Synthese (2007) 158:207–231
Knowledge, Rationality & Action 297–321
DOI 10.1007/s11229-006-9036-x

ORIGINAL ARTICLE

A semantics for means-end relations

Jesse Hughes · Peter Kroes · Sjoerd Zwart

Received: 6 October 2005 / Accepted: 20 April 2006 / Published online: 15 September 2006

Abstract There has been considerable work on practical reasoning in artificial intelligence and also in philosophy. Typically, such reasoning includes premises regarding means–end relations. A clear semantics for such relations is needed in order to evaluate proposed syllogisms. 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 work is primarily an exercise in conceptual analysis, aimed at clarifying and eventually evaluating existing theories of practical reasoning (pending a similar analysis regarding desires, intentions and other relevant concepts).

Keywords Means–end relations · Propositional dynamic logic · Formal semantics · 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

The aim of this paper is to improve our understanding of means–end relations in practical reasoning. We take practical reasoning to be the process of deriving prescriptions for actions, typically from premises including means–end relations. Millgram

J. Hughes (✉) · S. Zwart
Eindhoven University of Technology, P.O. Box 513, 3600 MD, Eindhoven, The Netherlands
e-mail: j.hughes@tm.tue.nl

P. Kroes · S. Zwart
Delft University of Technology, P.O. Box 5, 2600 AA Delft, The Netherlands
e-mail: p.a.kroes@tbm.tudelft.nl

S. Zwart
e-mail: s.d.zwart@tm.tue.nl

(2004) summarizes the primary issues in this field: “The current debate in practical reasoning focuses on the question: what inference patterns are legitimate methods of arriving at decisions or intentions to act, or other characteristically practical predecessors of actions such as evaluations, plans, policies, and judgments about what one ought to do?” Practical reasoning and the use of means–end relations are integral aspects of linguistic practices in general, and in particular, of engineering practices. We want to contribute to the understanding of these practices by providing a clear analysis of means–end relations.

The broad topic of practical reasoning has been the focus of much attention in artificial intelligence, largely starting with the seminal paper of McCarthy and Hayes (1969). This work aims at producing software agents capable of attaining goals by choosing appropriate actions. Some of this work has been cast explicitly in terms of means–end relations, including the engineering perspective presented in Bratman et al. (1988) and Pollock’s (2002) work. Central issues include the epistemological problem of knowledge representation, and the heuristical problem regarding decision making and achieving goals, discussed in McCarthy (1999).

Our work is primarily inspired by an older tradition in philosophical circles, namely the investigation of practical syllogisms, dating back to Aristotle and enjoying renewed interest due to the work of von Wright (1963). We are particularly indebted to von Wright’s work and take his analysis as a model for our semantics. Broadly, such syllogisms typically involve premises of the following types:

(1) assertions that an agent A desires some *end* φ,
(2) assertions that (possibly given some precondition ψ) the action α is related to the realization of φ,
(3) assertions of other matters of fact, such as that the precondition ψ is true.

Prominent among premises of type (ii) are ones, which 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 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.

Our working example of such syllogisms comes from von Wright (1963), stated in the third-person form here.

$$\frac{\begin{array}{l}A \text{ wants to make the hut habitable}\\ \text{Unless } A \text{ heats the hut, it will not become habitable}\end{array}}{\text{Therefore A must heat the hut}}.$$

Such syllogisms conclude either in an action, an intention to act or a prescription to act, depending on the author. The premises are supposed to be sufficient to justify the conclusion, of course, but how should one evaluate this claim? For this, one must have an unambiguous understanding of each of the premises involved. A clear means–end semantics, however, seems to be largely lacking in the literature. We aim to contribute by providing a semantics for means–end relations which may help to evaluate practical syllogisms. We also hope that our analysis will help clarify AI work on means–end relations, but our focus is primarily the premises of practical syllogisms. We have chosen to present a formal semantics, because such formalisms help ensure clarity and allow one to indicate precisely which features are taken to be relevant to the analysis. Any such formalization requires some idealization as well as deviation from natural language expressions,[1] but we hope that our semantics reasonably approximates natural language in the relevant features.

Means–end talk in natural language exhibits many different features and our aim is to represent these features as faithfully as possible in our formal system. We present the following list of features of means–end talk, but we make no claims regarding completeness of this list as follows:

(1) In general, means are different from ends.
(2) What is an end in one context may be a means in another.
(3) There is a fruitful distinction between necessary and sufficient means.
(4) Means and ends are causally connected.
(5) Means–end conditionals are defeasible.
(6) A means may be an end in itself, rather than a means for some distinct end.
(7) Entities of different types, such as objects and actions, may constitute means.
(8) Effective means are not necessarily efficient means.
(9) Good means may be distinguished from bad means.

These features should be taken as prima facie features of means–end talk, but our formal reconstruction of means–end relations will not accommodate all of these features For instance, we interpret feature (1) as suggesting that means and ends are distinct types, while features (2) and (6) claim that means can be ends and ends can be means. Any formal semantics distinguishing means and ends as different types will be unable to satisfy the latter two features, but we nonetheless regard our interpretation of actions-as-means and ends-as-formulas as a reasonable starting point for an analysis of means–end relations.

In the end, our semantics satisfies or explains features (1)–(5). We believe, we can give a good account of (6) and (7) and our work on efficacy in Hughes et al. (2005) forms a partial analysis of (8), but we save these considerations and an analysis of (9) for a later day.

Our most fundamental contribution comes in Sects. 2 and 3, where we provide the basic semantics for *local* means–end relations—relations, which express the sufficiency/necessity of a means to an end *now* (but maybe not later). We use models of *Propositional Dynamic Logic* (PDL) (Harel 1984) to provide the setting for these means–end relations. The PDL is a natural setting for means–end relations, since it is motivated by reasoning about outcomes of actions. The PDL has long been used for reasoning about program correctness but also has a healthy tradition in current AI research, surveyed in Meyer (2000), with additional examples in Castilho et al. (2002), Castilho et al. (1999), Giacomo and Lenzerini (1995), Giordano et al. (2000), Prendinger and Schurz (1996), and Zhang and Foo (2002, 2005). 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 certain natural language expressions via formal semantics.

An alternative tradition for practical reasoning involves temporal logic. Recently, Brown (2005) suggested a means–end semantics involving such logics with stit (sees-to-it-that) operators (see also Horty and Belnap 1995). His logic includes sophisticated

[1] See (Hansson, 2000) for an insightful discussion of the role of formalisms in philosophy.

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 is hard to see how it can represent means at all. 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 (2000) says, in PDL actions appear as "first-class citizens".

Following our discussion of local means–end relations, we introduce conditional relations in Sect. 4. Conditional means–end relations more closely approximate natural language usage and introduce certain epistemological issues. A typical agent will not know every fact about the current state of the actual world, but instead reasons about means to an end given certain features he believes to be true. A conditional operator serves to represent this limitation. Moreover, due to well-known issues in practical reasoning (notably the ramification problem), we prefer a *non-monotonic* conditional operator. We sketch some of the features that such an operator should have and in Sect. 5, we apply this operator to yield sufficient and necessary preconditions for various conditional means–end relations.

2 Means–end relations in PDL

Means–end reasoning is about the adjustment of the actual world to realize a sought-after situation that may fail to be the case at present. Consequently, it concerns doing something that brings about a change in the present state of affairs such that some sentence φ describing this favorable end will be true. As our description suggests, we find Kripke-style semantics to be appropriate for reasoning about propositions with changing truth values. Then a means to φ is a way to change the current state of the world to a state in which φ is realized. Thus, a means involves a transition from the current state to some φ-state and, inasmuch as the agent may choose to make the change or not, it is natural to think of means as actions in a dynamic logic. Hence, we follow suggestion in Segerberg's (1992) and choose PDL as our basic setting.

There are alternatives to PDL that may serve for a means–end semantics, including temporal logic (applied to means–end relations in Brown (2005) and the modal μ-calculus. The former does not seem well suited for our application, since it does not naturally include a syntax for actions. The latter is somewhat better suited, since it combines the explicit actions of PDL with many of the fixed point operators in temporal logic. But although some of these operators (while, until, etc.) may be useful for understanding complex ends, we felt that the simplicity of PDL sufficed for an introduction to means–end semantics.

Similarly, on grounds of simplicity, we rejected expansions of PDL found in the AI literature, including: EPDL used in Zhang and Foo (2001), $\mathcal{DIFR}$ in Giacomo and Lenzerini (1995) and AD in Giordano et al. (2000). Because, we are currently interested in practical syllogisms instead of automated agent reasoning, the motivations for these extensions are less pressing on our investigation. We prefer to focus presently on a conceptual analysis of means–end relations as they occur in practical reasoning rather than issues in agent-based, goal-directed artificial intelligence.

2.1 Propositional Dynamic Logic

The 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 term corresponds to an accessibility relation. The strong modal operator $[\alpha]\varphi$ expresses that φ will necessarily be realized after performing α (with no commitment that α *can* be performed). The weak modal operator $\langle\alpha\rangle\varphi$ is defined by duality as usual and means that φ might be realized if one does α. We refer the reader to Harel (1984), 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 more distracting than necessary.

The syntax of PDL is based on two disjoint types: the set Π_0 of atomic action terms and the set Φ_0 of atomic propositions. From these two sets, we inductively define the sets Π of action terms and Φ of formulas as follows:

(1) $\{\top\} \cup \Phi_0 \subseteq \Phi$;
(2) if $\varphi, \psi, \in \Phi$ then $\neg\varphi$ and $\varphi \wedge \psi$ are in Φ;
(3) if $\alpha \in \Pi$ and $\varphi \in \Phi$ then $[\alpha]\varphi \in \Phi$;
(4) $\Pi_0 \subseteq \Pi$;
(5) if $\alpha, \beta \in \Pi$ then $\alpha; \beta$ and $\alpha \cup \beta$ are in Π.
(6) if $\varphi \in \Phi$ then $\varphi? \in \Pi$.

We introduce the propositional constant $\bot$, the connectives $\vee$ and $\rightarrow$ and the weak operator $\langle\alpha\rangle$ as usual. The action term constructors are intended thus: the semicolon denotes sequential composition (do α followed immediately by β), the union $\alpha \cup \beta$ represents non-deterministic choice between α and β and the test operator $\varphi?$ allows one to form conditional action dependent on the truth value of φ.[2]

We should note that the formalism defined here has no explicit role for agents. Of course, we could introduce agents by indexing the atomic action terms (so that the action term m_x stands for "agent x does action m") and thereby introduce some elements of cooperation and interference. This simple fix introduces these issues only in a very elementary form, however, since, we did not include synchronous composition of actions and other useful constructions. In any case, our interest is to introduce some basic semantics for means–end relations here, so we will leave aside consideration of multi-agent issues for now.

A *PDL model* $\mathcal{M}$ is a triple consisting of

(1) A set $\mathcal{W}$ of states.
(2) A *dynamic interpretation* $[\![-]\!]: \Pi_0 \rightarrow (\mathcal{P}\mathcal{W})^{\mathcal{W}}$ of atomic action terms via non-deterministic transition systems and
(3) A *valuation* $[\![-]\!]: \Phi_0 \rightarrow \mathcal{P}\mathcal{W}$ of atomic propositions.

Here $\mathcal{P}$ denotes the powerset functor and exponentiation A^B denotes the set of functions $B \rightarrow A$. Consequently, the interpretation $\Pi_0 \rightarrow (\mathcal{P}\mathcal{W})^{\mathcal{W}}$ assigns to each $m \in \Pi_0$ a function $[\![m]\!]: \mathcal{W} \rightarrow \mathcal{P}\mathcal{W}$. For $w \in \mathcal{W}$, we interpret $[\![m]\!](w)$ as the set of possible outcomes[3] of doing m in w. Clearly, a dynamic interpretation is just a labeled transition

[2] The name "test operator" often creates more confusion than necessary. An action term $\varphi?$ does not represent checking the truth condition of φ, updating one's epistemic state or anything similar. Instead, $\varphi?$ is defined precisely by the semantics given in Table 1 and the reader should avoid inferences about $\varphi?$ based on observation, testing, and so on.

[3] In some situations, "possible outcomes" will be too broad for our purposes. It may be possible to kill a seagull by ricocheting a bullet off a church bell, but the probabilities of success are small enough

Table 1 Extension of valuation to Φ and interpretation to Π

On formulas

$$[\![\top]\!] = \mathcal{W}$$
$$[\![\neg\varphi]\!] = \mathcal{W}\backslash[\![\varphi]\!]$$
$$[\![\varphi \wedge \psi]\!] = [\![\varphi]\!] \cap [\![\psi]\!]$$
$$[\![[\![\alpha]\!]\varphi]\!] = \{w \in \mathcal{W} | [\![\alpha]\!](w) \subseteq [\![\varphi]\!]\}$$

On action terms

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

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 valuation assigns to each atomic proposition $P \in \Phi_0$ a set $[\![P]\!] \subseteq \mathcal{W}$ of states. We interpret $[\![P]\!]$ as the set of states, in which P is true. We extend the valuation of atomic propositions to a function $[\![-]\!]: \Phi \to \mathcal{PW}$ and the interpretation of atomic action terms to a function $[\![-]\!]: \Pi \to (\mathcal{PW})^{\mathcal{W}}$ recursively as shown in Table 1.

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 say that α is *impossible in w* if $[\![\alpha]\!](w) = \emptyset$. Intuitively, such actions cannot be performed in w. If α is impossible 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 $\top$).

We call a formula φ *attainable in w* if there is some action term α such that $w \models \langle\alpha\rangle\varphi$. Otherwise, φ is *unattainable in w*—there is no path from w to a state realizing φ. As this terminology indicates, we assume that each action relevant to our discussion of means and ends is described by some action term. We take it that practical reasoning works on linguistic entities, which describe actions, conditions, etc. Therefore, we treat those ends, which are attainable only by some unnamed action as if they are unattainable. It is hard to imagine how one reasons about actions, which cannot be described.

Remark 2.1 It is tempting but misleading to think that the function $[\![\alpha]\!]: \mathcal{W} \to \mathcal{PW}$ is the *action described by* α. In our view, $[\![\alpha]\!]$ encodes the various possible outcomes of doing α, but actions are not identified with their possible outcomes. After all, toggling a light switch 32 times may produce the same results as doing nothing (depending on how finely, we distinguish outcomes), but it is not the same action as doing nothing.

We tend rather to follow (Meyer, 1989) in interpreting the action described by α as a non-empty set of non-empty sequences of atomic and test actions, as we discuss in Sect. 3 and endnote 6. One could make this interpretation explicit by defining actions in terms of sequences, which are further interpreted as paths in a model and defining the dynamic operator in terms of these paths, but this would be rather more distracting than illuminating.

Footnote 3 continued
that, we wouldn't entertain this possibility as a realistic *means* of killing seagulls. Thus, one may prefer to interpret $[\![m]\!](w)$ as the set of "normal" or "reasonably expected" outcomes.

Table 2 Properties of the equivalence relation $\equiv$

$\alpha \cup \beta \equiv \beta \cup \alpha$	$\alpha; (\beta; \gamma) \equiv (\alpha; \beta); \gamma$	$(\alpha \cup \beta); \gamma \equiv (\alpha; \gamma) \cup (\beta; \gamma)$
$\alpha \cup \alpha \equiv \alpha$	$\alpha \cup (\beta \cup \gamma) \equiv (\alpha \cup \beta) \cup \gamma$	$\gamma; (\alpha \cup \beta) \equiv (\gamma; \alpha) \cup (\gamma; \beta)$
$\alpha \cup \bot? \equiv \alpha$	$\varphi? \equiv \psi?$ if $\vdash \varphi \leftrightarrow \psi$	

Instead, we will leave the semantic concept "action" implicit as far as possible. We will sometimes want to know, however, when two action terms describe the same action. For this, let $\equiv$ be the least equivalence relation satisfying the axioms in Table 2, and say that α and β describe the same action whenever $\alpha \equiv \beta$.

Note that if $\alpha \equiv \beta$, then $\vdash [\alpha]\varphi \leftrightarrow [\beta]\varphi$, but the converse is not true. Indeed, in Sect. 3, we will claim that there are some action terms α and β such that $\vdash [\alpha]\varphi \leftrightarrow [\beta]\varphi$ yet α and β do not share the same necessary means–end relations. See Remark 3.4. Equivalence of action terms α and β is not intended to express that the two actions yield the same outcomes. It is stronger, namely that the performance criteria of the two actions are the same, explicitly that the set of sequences defined α is the same set as those defined by β.

Hereafter, we will sometimes abuse terminology and refer to the *action* α rather than the *action term* α. In these cases, we literally mean the equivalence class of the term α,[4] i.e. the action described by α.

Example 2.2 Consider the example of a footrace about to begin.[5] 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 iff the race has started,
Loaded	true iff 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 term 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 states, so that each combination of atomic predicates is represented. See Fig. 1, in which an arrow $w \rightsquigarrow w'$ denotes that $w' \in [\![\mathsf{load}]\!](w)$ and $w \rightsquigarrow w'$ that $w' \in [\![\mathsf{fire}]\!](w)$. We assume that one cannot load an already loaded gun. Just to make the model more interesting, we assume that our starter pistol may misfire. When a loaded pistol misfires, nothing relevant 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. The interpretation of a number of sample formulas is given in the figure. The reader may confirm the equations in Table 3 for himself.

[4] More precisely, we have in mind the set seq(α) of sequences defined by α. This is our model-independent interpretation of the action described by α, discussed in Sect. 3.

[5] This example is superficially similar to the Yale shooting problem (Hanks and McDermott 1987) but we have a different purpose in mind. Where Hanks and McDermott are primarily interested in solutions to the frame problem, we postpone such considerations and instead concentrate on means–end semantics.

Fig. 1 A sample PDL model

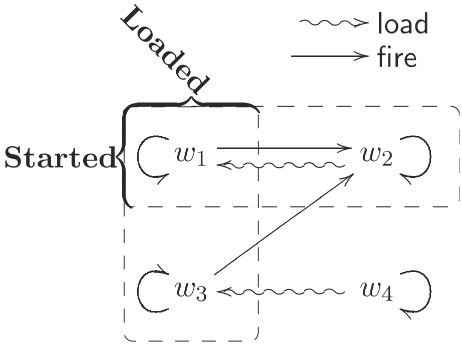

Table 3 Sample facts about the model in Fig. 1

$[\![\mathbf{Started}]\!]$	$= \{w_1, w_2\}$
$[\![\mathbf{Loaded}]\!]$	$= \{w_1, w_3\}$
$[\![[\mathsf{fire}]\mathbf{Started}]\!]$	$= [\![\mathbf{Started}]\!]$
$[\![\langle\mathsf{fire}\rangle\ \mathbf{Started}]\!]$	$= [\![\mathbf{Started}]\!] \cup [\![\mathbf{Loaded}]\!]$
$[\![[\mathsf{load}]\mathbf{Loaded}]\!]$	$= \mathcal{W}$
$[\![\langle\mathsf{load}\rangle\ \mathbf{Loaded}]\!]$	$= [\![\neg\ \mathbf{Loaded}]\!]$
$[\![[\mathbf{Loaded}?; \mathsf{fire}]\ \mathbf{Started}]\!]$	$= [\![\mathbf{Started}]\!] \cup [\![\neg\ \mathbf{Loaded}]\!]$
$[\![\langle\mathbf{Loaded}?; \mathsf{fire}\rangle\ \mathbf{Started}]\!]$	$= [\![\mathbf{Loaded}]\!]$

Table 4 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	$[\alpha \cup \beta]\varphi \leftrightarrow ([\alpha]\varphi \wedge [\beta]\varphi)$
Test	$[\psi?]\varphi \leftrightarrow (\psi \wedge \varphi)$
K	$[\alpha](\varphi \rightarrow \psi) \rightarrow ([\alpha]\varphi \rightarrow [\alpha]\psi)$
	Inference rules
Modus Ponens	$\varphi, \varphi \rightarrow \psi\ /\ \psi$
Necessitation	$\varphi\ /\ [\alpha]\varphi$

To complete our introduction to PDL, we present in Table 4 the standard axiom system for PDL, taken from Harel (1984). For rules of inference, we write φ/ψ to mean: From a proof of φ, infer ψ. We omit the proof that this system is sound and complete for our semantics, i.e. $\vdash \varphi$ iff $\models \varphi$.

2.2 Sufficient means for an end

There are at least three distinct kinds of means–end relations that are relevant for practical reasoning. They are:

weakly sufficient means: doing α *may* realize φ.
(strongly) sufficient means: doing α *will* realize φ.
necessary means: φ will not be realized unless the agent does α.

The different kinds of relations yield different motivational force for the agent who desires φ. In this section, we will provide semantics for the two sufficient means–end relations and sketch the kind of practical consequences they support.

When, we say that an action α is a (*strongly*) *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 impossible in w. If one cannot do α, then surely α is not a means to any end at all.[6] Thus, α is a sufficient means to φ in w if (i) doing α in w ensures that φ and (ii) one *can* do α.

An action α is *weakly sufficient* for φ in w just in case in w doing α *might* realize φ. But this is exactly captured by the weak operator $\langle\alpha\rangle$. Thus, we suggest the following definition.

Definition 2.3 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 propositions 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 $\emptyset \neq [\![\alpha]\!](w) \subseteq [\![\varphi]\!]$ and is weakly sufficient iff $[\![\alpha]\!](w) \cap [\![\varphi]\!] \neq \emptyset$. In case that one wants to realize φ, then one may be sure to do so by performing any sufficient means, but there may be reasons that he chooses not to perform any (strongly) sufficient means, of course. One *cannot* realize φ, however, *without* performing some weakly sufficient means.

Remark 2.4 As Sven Ove Hansson has noted, our definition of sufficient means neglects a common feature of natural language means–end talk: the relevancy of the means to its end. We should not call α a means to φ if *every* action realizes φ. For instance, the action fire is surely not a means to realizing the condition that $1 + 1 = 2$, since that condition is inevitable. It does not depend on anything we can do. Von Wright agrees that relevance of the action is a central feature in means–end relations. Indeed, he describes such relations as *causal*—as mentioned in (4) from the introduction—and fire surely does not cause the mathematical fact.

This feature is central to the stit operators discussed in Horty and Belnap (1995). An agent can see to it that φ if he can perform an action that realizes φ and also can perform an action that *might fail* to realize φ. Without the negative condition, the fact that φ is realized is irrelevant to the agent's actions. In that case, the agent can not see to it that φ; instead, φ is simply inevitable.

We may amend Definition 2.3 by adding a negative condition as well. An obvious choice is to require that α is a (weakly/strongly) sufficient means to φ in w only if there is some action $\beta \neq \alpha$ such that $w \models \neg[\beta]\varphi$ (equivalently $w \models \langle\beta\rangle\neg\varphi$). This first approximation is suitable for atomic actions α, if one restricts the quantifier to atomic β.

For complex action terms α, however, this simple fix can yield unintuitive results. Suppose that it is nearly dawn and let α be the action of flipping a switch 30 times,

6 Note that this is different than saying doing nothing cannot be a means at all. Letting things run their course may end in a desired situation; but the relevant action is "letting things run their course" (or something similar) and *not* some impossible action.

while β is the action of flipping the same switch once. We assume that α takes longer than β and indeed takes long enough that the sun will have risen by the time one has completed α (but not by the time one has completed β). Thus, doing α will result in a state, in which the sun has risen, while doing β will not, i.e. $w \models [\alpha]\mathbf{sunlit}$ and $w \models \neg[\beta]\mathbf{sunlit}$. Thus α passes the above test: it satisfies both the positive and negative conditions. It is not, however, relevant to the outcome: the sunlit state is inevitable if one waits long enough.

In sum, the proposed negative condition would eliminate *some* but not all means–end relations involving actions irrelevant to the end. It seems that further reflection is required to properly express the relevance condition for sufficient means–end relations. We suspect that a solution requires the right relation between our proposed means α and the action term β serving to satisfy the negative condition—in this case, that β takes about as long as α.

The practical consequences of sufficient means are somewhat difficult to analyze. It is not the case that an agent, on pain of practical irrationality, say, should 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 strongly sufficient means. One might try to explain the motivation of (weakly/strongly) sufficient means in terms of defeasible reasons to do α.

Alternatively, some might argue that our agent should give up his end or perform *some* weakly sufficient means. This latter claim is similar to reasoning involving necessary means, since if there are a finite number of weakly sufficient means $\alpha_1, \ldots, \alpha_k$ to φ, then their disjunction $\alpha_1 \cup \cdots \cup \alpha_k$ is a *necessary* means to φ. Generally, the motivational consequences of necessary means have seemed clearer than the consequences of sufficient means.

Thus, many treatments of practical reasoning, including important contribution von Wright's in (1963), 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". (von Wright, 1963)

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.

3 Necessary means and complex actions

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 are more complicated due, in part, to three features of such means.

(A) A necessary means α to φ need not be sufficient. Thus, necessity is not expressed simply in terms of $[\alpha]$ or $\langle\alpha\rangle$.

(B) A necessary means α to φ need not be *immediately* necessary. One may do other things (relevant to φ or not) prior to performing α.
(C) A sequential necessary means $\alpha_1;\alpha_2;\ldots;\alpha_n$ need not be performed "all at once" to realize its end. It may be the case that one can realize φ by performing $\alpha_1;\beta_1;\alpha_2;\ldots;\beta_{n-1};\alpha_n$ without refuting the necessity of $\alpha_1;\alpha_2;\ldots;\alpha_n$.

Features (A) and (B) are discussed explicitly in von Wright (1963). The third feature is not explicit there, but we believe that it is a reasonable feature of necessary means. Let us clarify this feature with a short example.

Suppose that we add an unload action to our footrace model from Example 2.2, so that one may unload a gun without firing it. We believe that in state w_4, it is reasonable to say that load; fire is a necessary means to **Started**. One cannot reach a **Started**-state unless she loads the gun and later fires it. It is in this sense that the sequence is necessary: each action must be done in order (but not necessarily contiguously).

Of course, some actions done between load and fire will interfere with our plans. For instance, if one does load; unload; fire, then she has done load and fire in order, but she cannot reach a **Started**-state this way. But this does not contradict the necessity of load; fire. When, we say that a sequential action is necessary, we mean that the end will not be realized unless each element in the sequence is done. We do *not* mean that if each element is done, then the end will (or even may) be realized: other actions (like unload) can interfere with our aim.

With these features in mind, let us present a rough working definition of necessary means.

Definition 3.1 (*Informal sketch*) An action α is a necessary means to φ in w if the following hold.

1. φ is attainable in w and
2. every weakly sufficient means to φ in w involves α.

Item (1) avoids vacuous necessary means to unattainable ends. Without it, *every* action would be a necessary means to any unattainable end, but surely we do not conclude that an agent desiring an unattainable end ought to do *everything*. The second item in Definition 3.1 depends on the undefined term "involves".[7] Defining this term is our primary duty in the next section.

3.1 Involvement: a sketch

In defining involvement, we must ensure that features (A)–(C) are satisfied. This places certain restrictions on involvement and its interaction with composition. We will begin with an informal sketch for the semantics of involvement. A full technical development of this sketch would require the introduction of rather complicated formal definitions and we try to avoid this as far as possible. Thus, we use the sketch

[7] We find a notion of involvement in Dignum et al. (1994), which the authors indicate can be easily extended to include sequential composition, but the obvious extension does not satisfy our requirements for composition. As well, the authors assume that every action is possible in every state. This assumption would greatly simplify our definition of involvement, but, we believe it would limit our models too much. In particular, one would have to omit the test operator from the language.
Note also that—unlike (Dignum et al. 1994)—we have no synchronous composition constructor. As one of our reviewers pointed out "one could also imagine a natural notion that includes a form of concurrency".

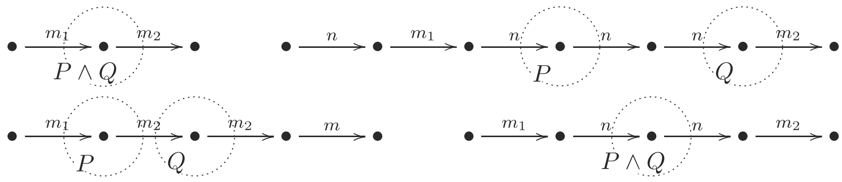

Fig. 2 Four different paths that do $m_1; P?; Q?; m_2$ "along the way"

to motivate an axiomatization of involvement (sound for the intended semantics) which should be accessible to a wider audience. Nonetheless, the resulting development is more complicated than, we would like, but we believe that this is a natural consequence of the features (A)–(C).

In Definition 3.2, when, we say that an action β *involves* an action α, we mean that one cannot do β without doing α "along the way". Let us clarify what, we mean by first supposing that α is just a sequence of atomic actions, say $\alpha = m_1; m_2; \ldots; m_n$. By doing β, one has done α "along the way" if one has performed each of the atomic actions $m_1, m_2, \ldots, m_n$ (in order, but not necessarily contiguously) during his performance of β. In this way, he has done each of the actions necessary to do α and he has done them in order (first m_1, later m_2,...), but he may have done other things before, after or in between. For example, the action described by load; unload; fire involves the action described by load; fire in this sense. In doing load; unload; fire, one does load; fire along the way (in order, but not contiguously).

Generalizing a bit, suppose that α also includes test actions, say $\alpha = m_1; \varphi?; m_2$. Then, one does α along the way if one does m_1, later "does" $\varphi?$ and eventually m_2, but how does one "do" a test action? The answer is fairly obvious: one does a test action $\varphi?$ by passing through a φ-state. But this obvious answer also suggests that involvement should be an indexed relation. Whether β involves α depends, on which state one is in when he does β. Explicitly, we can interpret doing β in w as a set paths(β, w) of paths in $\mathcal{M}$.[8] A path σ does $m_1; \varphi?; m_2$ along the way if σ has an edge labeled m_1, a φ-state w some point after that edge and an edge labeled m_2 some point after w. The action β involves $m_1; \varphi?; m_2$ if every path $\sigma \in \mathtt{paths}\ (\beta, w)$ does $m_1; \varphi?; m_2$ along the way.

As indicated in Fig. 2, we interpret a path σ as doing $\varphi?; \psi?$ along the way if σ includes a φ-state w and a ψ-state w' where w' occurs *no earlier* than w. Thus, if σ does $(\varphi \wedge \psi)?$ along the way, it also does $\varphi?; \psi?$ along the way (and also $\psi?; \varphi?$). Hence, if β involves $(\varphi \wedge \psi)?$ in w, it also involves $\varphi?; \psi?$ (and $\psi?; \varphi?$) in w, but not conversely.

Note that "$\varphi?$ is a necessary means to ψ" is a pretty good approximation of the natural language assertion "φ is a necessary means to ψ," although the latter involves a proposition rather than an action as a means. In both cases, one means that ψ cannot be realized without realizing φ either before or simultaneously.

8 The definition of paths(β, w) is fairly straightforward. For atomic action terms m, paths(m, w) is the set of all arrows $w \stackrel{m}{\to} w'$ such that $w' [\![m]\!](w)$. For test actions $\varphi?$, paths$(\varphi?, w)$ is the singleton set $\{w\}$ consisting of the zero-length path w if $w \models \varphi$ and otherwise is empty. For non-deterministic choice, paths$(\alpha \in \cup \beta, w) =$ paths$(\alpha, w) \cup$ paths(β, w) and for sequential composition, paths$(\alpha; \beta, w)$ consists of the concatenation of all compatible paths in paths(α, w) and paths(β, w). Two paths σ and τ are compatible in this sense if the last state in σ is the first state in τ, and concatenation identifies these states.

More generally, one can represent any action term α as a set seq(α) of sequences of atomic and test action terms.[9] For example, the action term $\alpha = \varphi?; (m \cup n)$ corresponds to the set $\{\langle \varphi?, m\rangle, \langle \varphi?, n\rangle\}$. Indeed, each action can be named by a term, which is a disjunction of sequences of test and atomic action terms: in the case above, the action described by α is also described by $(\varphi?; m) \cap (\varphi?; n)$. Let us call this a *normal form* for α. Each disjunct specifies a way of doing α and one has not done α unless he performs one of the disjuncts. Similarly, one has done α *along the way* iff he has done one of the disjuncts (i.e., an element of $\mathsf{seq}(\alpha)$) along the way.[10] In the case of $\alpha = \varphi?; (m \cup n)$, for instance, a path σ does α along the way just in case σ passes through a φ-state and later has either an m- or an n-labeled arrow. Pictorially, σ does α along the way just in case σ looks like one of the two paths below.

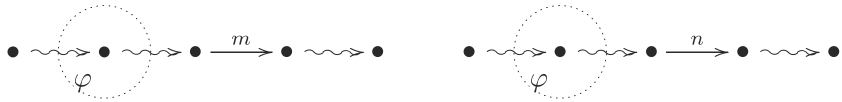

The squiggly arrows here denote any path (including a zero-length path).

Finally, the action β *involves* α in w just in case each $\sigma \in$ paths(α, w) does β along the way.

As, we can see, our definition of involvement includes some subtle features, particularly the asymmetry in the treatment of the two arguments. Whether an action β involves α depends on the state in which β is performed. This is due to the presence of impossible actions generally and in particular our need to evaluate test actions. Thus, we compare the model-independent interpretation of α as a set of sequences to the model- and state-dependent interpretation of β as a set of paths. The action β involves α at w if every way of doing β at w also does α along the way.

Let us write $\beta \rhd_w \alpha$ for "β involves α in w". This relation clarifies our intended meaning of necessary means: α is a necessary means to an attainable end φ in w iff every action β such that $w \models \langle\beta\rangle\varphi$ does α along the way. In fact, as we will see, some of the rules of inference depend on all of the states reachable by some action, so it is more natural to index the involvement relation by sets of states. If $S \subseteq \mathcal{W}$, we write $\beta \rhd_S \alpha$ for $\beta \rhd_w \alpha$ for every $w \in S$ and we define

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

We turn now to some of the basic properties of involvement, so that, we can present axioms for $\rhd_S$ in lieu of a complicated presentation of our formal semantics.

3.2 Involvement axiomatized

In this section, we present a first axiomatization of the involvement relation introduced above. Our aim is to sketch a few initial results, which clarifies the relation and

[9] The definition of $\mathsf{seq}(\alpha)$ can also be made explicit. For atomic and test action terms, $\mathsf{seq}(\alpha) = \langle\alpha\rangle$. For choice, $\mathsf{seq}(\alpha \cup \beta) = \mathsf{seq}(\alpha) \cup \mathsf{seq}(\beta)$ and for composition, $\mathsf{seq}(\alpha; \beta)$ is the set of all concatenations $s * t$ where $s \in \mathsf{seq}(\alpha)$ and $t \in \mathsf{seq}(\beta)$. Indeed, it is natural to define $\mathsf{seq}(\alpha)$ first and define $\mathsf{paths}(\alpha, w)$ in terms of $\mathsf{seq}(\alpha)$. One may regard $\mathsf{seq}(\alpha)$ as a model-independent interpretation of the action described by α while $\mathsf{paths}(\alpha,–)$ represents the action as interpreted in a model.
In interpreting action terms as sequences, we are following in the traditions of Meyer (1989).

[10] Explicitly, then, a path σ does a singleton sequence $\langle m\rangle$ along the way just in case there is an m-labeled arrow in σ and it does $\langle\varphi?\rangle$ along the way just in case there is a φ-state in σ. Also, if s and t are non-empty sequences, σ does the concatenation $s * t$ along the way just in case there are compatible paths σ_1 and σ_2 such that $\sigma = \sigma_1 * \sigma_2, \sigma_1$ does s along the way and σ_2 does t along the way.

Table 5 The deductive system for $\rhd_S$

Axioms for involvement		
$\beta \rhd_W \beta$	$\beta \rhd_W \top ?$	$\varphi ? \rhd_S \psi ?$ if $S \models \varphi \wedge \psi$
$\beta \rhd_S \alpha$ if $S \models [\beta] \perp$		
Rules of inference for involvement		
$\frac{\beta \rhd_S \alpha}{\top ?; \beta \rhd_S \alpha}$		$\frac{\beta \rhd_S \alpha}{\beta \rhd_S \top ?; \alpha}$
$\frac{\beta \rhd_S \alpha}{\beta; \top ? \rhd_S \alpha}$		$\frac{\beta \rhd_S \alpha}{\beta \rhd_S \alpha; \top ?}$
$\frac{\beta \rhd_S \alpha}{\beta; \gamma \rhd_S \alpha}$		$\frac{\beta \rhd_S \alpha; \gamma}{\beta \rhd_S \alpha}$
$\frac{\beta \rhd_{[\gamma](S)} \alpha}{\gamma; \beta \rhd_S \alpha}$	$\frac{\beta \rhd_S \alpha \gamma \rhd_{[\beta](S)} \delta}{\delta \beta; \gamma \rhd_S \alpha; \delta}$	$\frac{\beta \rhd_S \alpha; \gamma; \delta}{\beta \rhd_S \alpha; \delta}$
		$\frac{\beta \rhd_S \gamma; \alpha}{\beta \rhd_S \alpha}$
$\frac{\beta \cup \gamma \rhd_S \alpha}{\beta \rhd_S \alpha}$	$\frac{\beta \rhd_S \alpha \gamma \rhd_S \alpha}{\beta \cup \gamma \rhd_S \alpha}$ $\frac{\beta \rhd_S \alpha}{\beta \rhd_S \alpha \cup \gamma}$	
$\frac{\varphi ?; \psi ? \rhd_S \alpha}{(\varphi \wedge \psi) ? \rhd_S \alpha}$	$\frac{(\varphi \wedge \psi) ? \rhd_S \alpha}{\varphi ?; \psi ? \rhd_S \alpha}$	$\frac{\beta \rhd_S (\varphi \wedge \psi) ?}{\beta \rhd_S \varphi ?; \psi ?}$
$\frac{\beta \rhd_T \alpha}{\beta \rhd_S \alpha}$ if $S \subseteq T$	$\frac{\beta \rhd_S \alpha \beta \rhd_T \alpha}{\beta \rhd_{S \cup T} \alpha}$	$\frac{\beta \rhd_S \alpha}{\gamma \rhd_S \delta}$ if $\beta \equiv \gamma, \alpha \equiv \delta$

hence our definition of necessary means (to which we return in the next section). The theory presented in Table 5 is sound but not complete with respect to the semantics discussed in Sect. 3.1. We hope that later revisions will produce a more elegant and complete axiomatization.

Let us begin our discussion with a surprising result: involvement is *not* transitive. An action β may involve α in w and α involve γ in w, while β does not involve γ in w. This is a result of the basic asymmetry in the treatment of the arguments: $\beta \rhd_w \alpha$ says something about β-paths at w and α-sequences. But $\alpha \rhd_w \gamma$ expresses something about α-paths at w, not α-sequences, and so these two facts do not ensure $\beta \rhd_w \gamma$. For a concrete counterexample, consider the following model:

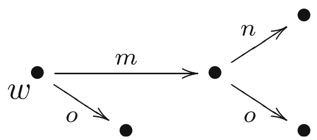

In w, the action n is impossible, so the action $n \cup o$ involves o: one cannot do $n \cup o$ in w without doing o. Also, the action $m; (n \cup o)$ involves $n \cup o$ (in w and elsewhere). But it is not the case that $m; (n \cup o)$ involves o at w: one can do the former action without doing o along the way since one can do $m; n$. Thus, we see that $m; (n \cup o) \rhd_w n \cup o$ and $n \cup o \rhd_w 0$ but $m; (n \cup o) \not\rhd_w o$, i.e. $\rhd_w$ is not transitive.

Because $\rhd_w$ is not transitive, our axiomatization requires a large number of inference rules. We cannot rely on transitivity to carry the consequences of axioms forward, so to speak.

Our short list of axioms is fairly self-explanatory, with one exception, namely, $\beta \rhd_S \alpha$ for all α whenever β is impossible in S. In other words, an impossible action vacuously involves every other axiom. This axiom can be understood as a vacuous implication: if one does an impossible action, he does every other action along the way (because he *cannot* do an impossible action). In terms of our path-sequence semantics, if β is impossible in w, then paths (β, w) is empty and so β vacuously involves α.

A natural consequence of this axiom can be found in our discussion of the counterexample to transitivity. As our rules indicate, if $\beta \rhd_S \alpha$ and $\gamma \rhd_S \alpha$ then $\beta \cup \gamma \rhd_S \alpha$. Now, because n is impossible in w, we see that $n \rhd_w o$. Trivially $o \rhd_w o$, too, so (as we claimed) $n \cup o \rhd_w o$. If one cannot do n, then he cannot do $n \cup o$ without doing o.

Our rules of inference include nine rules involving sequential composition. Most of these come with a left-hand and right-hand version, which inflates our list somewhat. This includes four trivial rules regarding composition with $\top$? and three right-hand "weakening" principles:

$$\frac{\beta \rhd_S \alpha; \gamma}{\beta \rhd_S \alpha}, \qquad \frac{\beta \rhd_S \alpha; \gamma; \delta}{\beta \rhd_S \alpha; \delta}, \qquad \frac{\beta \rhd_S \gamma; \alpha}{\beta \rhd_S \alpha}.$$

More importantly, we include three rules related to condition (C), namely:

$$\frac{\beta \rhd_S \alpha}{\beta; \gamma \rhd_S \alpha}, \qquad \frac{\beta \rhd_{[\![\gamma]\!](S)} \alpha}{\gamma; \beta \rhd_S \alpha}, \qquad \frac{\beta \rhd_S \alpha \quad \gamma \rhd_{[\![\beta]\!](S)} \delta}{\beta; \gamma \rhd_S \alpha; \delta}.$$

The first two of these rules are the left-hand analogues of the right-hand weakening rules and the third expresses that $\rhd_w$ is monotonic, in a certain sense. Let us sketch how the rules relate to condition (C) (although, we use only two rules for this, we are free to pick, which left-hand weakening rule to apply). This condition essentially says (in light of Definition 3.1) that $\alpha_1; \beta_1; \alpha_2; \ldots; \beta_{n-1}; \alpha_n$ should involve $\alpha_1; \alpha_2; \ldots; \alpha_n$. Indeed, this is easy to see: because $\alpha_1 \rhd_{\mathcal{W}} \alpha_1$, we may apply the rule

$$\frac{\beta \rhd_S \alpha}{\beta; \gamma \rhd_S \alpha}$$

to conclude $\alpha_1; \beta_1 \rhd_{\mathcal{W}} \alpha_1$. Similarly, $\alpha_2; \beta_2 \rhd_{\mathcal{W}} \alpha_2$ and hence $\alpha_2; \beta_2 \rhd_{[\alpha_1;\beta_1](\mathcal{W})} \alpha_2$. Thus by the rule

$$\frac{\beta \rhd_S \alpha \quad \gamma \rhd_{[\![\beta]\!](S)} \delta}{\beta; \alpha \rhd_S \gamma; \delta},$$

we conclude $\alpha_1; \beta_1; \alpha_2; \beta_2 \rhd_{\mathcal{W}} \alpha_1; \alpha_2$ and so on.

The rules for $\cup$ are fairly straightforward, as are the two structural rules

$$\frac{\beta \rhd_T \alpha}{\beta \rhd_S \alpha} \text{ if } S \subseteq T, \qquad \frac{\beta \rhd_S \alpha \beta \rhd_T \alpha}{\beta \rhd_{S \cup T} \alpha}.$$

The rule regarding $\equiv$ (introduced in Remark 2.1) depends on an easy theorem: if $\beta \equiv \gamma$ then $\mathsf{seq}(\beta) = \mathsf{seq}(\gamma)$ (up to provably equivalent test actions) and also $\mathsf{paths}(\beta, -) = \text{paths}(\gamma, -)$.

Finally, we have the two left-hand and one right-hand rule for $(\varphi \wedge \psi)$?. The left-hand rules simply reflect that $\mathsf{paths}((\varphi \wedge \psi)?, -) = \mathsf{paths}(\varphi?; \psi?, -)$. There is an asymmetry for the right-hand rule: if β involves $(\varphi \wedge \psi)$? in w, then every β-path rooted at w passes through a $\varphi \wedge \psi$-state. Consequently, every such path passes through a φ-state and a ψ-state (which is *not earlier* than at least one of the φ-states). Thus, we see that

$\beta \rhd_w \varphi?;\psi?$, but the converse rule does not hold. An action β that involves $\varphi?;\psi?$ in w does not necessarily involve $(\varphi \wedge \psi)?$.

A final word about involvement: one should note that $\rhd_w$ is a relation in our meta-theory. We do not introduce this partial order in the syntax of PDL, because we see no convenient means of extending the language to do so. Thus, our definition of necessary means remains in the meta-theory, unlike the definitions of sufficient means in Sect. 2.2. But our goal here is to provide a semantics for necessary means to evaluate practical syllogisms and meta-theoretical definitions will suffice for this.

3.3 Necessary means summarized

We are now prepared to give our definition of necessary means.

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 $\beta \rhd_w \alpha$.
2. For every $\cup$-free action β, if $w \models \langle\beta\rangle\varphi$ then $\beta \rhd_w \alpha$.

Each item above expresses explicitly the corresponding item from Definition 3.1, but (2) includes an additional technical constraint: it applies only to "$\cup$-free" actions β. An action β is $\cup$-free if it can be described without using the choice constructor $\cup$, i.e. if there is some term α such that (1) α is constructed from atomic and test actions using sequential composition only and (2) $\alpha \equiv \beta$. Without this constraint, actions $m \cup n$ would trivially refute the necessity of m as a means to an end since $m \cup n$ does not in general involve m.

The restriction to $\cup$-free β avoids this kind of trivial refutation, but one may worry that the restriction is too strong. Have, we thrown out too many counterexamples by restricting to $\cup$-free actions? We claim this is not the case, for suppose there is a weakly sufficient means β to φ. Without loss of generality, we may suppose that β is in normal form, say $\beta = \beta_1 \cup \beta_2 \cup \cdots \cup \beta_n$ where each β_i is $\cup$-free. Then at least one of the β_i's is also weakly sufficient for φ. If every weakly sufficient β_i involves α, then β was a faux counterexample: it did not represent a way to realize φ without doing α. On the other hand, if some β_i does not involve α, then β_i satisfies condition (2) and thus refutes the necessity of α, as expected.

To put it another way, one can rephrase the conditions of Definition 3.2 as follows:

1. There is a path σ such that
   (a) the first state of σ is w;
   (b) the last state of σ is a φ-state;
   (c) σ does α along the way.
2. Every path σ satisfying (1a) and (1b) also satisfies (1c).

This equivalent formulation seems right for necessary means. It says, in essence, there is a way to a φ-state from w and every such way involves α.

Because our counterexample actions in (2) occur only $\cup$-free actions, we could have modestly simplified the axioms from Sect. 3.2. We do not need the two axioms involving $\cup$ on the left in order to define necessary means, but we believe that our notion of involvement is more naturally stated in general terms. In any case, most of our complications are due to impossible actions (including test actions) and *not* non-deterministic choice.

This completes our semantics for necessary means. It is unfortunate that this analysis has introduced such a complicated notion (involvement), but we believe that this is largely unavoidable. In order to understand necessity, we need an analysis of when doing one action entails doing another (the proposed necessary means). This is the relation we called involvement. But because of von Wright's motivating work, we are not limiting ourselves to immediate necessity nor to plans which must be executed without interruption. Thus, the notion of involvement is subtle and the resulting semantics and axiomatization a bit more complicated than one might like. This outcome is indeed striking: practical reasoning involving necessary means seems much simpler than that involving sufficient means, but the situation is reversed in the semantics. Sufficient means are easy to understand, while necessary means are comparatively difficult.

Example 3.3 We return to the footrace from Example 2.2 and investigate some sufficient and necessary means for starting the race.

For w_1 and w_2, the action $\top?$ is a necessary means for **Started** since any action involves $\top?$.[11] Furthermore, *any* (not impossible) action, $\top?$ included, is a sufficient means for **Started** in these states since no actions lead to states, in which **Started** is false. Thus, in w_2, the action load is a sufficient means for **Started**, since it leads to state w_1 but in w_1, load is not sufficient since it is impossible there.

For w_3 and w_4, there is *no* (strongly) sufficient means for starting the race, since the possibility of misfire precludes any guarantee that **Started** will be realized. In w_4, the composite load; fire is weakly sufficient, since it first leads to w_3 where fire is weakly sufficient. In w_3 and w_4, the action fire is a necessary means to **Started**, and in w_4, the action load is also a necessary means, as is the sequence load; fire.

Remark 3.4 We said in Remark 2.1 that there are terms α and β such that (i) for all $\varphi, \vdash [\alpha]\varphi \leftrightarrow [\beta]\varphi$ and (ii) α is a necessary means to some ψ, while β is not. Indeed, consider the following PDL model:

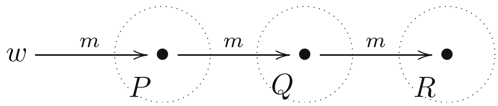

It can be shown that the action $P?; Q?$ is a necessary means for R (as is the action $m; P?; m; Q?; m$ and so on). On the other hand, $(P \wedge Q)?$ is *not* necessary for R—indeed, it is an impossible action in every state in this model! Clearly, $P?; Q?$ and $(P \wedge Q)?$ satisfies condition (i) above.

4 Adding conditionals

Our semantics so far has involved *local* means–end relations: they have expressed causal relations between states (and hence each definition included a state w as a parameter). But this is a very narrow sense of means–end relations. Its primary advantage is simplicity (*relative* simplicity in the case of necessary means!), but it is in many ways insufficient as a semantics for natural language means–end relations. There are at least three shortcomings of local means–end relations as follows:

[11] This holds generally: if φ is attainable in w, then $\top?$ is trivially a necessary means to φ in w. Admittedly, there is no natural language analogue to make this trivial ascription plausible, but we accept it as an artifact of our formalization.

(1) Local means–end relations are not useful in understanding the reasoning of epistemically limited agents. An agent may know only some of the facts that hold in the world currently. In this case, he will not know, which state is the actual state, although he may be able to eliminate certain states from consideration.
(2) Local means–end relations express a causal relation only about one particular state, but means–end relations in natural language are intended more broadly.
(3) Local means–end relations obscure the important role of intermediate ends.

This last limitation is of particular importance to us and is relevant to feature (2) from the introduction. An intermediate end is an end adopted by the agent primarily so that he can achieve another (primary or intermediate) end. For instance, an agent in state w_4 from Example 2.2 knows that fire is a necessary means to **Started**, but it is not weakly sufficient unless **Loaded** holds. Thus, he adopts **Loaded** as an intermediate end. Such reasoning is the subject of much interest in the A.I. literature and is treated in some detail in Pollock (2002) and elsewhere.

It is fairly common to call such intermediate ends *means* (either to a final end or to other intermediate ends). For example, an academic degree is sometimes called a means to better employment as well as an end. We believe that it is better to clearly distinguish the two basic types here: a means is—as says in von Wright (1963)—something one does and an end is a condition to be brought about. In this analysis, then, a bachelor's degree is not a means (since it is not an action that causes a desired condition), but it is a *precondition* for a related means–end relation; hence it is an *intermediate* end. There is a natural tension between this account and feature (2)—an end can be a means in different context—but, we believe that this is a useful and clarifying distinction between means and ends. An end is something one aims for (perhaps to further a later goal) and a means is what one does.

Intermediate ends are a basic feature of practical reasoning because our means–end relations are neither *local* nor *global*[12] but *conditional:* given **Loaded**, the action fire is weakly sufficient for **Started**. Reasoning about intermediate ends is fraught with difficulties and these difficulties should be reflected in our semantics. In particular, we are interested in the *ramification* problem: the problem of indirect consequences of action. We motivate our interest via an example, we call *The Shortsighted Suitor*.

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

This argument can be represented thus:

[12] A *global* means–end relation is a local means–end relation that is true in every state.

$$\frac{\begin{array}{l}\textbf{Money} \Rightarrow \langle \mathsf{ask} \rangle \textbf{Married} \\ \langle \mathsf{rob} \rangle \textbf{Money}\end{array}}{\langle \mathsf{rob}; \mathsf{ask} \rangle \textbf{Married}}.$$

The argument fails, of course, because if I rob my sweetheart, she will hate me (let us assume such a clueless suitor will not mask his identity). We assume that

$$\textbf{Hate} \Rightarrow \neg\langle \mathsf{ask} \rangle \textbf{Married}.$$

But this conditional is inconsistent with our first premise, unless we use a non-monotonic conditional. If our conditional is monotonic, then the first premise implies $(\textbf{Money} \wedge \textbf{Hate}) \Rightarrow \langle \mathsf{ask} \rangle \textbf{Married}$ and our assumption yields $(\textbf{Money} \wedge \textbf{Hate}) \Rightarrow \neg\langle \mathsf{ask} \rangle \textbf{Married}$. Since, we also suppose $\langle \mathsf{rob} \rangle (\textbf{Money} \wedge \textbf{Hate})$, we would reach inconsistency.

Thus, in order to represent intermediate ends, we add a conditional operator to our language and in order to represent some of the well-known problems of practical reasoning, we allow non-monotonicity. The literature on such conditional operators is broad, but we hope that a few simple definitions will satisfy our purposes. At present, we value flexibility over logical commitments. We propose the following (tentative) semantics for our conditional operator. We add to our PDL models a "relevance" function r mapping pairs (w, S) to a set $T \subseteq S$ of "normal" S-states (from the perspective of w), explicitly $r: \mathcal{W} \times \mathcal{P}\mathcal{W} \rightarrow (\mathcal{P}\mathcal{W})$ satisfying the constraint that for every state w and set $S \subseteq \mathcal{W}$,[13]

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

We interpret $r(w, S)$ to be the set of S-states that are reasonably close to (or *normal* from the perspective of) w. We call $r(w, S)$ the *normal*$_w$ *S-states*. The idea is similar to the minimal-change or small-change conditionals discussed in Nute (1984), but one important difference is that, we do not require that $w \in r(w, S)$ if $w \in S$—there is no requirement that the actual state is "normal".

Our conditionals are intended to capture a sense of normality: $w \models \psi \Rightarrow \varphi$ iff, *normally*, given ψ, φ is true, but the sense of "normally" may depend on the state w. We extend the semantics of Sect. 2.1 to include

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

Thus, $\psi \Rightarrow \varphi$ evaluates to true at w iff all the normal$_w$ ψ-states also satisfy φ.

Our models satisfy the following axioms and inference rules shown in Table 6, taken from Nute (1984, 1994). This list is not minimal: axioms CC and CM, for instance, are derivable from the remainder. Again, a rule of inference φ/ψ should be read: if φ is provable (and hence true in every state), then ψ is provable (and hence true in every state).

Clearly, one would like a more thorough investigation 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 present bias for flexibility.

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

Table 6 Logical properties of $\Rightarrow$

Axioms		
ID:	$\varphi \Rightarrow \varphi$	
CC:	$((\psi \Rightarrow \varphi) \wedge (\psi \Rightarrow \chi)) \rightarrow (\psi \Rightarrow (\varphi \wedge \chi))$	
CM:	$(\psi \Rightarrow (\varphi \wedge \chi)) \rightarrow ((\psi \Rightarrow \varphi) \wedge (\psi \Rightarrow \chi))$	
	Inference rules	
RCEC:	$\varphi \leftrightarrow \chi \;/\; (\psi \Rightarrow \varphi) \leftrightarrow (\psi \Rightarrow \chi)$	
RCK:	$(\varphi_1 \wedge \ldots \wedge \varphi_n) \rightarrow \chi \;/\; ((\psi \Rightarrow \varphi_1) \wedge \ldots \wedge (\psi \Rightarrow \varphi_n)) \rightarrow (\psi \Rightarrow \chi)$	$(n \geq 0)$
RCEA:	$\psi \leftrightarrow \chi \;/\; (\psi \Rightarrow \varphi) \leftrightarrow (\chi \Rightarrow \varphi)$	
RCE:	$\psi \rightarrow \varphi \;/\; \psi \Rightarrow \varphi$	

5 Sufficient and necessary pre-conditions

We have seen that in the presence of non-monotonic conditionals we should read a sentence $\psi \Rightarrow [\alpha]\varphi$ as, "*Normally*, ψ implies that doing α will realize φ". Combining the previous results, in this section, we will use such conditional formulas to define conditional means–end relations.

Definition 5.1 In a state w, an action α *is* (*normally*) *a sufficient means to* φ *given* ψ if

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

In w, α *is* (*normally*) *a weakly sufficient means to* φ *given* ψ if

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

Finally, α *is (normally) a necessary means to* φ *given* ψ if the following hold:

1. there is some $w' \in r(w, [\psi])$ and $\beta \rhd_{w'} \alpha$ such that $w' \models \langle\beta\rangle\varphi$;
2. for every $\cup$-free action β and state $w' \in r(w, [\psi])$, if $w' \models \langle\beta\rangle\varphi$ then $\beta \rhd_{w'} \alpha$.

The formula ψ in above definitions is a *sufficient precondition* for the (sufficient or necessary) means–end relation involving α and φ.

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 indication as to why such reasoning is naturally defeasible.

Let us sketch where our shortsighted suitor went awry. We may suppose that in *every* state satisfying **Hate**, we have $\neg\langle \mathsf{ask}\rangle$ **Married**, i.e.,

$$\mathcal{M} \models \textbf{Hate} \rightarrow \neg\langle \mathsf{ask}\rangle\textbf{Married}$$

(where $\rightarrow$ is material implication, as usual). We further suppose that in the current state w, it is the case that every *normal* **Money**-state satisfies $\langle \mathsf{ask}\rangle$ **Married**, that is,

$$r(w, [\textbf{Money}]) \subseteq [\langle \mathsf{ask}\rangle\textbf{Married}].$$

Thus, $w \models$ **Money** $\Rightarrow \langle \mathsf{ask}\rangle$**Married**. These suppositions jointly satisfy our analysis from Sect. 4 and lead to no contradictions, provided that $r(w, [\![\textbf{Money}]\!])$ is disjoint from [**Hate**], i.e. that normal$_w$ **Money**-states are not **Hate**-states. Thus, we see how our non-monotonic conditional supports phenomena like the ramification problem.

Table 7 A summary of our means–end relations

	α is a weakly sufficient means for φ	α is a sufficient means for φ	α is a necessary means for φ
Unconditionally	$w \models \langle\alpha\rangle\varphi$	$w \models [\alpha]\varphi \wedge \langle\alpha\rangle\top$	$\exists\beta \rhd_w \alpha,$ $w \models \langle\beta\rangle\varphi.$ $\forall\cup\text{-free}\beta \not\rhd_w \alpha$ $w \not\models \langle\beta\rangle\varphi.$
ψ is a sufficient precondition for it to be true that:	$w \models \psi \Rightarrow \langle\alpha\rangle\varphi$	$w \models \psi \Rightarrow$ $([\alpha]\varphi \wedge \langle\alpha\rangle\top)$	$\exists w' \in r(w, [\psi])$ $\exists\beta \rhd_w \alpha,$ $w' \models \langle\beta\rangle\varphi.$ $\forall w' r(w, [\psi])$ $\forall\cup\text{-free}\beta \not\rhd_w \alpha,$ $w' \not\models \langle\beta\rangle\varphi.$
ψ is a necessary precondition for it to be true that:	$w \models \langle\alpha\rangle\varphi \Rightarrow \psi$	$w \models ([\alpha]\varphi \wedge \langle\alpha\rangle\top) \Rightarrow \psi$	n/a

Let us now turn to *necessary preconditions* for means–end relations. It is easy to define this concept for sufficient means–end relations: ψ is a necessary precondition for α to be a sufficient means to φ, just in case *normally* α is a sufficient means for φ only if ψ.

Definition 5.2 We say that ψ is a *necessary precondition* for α to be a (weakly, resp.) sufficient means for φ in a state 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 yet found a suitable corresponding definition of "necessary preconditions for necessary means–end relations", but this notion does not arise as naturally in means–end talk as the other conditional means–end relations.

In Table 7 we summarize our taxonomy of means–end relations. To illustrate the flexibility and consequences of our conditional means–end relations, we return to the footrace and starter pistol example, and define three different relevance functions which may be added to the original example. In the third case, the relevance function is also accompanied by new states, in which the gun is malfunctioning.

Example 5.3 The relevance function may be used to interpret the conditional as material implication, to express epistemic limitations or to express the abnormality of complications like broken artifacts.

Strict implication: First, we may define $r_s(w, S) = S$ for every set $S \subseteq W$ and $w \in \mathcal{W}$. In that case, the conditional connective $\Rightarrow$ coincides with strict implication (the subscript s stands for "strict implication"). Thus, in every state 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.

Epistemic limitation: Second, we may define r so that it reflects epistemic limitations of our agent. In our case, we suppose that the agent knows whether the race has

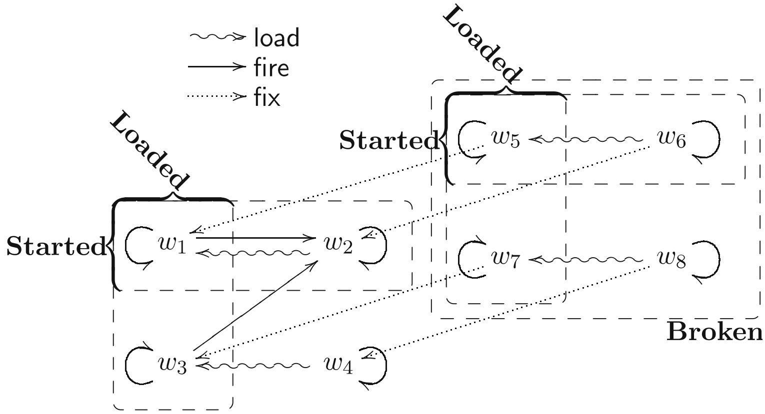

Fig. 3 Additional states for the footrace model

started or not, but he does not know, whether the gun is loaded. A state w is more relevant than another state w' if w is consistent with his current knowledge and w' is not. As a result, in w_4 (say), he regards w_3 as more relevant than w_1 or w_2 and as relevant as w_4 itself. Our approach here is essentially a "nearest relevant states" system. In interpreting $\varphi \Rightarrow \psi$ in w, we consider only the set of states *most relevant* to w that satisfy φ. Hence, we define

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

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

In this example, one can show that in w_4, fire is a necessary means to **Started** given the trivial precondition $\top$. This is not true for r_{s}, since $r_{\mathsf{s}}(w_4, [\top]) = \mathcal{W}$ and fire is not a necessary means to **Started** in w_1 or w_2.

Broken gun: For the third part of the example, we 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 states and transitions to our model as presented in Fig. 3; we also include the new action fix.

We may suppose that the gun is not "normally" broken, regardless of which state is the actual state. Thus, we define for all $w \in W$

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

With this definition, we assume that *even in states, in which the gun is broken,* it is not "normally" broken. For instance, $w_5 \notin r_b(w_5, [\![\top]\!])$.

This model agrees with the material implication conditional defined by r_{s} whenever the antecedent ψ satisfies either $[\![\psi]\!] \subseteq [\![\textbf{Broken}]\!]$ or $[\![\psi]\!] \subseteq [\![\neg\textbf{Broken}]\!]$. It differs from r_{s} just when ψ includes both **Broken** and $\neg$ **Broken** states. The model also satisfies that, in every state w, fix is a necessary means to **Started**, given **Broken**. Moreover, in every state, fire is a weakly sufficient means to **Started** given **Loaded**, since then the "normal"-**Loaded** states regarding w are $r_b(w, [\![\textbf{Loaded}]\!]) = \{w_1, w_3\}$, and in those states fire might lead to starting the race.

We have shown the considerable flexibility of our conditional semantics in these three examples. Of course, it may be that our restrictions are currently too loose—and the conditional semantics *too* flexible—to capture important features of conditional means–end relations. Nonetheless, we prefer to err on the side of flexibility for now.

This concludes our initial presentation of a semantics for means–end relations. In this paper, we have paid particular attention to the first five items on our list of features of means–end relations. In particular, we have argued that a proper analysis of means and ends involves taking the former as actions in a dynamic logic and the latter as formulas. Thus, since the two types are distinct (as required by (1)), it is not literally the case that an end may also be a means, as alleged in (2). This "feature" is an artifact of natural sloppiness in informal language, in our opinion. We have worked to distinguish sufficient means from necessary in Sects. 2 and 3, in order to fulfill (3). We have touched on the causal impact of means (feature (4), especially regarding the relevance condition in Remark 2.4), although clearly more could be said. Finally, we have selected a non-monotonic conditional operator so that conditional means–end relations are defeasible, as anticipated by (5).

6 Concluding remarks

We have introduced a formal semantics for means–end relations which, we hope, approximates the meaning of certain premises in practical syllogisms. Our primary aim is not formalization but clarification, but we prefer the relatively clear consequences of a formal semantics. Admittedly, this makes the approximation a bit rougher, since natural meanings are only crudely represented by formalisms.

Our own work misses its mark in a few ways, which we hope are reparable. First (as noted by our reviewer), human actions are not much like computer programs: there is no clear set of atoms, from which other actions are built. Perhaps our set of action terms should more closely resemble the alternative means humans consider, but the set of alternatives is very much sensitive to context and expertise (see e.g. Bratman (1983)). In this respect, our formalization is not very similar to practical reasoning and it is indeed difficult to see how to fix this while remaining in PDL.

We would also like to revisit our decision to distinguish means and ends as separate types. One could consider actions such as "bring about φ" so that intermediate ends could also count as means. Indeed, even final ends would correspond to actions in this way and one could reason about actions which "realize" these kinds of actions. One hopes that this approach would yield an analysis of constitutive ends (discussed e.g. in Schmidtz (1994)).

Finally, means–end reasoning often depends on temporal constructions (including "while", "until" and others). Moreover, one often performs an act now for an effect to be realized later (such as calling for dinner reservations). This kind of delayed gratification is only awkwardly represented in PDL. It may be possible to introduce the temporal constructions by enriching our syntax and semantics (so that it is closer to the modal μ-calculus), but even then issues of delayed gratification will most likely prove subtle.

Aside from these issues are the features in our introduction. We would like in particular to address objects-as-means in our future work.

There are, therefore, many outstanding issues in our semantics of means–end relations. We regard this as a fruitful first analysis which is so far surprisingly absent in the literature.

Acknowledgements Our work and presentation has been greatly influenced by conversations with Frank Dignum, Sven Ove Hansson, John-Jules Meyer and Lambér Royakkers. An earlier version of this material was presented at the 2005 meeting of the Society for Exact Philosophy, where we received several helpful comments. We would also like to thank our reviewers for their remarkably detailed and insightful comments.

References

Bratman, M. (1983). Taking plans seriously. *Social Theory and Practice*, *9*, 271–287.

Bratman, M. E., Israel, D. J., & Pollack, M. E. (1988). Plans and resource-bounded practical reasoning. *Computational intelligence*, *4*(4), 349–355.

Brown, M. A. (2005). Means and ends in branching time. *Presented at the norms, reasoning and knowledge in technology workshop*.

Castilho, M., Herzig, A., & Varzinczak, I. (Eds.) (2002). It depends on the context! A decidable logic of actions and plans based on a ternary dependence relation. *9th Intl. workshop on non-monotonic reasoning NMR'2002*.

Castilho, M. A., Gasquet, O., & Herzig, A. (1999). Formalizing action and change in modal logic I: the frame problem. *Journal of logic and computation*, *9*(5).

Dignum, F., Meyer, J.-J. Ch., & Wieringa, R. (1994). Contextual permission: a solution to the free choice paradox. In A. J. Jones & M. Sergot (Eds.), *DEON'94, Second internation workshop on deontic logic in computer science*, University of Oslo Complex series, pp 107–135.

Giacomo, G. D., & Lenzerini, M. (Eds.) (1995). PDL-based framework for reasoning about actions. In *LNAI 1992*. pp 103–114.

Giordano, L., Martelli, A., & Schwind, C. (2000). Ramification and causality in a modal action logic. *Journal of logic and computation*, *10*(5), 615–662.

Hanks, S., & McDermott, D. (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, S. O. (2000), Formalization in philosophy. *The Bulletin of symbolic logic*, *6*(2), 162–175.

Harel, D. (1984). Dynamic logic. In D. Gabbay, & F. Guenthner (Eds.), *Handbook of philosophical logic*, Vol. II. D. Dordrecht: Reidel Publishing Company, pp 497–604.

Horty, J. F., & Belnap, N. (1995). The deliberative stit: a study of action, omission, ability and obligation. *Journal of philosophical logic*, *24*, 583–644.

Hughes, J., Esterline, A., & Kimiaghalam, B. (2005). Means–end semantics and a measure of efficacy. *Journal of logic, language and information*. Forthcoming.

McCarthy, J. (1999). Concepts of logical AI. `http://www-formal.stanford.edu/jmc/concepts-ai/ concepts-ai.html`.

McCarthy, J., & Hayes, P. J. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer, & D. Michie (Eds.), *Machine intelligence 4*, Edinburg: Edinburgh University Press, pp 463–502.

Meyer, J.-J. C. (1989). Using programming concepts in deontic reasoning. In R. Bartsch, J. van Benthem, & P. van Emde Boas (Eds.), *Semantics and contextual expression*, Dordrecht: Riverton, pp 117–145.

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.

Millgram, E. (2004). Practical Reasoning entry in the online Dictionary of Philosophy of Mind, (Eds.), Chris Eliasmith, `http://philosophy.uwaterloo.ca/MindDict/practicalreasoning.html`.

Nute, D. (1984). Conditional logic. In D. Gabbay, & F. Guenthner (Eds.), *Handbook of philosophical logic*, Vol. II. D. Dordrecht: Reidel Publishing Company, pp 387–439.

Nute, D. (1994) Defeasible logic. In D. Gabbay, C. J. Hogger, & J. A. Robinson (Eds.), *Handbook of philosophical logic*, Vol. III, D. Dordrecht: Reidel Publishing Company, pp 353–395.

Pollock, J. L. (2002). The logical foundations of means–end reasoning. In R. Elio (Ed.), *Common sense, reasoning and rationality*, Oxford: Oxford University Press.

Prendinger, H., & Schurz, G. (1996). Reasoning about action and change. a dynamic logic approach. *Journal of logic, language and information*, *5*(2), 209–245.
Schmidtz, D. (1994). Choosing ends. *Ethics*, *104*(2), 226–251.
Segerberg, K. (1992). Getting started: beginnings in the logic of action. *Studia logica*, *51*, 347–378.
Von Wright, G. H. (1963). Practical inference. *The Philosophical review*, *72*(2), 159–179.
Zhang, D., & Foo, N. (2002). Dealing with the ramification problem in the extended propositional dynamic logic. In F. Wolter, H. Wansing, M. de Rijke, & M. Zakharyaschev (Eds.), *Advances in modal logic*, Vol. 3. Singapore: World Scientific, pp 173–191.
Zhang, D., & Foo, N. Y. (2001). EPDL: a logic for causal reasoning. *Proceedings of the seventeenth international joint conference on artificial intelligence, IJCAI 2001, Seattle, Washington, USA, August 4-10*. pp 131–138.
Zhang, D., & Foo, N. Y. (2005). Frame problem in dynamic logic. *Journal of applied non-classical logics*, *15*(2), 215–239.