means-end relation involving α and φ. When ψ holds in w, then one expects the corresponding non-conditional (sufficient or necessary) means-end relation to hold in w, too, but this expectation may be thwarted.

For sufficient means-end relations, one may also talk of *necessary preconditions*. If ψ is a necessary precondition for α to be a sufficient means to φ, then *normally* α is *not* a sufficient (necessary, resp.) means to φ unless ψ holds.

Definition 5.2. We say that ψ is a *necessary precondition* for α to be a (weakly, resp.) sufficient means for φ in a world w iff

$$w \models ([\alpha]\varphi \wedge \langle\alpha\rangle\top) \Rightarrow \psi.,$$

($w \models \langle\alpha\rangle\varphi \Rightarrow \psi$, resp.)

We have not found a suitable corresponding definition of "necessary preconditions for necessary means-end relations," but neither have we felt much need for this concept. It does not seem to arise naturally in reasoning about means and ends.

The distinction between pre-conditions and means is crucial to our analysis and leads to much confusion in natural language. In natural language, one sometimes calls a pre-condition (necessary or sufficient) a *means* and this leads to claims that a particular fact is both a means and an end (as mentioned in (2) in the introduction). For instance, the attainment of a bachelor's degree is an end but (some say) it is also a means to a better career. This terminology is inconsistent with our usage. The end for which one attends college is the condition, "He has a bachelor's degree." But a proposition cannot be a means: a means is something that can change the world so that an end is realized, and a proposition isn't such a thing.

So what *does* one "really" mean when he says the degree is a means? We agree that the end of having a degree plays a dual role in this example, but it is not the role of end and means. It is the role of (intermediate) end and *precondition*. Having a bachelor's degree is a precondition[6] for a means-end relation. If one has a bachelor's degree and applies for good careers, he may get one[7]. Our analysis forces a sharp distinction between preconditions for means-end relations and means themselves, a distinction which is much fuzzier in natural language, but we regard this feature as clarifying the situation crudely expressed in natural language.

We summarize our taxonomy of means-end relations in Table 5.

Example 5.3. We return to the footrace and starter pistol example to show the flexibility and consequences of our conditional means-end relations. We will do so by augmenting our example with three different functions $r : \mathcal{PW} \to (\mathcal{PW})^{\mathcal{W}}$ and also by adding some new worlds in which the gun is malfunctioning for the third treatment.

For the first treatment, define $r_{\mathsf{m}}(S)(w) = S$ for every set $S \subseteq \mathcal{W}$ and $w \in \mathcal{W}$. In this case, the conditional connective $\Rightarrow$ coincides with material implication (the subscript m stands for

[6] A sufficient precondition? A necessary precondition? Only context can tell.

[7] This example is not very well-analyzed in terms of sufficient or necessary means. The reason to get a degree is that it makes a better career more likely, not that it makes it either a possible or necessary outcome of applications. Thus, this example is better analyzed in a semantics with a measure of efficacy, like our fuzzy set semantics for PDL, briefly discussed in Section 6.

In w, α is a .?. means to φ iff...	Unconditional	Conditional (w.r.t ψ)	
		Sufficient	Necessary
Sufficient	$w \models [\alpha]\varphi \wedge \langle\alpha\rangle\top$	$w \models \psi \Rightarrow ([\alpha]\varphi \wedge \langle\alpha\rangle\top)$	$w \models ([\alpha]\varphi \wedge \langle\alpha\rangle\top) \Rightarrow \psi$
Weakly sufficient	$w \models \langle\alpha\rangle\varphi$	$w \models \psi \Rightarrow \langle\alpha\rangle\varphi$	$w \models \langle\alpha\rangle\varphi \Rightarrow \psi$
Necessary	$\exists\beta \preccurlyeq_w \alpha,$ $w \models \langle\beta\rangle\varphi.$ $\forall\beta \not\preccurlyeq_w \alpha,$ $w \not\models \langle\beta\rangle\varphi.$	$\exists w' \in r([\![\psi]\!])(w), \beta \preccurlyeq_w \alpha,$ $w' \models \langle\beta\rangle\varphi.$ $\forall w' \in r([\![\psi]\!])(w), \beta \not\preccurlyeq_w \alpha,$ $w' \not\models \langle\beta\rangle\varphi.$	n/a

Table 5: A summary of our means-end relations.

"material implication"). Thus, in every world w, a formula ψ is a sufficient precondition for some means-end relation just in case every $w' \in [\![\psi]\!]$ satisfies the means-end relation.

In the second treatment, we define r so that it reflects epistemic limitations of our agent. We suppose that the agent knows whether the race has started or not, but he does not know whether the gun is loaded or not. As a result, in w_4 (say), he regards w_3 as more relevant than w_1 or w_2 and equally relevant as w_4 itself. Hence, we define

$$r_{\mathsf{e}}(S)(w_1) = r_{\mathsf{e}}(S)(w_2) = \begin{cases} S & \text{if } S \subseteq \{w_3, w_4\}; \\ S \cap \{w_1, w_2\} & \text{else} \end{cases}$$

$$r_{\mathsf{e}}(S)(w_3) = r_{\mathsf{e}}(S)(w_4) = \begin{cases} S & \text{if } S \subseteq \{w_1, w_2\}; \\ S \cap \{w_3, w_4\} & \text{else} \end{cases}$$

In this example, we see that, in w_4, fire is a necessary means to **Started** given the trivial precondition $\top$, since $r_{\mathsf{e}}([\![\top]\!])(w_4) = \{w_3, w_4\}$. This is not true in our first example r_{m}, since $r_{\mathsf{m}}([\![\top]\!])(w_4) = \mathcal{W}$ and fire is not a necessary means to **Started** in w_1 or w_2.

For the third example, let us complicate our model by supposing that the gun may be broken. When the gun is broken, it always fails to fire. Thus, we add the following worlds and transitions to our model as presented in Figure 1, including the new action fix.

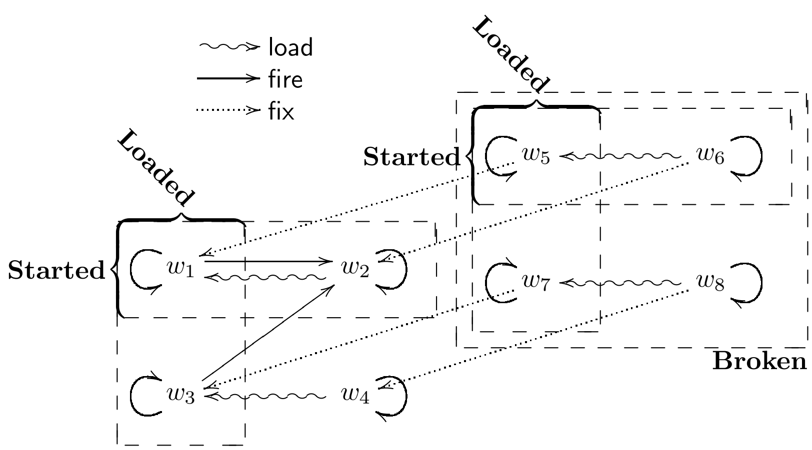

Figure 2: Additional worlds for the footrace model.

We may suppose that the gun isn't "normally" broken. Thus, we define

$$r_{\mathsf{b}}(S)(w) = \begin{cases} S & \text{if } S \subseteq [\![\textbf{Broken}]\!]; \\ S \setminus [\![\textbf{Broken}]\!] & \text{else.} \end{cases}$$

With this definition, we assume that *even in worlds in which the gun is broken,* it isn't "normally" broken. Thus, $w_5 \notin r_{\mathsf{b}}(\mathcal{W})(w_5)$.

This model agrees with the material implication conditional defined by r_{m} on sufficient preconditions ψ, provided $\not\models \psi \to \textbf{Broken}$. The model also satisfies that, in every world w, fix is a necessary means to **Started**, given **Broken**. Moreover, in every world, fire is a weakly sufficient means to **Started** given **Loaded**, but *not* given **Loaded** $\wedge$ **Broken**.

6 Further developments of means-end semantics

In this section, we sketch some further developments of our semantics. This includes material alluded to in our introduction, but for which a full treatment cannot be given here. In particular, we would like to indicate how to (1) understand objects as means, (2) include ends-in-themselves and (3) represent efficacy as an important feature of means. These three aims correspond to items (6), (3) and (7), respectively, from our introduction.

It is very common to speak of objects as means, in contrast with our semantics in which means are actions in a dynamic logic. For instance, one may say that a particular key is a means to gain entry to a house. From our perspective, this is a linguistic shorthand. Keys don't change the world, and so keys cannot bring about ends, but keys can be *used* to achieve an end. Thus, for each object o, we may introduce an action "use o" and say that o is a means to φ iff the action "use o" is a means to φ in the usual sense.

There is one important difference between such actions and the models we've considered up to now, however. Keys can lock and unlock doors and thermostats can be used to heat and cool a room. Actions like "use o" can bring about mutually exclusive ends, even reliably (as sufficient means). This violates the distributivity axiom of PDL (and hence also the K axiom). Thus, we will have move to monotone neighborhood semantics for our dynamic logic.

Monotone neighborhood semantics changes the dynamic interpretation so that $[\![m]\!](w)$ is a *(up-closed) set of* set of possible worlds. This set captures the alternative non-deterministic outcomes of an action. Such semantics have been used in game logic [11] and elsewhere, including Mark Brown's logic of ability [1] (closely related to our means-end semantics and using neighborhood semantics for much the same reason). See [2] for an introduction to neighborhood semantics (called *minimal models* there).

The second item is fairly simple to accommodate. Some actions we perform because of the pleasure they give us. For instance, one might swim just because he enjoys swimming. Thus, swimming is evidently a means and an end, contrary to our claim that means and ends are distinct types. But this situation is not so difficult: the action of swimming is distinct from the condition that one is currently swimming. We simply introduce an action swim and a related atomic proposition **IsSwimming** and continue as before.

In fact, this situation can lead one to neighborhood semantics independently of objects-as-means considerations. The action swim is a sufficient means to **IsSwimming**, but it is also a sufficient means to the proposition **IsAcross** that I have crossed to the opposite bank of the river. Clearly **IsAcross** and **IsSwimming** are mutually exclusive: if I am swimming, then I have not reached the opposite shore and vice versa. We are therefore pushed to neighborhood semantics regardless of our stance on objects-as-means.

Finally, efficacy is the propensity of a means to realize its end. Efficacy is a measure of how effective a means is, as opposed to its *efficiency*, which we take to roughly indicate the undesirable consequences of a means. The non-deterministic semantics discussed up to now can only crudely represent efficacy (a sufficient means is more effective than a means that does not ensure the realization of its end) but one may better represent efficacy by including probabilistic elements in our semantics. In our footrace example, we pretended that nothing could be said about the frequency or expectation of a misfire. This is unrealistic: surely, the misfire is less likely than a successful discharge (or one should buy a new gun or new blanks). This fact says something about the efficacy of fire in realizing **Started**. It is more effective than not.

One of the ways in which one compares means (and hence forms plans) is their efficacy. Thus, one would like a semantics in which this feature is explicitly represented. We can do this by attaching probabilities to the outcomes of actions and describing the resulting models in terms of fuzzy sets, so that the set $[\![\varphi]\!]$ of worlds satisfying φ is *fuzzy* and membership is a matter of degree. To do this, we introduce a new interpretation of the dynamic operators in terms of weighted averages. This yields a new fuzzy set semantics for PDL and is the subject of a forthcoming article.

References

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

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

[3] Daniel C. Dennett. Cognitive wheels: The frame problem of AI. In Margaret A. Boden, editor, *The Philosophy of Artificial Intelligence*, Oxford Readings in Philosophy, pages 147 – 170. Oxford University Press, 1990.

[4] F. Dignum, J.-J.Ch.Meyer, and R.J. Wieringa. Contextual permission: A solution to the free choice paradox. In Andrew J.I. Jones and Marek Sergot, editors, *DEON'94, Second Internation Workshop on Deontic Logic in Computer Science*, pages 107–135, 1994.

[5] Clive L. Dym and Philip Brey. Languages for engineering design: Empirical constructs for representing objects and articulating processes. In Peter Kroes and Anthonie Meijers, editors, *The Empirical Turn in the Philosophy of Technology*, pages 119–148. JAI/Elsevier, 2000.

[6] David Harel. Dynamic logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume II, pages 497–604. D. Reidel Publishing Company, 1984.

[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] W. Houkes and P.E. Vermaas. Actions versus functions: A plea for an alternative metaphysics of artifacts. *Monist*, 87:52–71, 2004.

[9] Donald Nute. Conditional logic. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic*, volume II, pages 387–439. D. Reidel Publishing Company, 1984.

[10] Donald Nute. Defeasible logic. In D. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Philosophical Logic*, volume III, pages 353–395. D. Reidel Publishing Company, 1994.

[11] H. Parikh. The logic of games and its applications. In M. Karpinski and J. van Leeuwen, editors, *Topics in the Theory of Computation*, volume 24 of *Annals of Discrete Mathematics*. Elsevier, 1985.

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

[13] Georg Henrik von Wright. Practical inference. *The Philosophical Review*, 72(2):159–179, Apr. 1963.