Practical Inference-A Formal Analysis

Zwart, Sjoerd; Franssen, Maarten; Kroes, Peter

DOI
10.1007/978-3-319-91029-1_3

Publication date
2018

Document Version
Final published version

Published in
The Future of Engineering

Citation (APA)
Zwart, S., Franssen, M., & Kroes, P. (2018). Practical Inference-A Formal Analysis. In *The Future of Engineering: Philosophical Foundations, Ethical Problems and Application Cases* (Vol. 31, pp. 33-52). (Philosophy of Engineering and Technology). Springer. https://doi.org/10.1007/978-3-319-91029-1_3

Practical Inference—A Formal Analysis

Chapter · July 2018

DOI: 10.1007/978-3-319-91029-1_3

Sjoerd D. Zwart
Delft University of Technology

Maarten Franssen
Delft University of Technology

Practical Inference—a Formal Analysis

Sjoerd Zwart (Delft and Eindhoven Universities of Technology)

Maarten Franssen (Delft University of Technology)

Peter Kroes (Delft University of Technology)

Abstract

Most engineering reasoning in practice is about how to achieve some predetermined end. Despite its paramount importance, this form of reasoning has hardly been investigated in the literature. The aim of this paper is therefore to explore the question to what extent technical norms can be said to have a truth-value, and under what conditions practical inferences are deductively valid. We take technical norms to be sentences of the form 'If you want A, and you are in a situation B, then you ought to do X'. Von Wright's standard example of making a hut habitable is our paradigm for practical inferences, where an obligation to act is deduced from an intention to realize an end, and an empirical constraint on how this end can be achieved. Our instrument of analysis is dynamic logic (PDL), since actions are aimed at changing the world. PDL already suffices to provide truth-conditions for technical norms. To accommodate the obligation in practical inferences we draw on John Jules Meyer's deontic version of PDL. By paraphrasing 'person *P* wants' with 'person *P* imposes an obligation on herself,' we can give a plausible definition of the validity of practical inferences. In the discussion section, we address the issues of the reliability instead of truth-value of technical norms, and of the defeasibility of practical inferences as they occur in engineering practice.

Introduction

Reasoning about means and ends is part and parcel of engineering practice. The aim of this paper is to explore to what extent this kind of reasoning may be deductively valid. Means-ends inferences are a form of what is called in logic 'practical inference' and the question whether this form of reasoning is deductively valid is still an unresolved issue. The locus classicus for a definition and analysis of practical inference is Von Wright's 'Practical Inference' (1963PI) and that is where we start our discussion. The (by now standard) example of a practical inference that he uses is the following:

(PI) A wants to make the hut habitable
Unless the hut is heated it will not become habitable
- -
A must heat the hut

Von Wright characterizes the structure of this practical inference in the following way. The first premise states an end of *action*, that is, "we want to attain the end as a result or consequence of something which we do" (1963PI, p. 160). The second premise states that a causal relationship of a particular kind exists between this end and some action, namely that this action is necessary to realize the end. This makes the action a (necessary) means to the end stated in the first premise. The conclusion, finally, expresses what Von Wright calls a practical necessity, which he describes as a "necessity of doing something under which an agent is, if he is to attain some end *of his own*" (1972, p. 43).[1]

Is this kind of means-end inference logically valid? Intuitively, the answer to this question is not straightforward. The first premise describes a goal or an end of A, the second an empirical necessity and the conclusion a means. How can a practical argument, relating such different kind of elements be conclusive in terms of truth values? What makes practical inferences logically interesting is that a causal necessity between physical facts or events is "transferred" to a practical necessity for an agent to perform an action. In other words, a necessity in the physical world is lifted to a necessity in the intentional domain. This is achieved by combining a premise that

---

[1] In 'Practical inference' Von Wright writes that he deliberately uses the word 'must' instead of 'ought' because it is "somehow stronger" (1963PI, p. 161). His intuitions on this point seem not to have been very clear, however, since elsewhere he downplays this distinction between 'must' and 'ought to' (e.g. 1963NA, p. 101).

describes a state of affairs in the intentional domain – "One wants to make the hut habitable" – with a premise that describes a state of affairs in the physical domain – "Unless the hut is heated, it will not become habitable". The conclusion states a necessity in the intentional domain.[2]

Von Wright maintains that practical inferences are logically conclusive, albeit that they are "logically valid pieces of argumentation in their own right" (1963VG, p. 167). So, practical inferences have to be added to the traditional forms of logically valid reasoning and cannot be reduced to them. He explains their peculiar form of logical validity by invoking the notion of practical necessitation, that is, "necessitation of the will to action through want and understanding" (idem, p. 170). In a practical inference, a want, an understanding and a decision to act are united in such a way that "the practical necessitation of the will to action must at the same time be a logical necessitation" (idem, p. 171). Thus, Von Wright appears to solve the issue of the logical validity of a practical inference by fiat, by introducing a new kind of logical validity. This may be considered an ad hoc solution. Can we do better? We argue that we can. But before going into the details of how this may be done, we have a brief look at Niiniluoto's work on technical norms; he is one of the few people who have taken up Von Wright's work on practical inferences.

In his (1993) paper 'The aim and structure of applied research', Niiniluoto identifies *rules of actions* as forming the results of design science. A problem for this position, however, is how results of this kind can classify as knowledge, given that rules are generally taken to be normative statements. To answer this question, Niiniluoto identifies such rules with *technical norms* as introduced by Von Wright in *Norm and action* (1963NA). The example of a technical norm given by Von Wright, as a match to his stock example of a practical inference, is

(TN) 'If you want to make the hut habitable, you must heat it'.

Von Wright's notion of technical norms forms an interesting parallel to his conception of practical inference, since it is precisely a statement of this sort that would, when combined with the statement 'You want to make the hut habitable', give, by the deductively valid argument of

---

[2] It could be objected that 'habitable' is a concept that involves notions from the intentional domain, but this is not what Von Wright seems to have had in mind, for he writes that the relation between temperature and habitability of a hut "is a causal fact about the living conditions of men" (1963PI, p. 160). If we interpret "living conditions" as "physical living conditions", the second premise articulates a necessity in the physical world. Thus "habitable" should be read here as a physical property rather than a means to get the hut inhabited.

modus ponens, the conclusion 'You must heat the hut'. But the truth of the conclusion follows only if the premises are true, and for that to be the case, they must be of the right kind to have a truth-value. Von Wright was not convinced that technical norms have a truth-value, and perhaps for that reason rejected this approach to settle the validity of a practical inference. In *Norm and action* Von Wright pronounced "the relation of the technical norm to truth and falsehood" to be "not clear to me" (p. 103). Niiniluoto did not follow Von Wright in this. According to Niiniluoto technical norms classify as 'some kind of rule of action' *and* have a truth-value, so that they can be regarded as knowledge, the sort of knowledge produced by design science. He does not, however, explain on what grounds he accepts what Von Wright could not accept. Niiniluoto simply claims, without further clarification, that a technical norm of the form 'If you want to make the hut habitable, you must heat it' is true if and only if a statement expressing an empirical necessity of the corresponding form 'heating the hut is necessary for making the hut habitable' is true (p. 12). If this is accepted, then an argument that has as premises 'You want *A*', 'You are in situation *B*' and 'If you want *A* and you are in situation *B* then you ought to do X' and as conclusion 'You ought to do X' is deductively valid simply for being an instance of modus ponens. Indeed Niiniluoto also classifies statements like the conclusion of this argument as descriptive and either true or false (idem, p. 12).[3]

In the following we chose to address the question whether a practical inference can be characterized as a deductively valid form of reasoning, and under which assumptions it is valid, rather than to be satisfied with declaring it to be a valid kind of syllogism *sui generis*. We interpret the 'want', 'understanding' and 'decision to act' statements such that the syllogism is deductively valid, and we discuss how faithful these interpretations are to the statements as formulated by Von Wright. The basic idea, upon which the validity of the practical inferences is based, runs parallel to validity of standard syllogisms, which is based on intuitive set theory. For instance, if all elements of a set have some property, then if you select one element it must have this property, too. Practical inference appears to follow the same pattern. If all paths from our state of the world to a state in which φ is the case include some world change event α, then any individual path realizing φ will contain α; thus, in those circumstances, an obligation in our state

[3] Niiniluoto additionally considers other forms of technical norms, which correspond to empirical connections between means and ends that are weaker than necessity, e.g. connections where a mean is sufficient but not necessary for an end and connections where a mean is only probabilistically sufficient.

of the world to realize φ implies an obligation to do α. The validity of this reasoning aspect of practical inference rests on the same ground as some validity claims in first-order logic.

We proceed as follows. In section 1 we briefly discuss the complexity of the conceptual framework that is required for representing the various ingredients of practical inference and for formally interpreting the notion of practical necessity and the character of a technical norm. In section 2 we clarify the character of statements that refer to actions, such as technical norms and practical necessities. For this, we put Segerberg's (1980) advice to practice and take recourse to dynamic logic. In the following section we address the normative or intentional aspects of practical inferences and technical norms. We show that in a modified form practical inferences are logically valid. For this we use dynamic deontic logic. In the concluding section we discuss our results and the limitations of our approach regarding practical inference and means-end reasoning in engineering practices.

1 Conceptual Choices

In proceeding to model practical inference, at least three different 'depths' can be chosen for the level of analysis. The simplest formalization does not include actions as a logical category and takes the world to exist in just a single state. It does, however, include a distinction between causal or empirical *necessity*, of the sort that is expressed in natural laws, and practical necessity or *obligation*. We illustrate the various options for formal analysis by various adaptations of Von Wright's stock example (PI). In the simplest version, which we call the *Static-State version*, all explicit reference to action falls out, and so does reference to wanting, since wanting amounts to an anticipation of a change of the state of the world, whereas there is no room for such change in the analysis. Therefore we get:

(PI^S) The hut should be habitable

Unless the hut is heated is will not be habitable

\- -

The hut should be heated

In this version practical inference can be formalized using standard deontic logic.The formal representation of the *Static State* version of practical inference comes out as:

(PI^S) $\quad O(\text{Hab}(h))$

$\Box(\text{Hab}(h) \rightarrow \text{Heat}(h))$

$O(\text{Heat}(h))$

Here $\Box$ is the modal operator signifying empirical necessity and O the modal operator signifying practical necessity or obligation. Further 'Hab' and 'Heat' refer to 'being habitable' and 'being heated', respectively, and '*h*' refers to the hut in question. Note that, because at this level of analysis action is not a distinct logical category 'being heated' is ambiguous between 'being heated by someone who is actively busy with heating the hut' and 'being equipped with a functioning heating installation'.

This formal argument can easily be made valid by adding to the logic a bridging axiom that links the two modal operators. The following axiom, which is widely accepted in deontic logic and called the *principle of deontic logic*, will do:

(B_O) : $\quad$ For any A, B: $\Box(A \rightarrow B) \rightarrow (OA \rightarrow OB)$

Thus, by paraphrasing practical inference in classical modal-deontic logic in this way, causal or empirical necessity is transferred to practical necessity or obligation. Von Wright introduces the notion of *anankastic statements* to express empirical necessities, such as expressed by the second premise. Anankastic statements do not include intentional actions *as intentional*.[4] They may include actions but then, only the physical manifestation of the action matters and not the intention that accompanies the action.

This Static-State approximation of practical inference implemented in classical deontic logic we consider too crude, because it does not include changing states of the world according to ones will or plan. In (PI^S) the inference fails to represent a *practical* necessity, which Von Wright describes as a "necessity of doing something under which an agent is, if he is to attain some end of his own" (1972, p. 43). To come closer to the intended meaning of 'practical inference', minimally it must be possible for the state of the world to change, to accommodate the idea of

[4] Von Wright: "A statement to the effect that something is (or is not) a necessary condition of something else I shall call an anankastic statement." (1963NA, p. 10); "Laws of nature and other anankastic propositions are, on the whole, not concerned with action; but these we have decided not to call 'norms'." (idem, p. 13).

achieving a desired state. One version of Von Wright's practical inference which allows for changing states of the world is:

(PI^D) The hut should become habitable

Unless the hut is heated it will not become habitable

- -

The hut should be heated

This *Dynamic-State version* is an improvement with respect to the Static-State one as it allows us to take account of changes in the actual state of the world(PC^D, first premise), but it is not expressible in standard modal and deontic logic. For this version we need a form of *dynamic logic*,[5] which provides means for an exact paraphrase of anankastic statements.

The most complicated version of practical inference, which does include actions of individuals, we call *Action versions*. Von Wright's paradigmatic inference (PI) presented in the Introduction corresponds to the Action version. A complete formalization requires that individuals may change states of the world by performing goal-directed actions. Beyond this version lie still completer versions of practical reasoning that contain goals, agenda's, planning, deliberations of advantages and disadvantages, knowledge, belief and so forth. In the next section we will sketch a formalization of an Action version of practical inference. In the final Discussion section, we will address the issue to what extent this version can take care of the aspect of wanting or of the goal-directedness of action.

The distinction that we have made here between Static-State, Dynamic-State and Action versions of practical inference also applies to technical norms. Just as for practical inference, the version of a technical norm presented in the Introduction is an Action version of the notion of a technical norm as introduced by Von Wright. The simplest or Static-State version of a technical norm is:

(TN^S) If the hut *has to be* habitable it *should* be heated

[5] Von Wright characterized his work in the 1960s as "a turn in logic away from a traditional interest in what is, the static, to that which comes to be, the dynamic" (1993, p. 30).

The corresponding expression in the language of deontic logic is: O(Hab(*h*)) → O(Heat(*h*)). In this logic, the truth of the anankastic statement □(Hab(*h*) → Heat(*h*)) in combination with the bridge axiom (B_O) is sufficient for the truth of (TN^S).

The intermediate *Dynamic-State version* of Von Wright's technical norm example is:

(TN^C) If the hut should *become* habitable, then it *should be heated*

(TN^C) strongly suggests that at the moment the hut in question is not habitable and some action should be performed to change that situation.

There is a significant difference between Von Wright's formulation of a technical norm and Niiniluoto's formulation, which is (our emphasis):

(TN_N) If you want *A and you believe you are* in situation *B* then you ought to do X.

Niiniluoto adds a clause with an epistemic operator to Von Wright's technical norm, which is understandable but complicates matters considerably. What if your beliefs are wrong and you are actually in situation *C*? Should you then not do *X*? It might be that in *C,* X is also a necessary action for *A*. And what description of *B* should suffice to do X? If Oedipus wanted to become king of Thebe and if it was necessary for that to marry the widowed queen of Thebe, should Oedipus marry Jocasta or should he marry his mother? With an epistemic operator included in technical norms, a complete belief-desire-intention framework, such as KARO[6], is required for an analysis, but such frameworks are too complex for our purposes. As our goal is to formulate a semantics of technical norms as part of as objective as possible engineering rules and know-how, we leave the epistemic side of things out from our account. This is in line with the treatment of Von Wright, who takes the *truth* of an anankastic statement to be what places the practical necessity on the actor, irrespective of whether the actor knows that statement to be true. Our *Action* version of the technical norm reads therefore:

(TN'_N) If you want *A*, and you are in situation *B*, then you ought to do X.[7]

---

[6] cf Meyer et al (1999).

[7] Note that one may also achieve *A* by *satisficing* some design criteria. Herbert Simon's notion does not change our approach fundamentally.

Let us turn to a more detailed formal analysis of practical inferences and technical norms.

2 PDL and Niiniluoto's Technical Norms

With respect to technical norms we saw, in the introduction, that Niiniluoto claims that the (slightly rephrased) technical directive 'if you want *A*, and you are in situation *B*, then you should do α' is true if and only if

(*) Doing α is a necessary cause of *A* in situation *B*.

Our first task is to explicate the definiens (*) of Niiniluoto's proposal. To do so we need at least three types of elements: (*i*) individual states S_i of the world; (*ii*) (performances of) actions α, β etc., which transform one state of the world into another; and (*iii*) sentences φ, ψ etc. that may have different truth-values in different states of the world. With these three types of elements we are able to model the fact that some actions may change the truth-value of a statement and therefore can be seen as being done for *ends*. The three elements are the main ingredients of *dynamic logic*. We use propositional dynamic logic (PDL) to model technical norms and DDeL (a deontic extension of PDL) for the modelling of practical inference.[8] Basing ourselves mainly on Meyer (1988), we present here only those parts of PDL and DDeL that are relevant for explaining technical norms and practical inference.

Within the syntax of PDL, α, β etc., signify the executions or performances of actions[9], which are recursively built from performances of atomic actions *a*, *b,* etc. Actions can be combined in three ways. One is *sequential composition*, expressed as $\alpha ; \beta$, which means that action α is followed by action β. (Rubbing down wood with sandpaper and then painting it.) Another is *joint action*, expressed by $\alpha \,\&\, \beta$, which means that α and β are performed simultaneously (for instance: lubricating a camshaft while moving the pushrods up and down). Finally there is *choice*, expressed by $\alpha \underline{\cup} \beta$, which means that one choose between doing either α or β or possibly both (for example reading a book or listening to music (or both) may turn your evening in a happy one). An action is $\underline{\cup}$-free if it does not contain a choice aspect.

[8] Because we look upon *actions* as a logical category of its own, we do not opt for a STIT logic.

[9] In PDL an action α is always successful – it changes a state in a predefined way.

Secondly, to define obligations we also need the *negation* of an action, expressed as $\sim\alpha$. This notion is fraught with philosophical difficulties, but here we use it in the intuitive sense of refraining from an action, that is, of doing anything (not necessarily nothing) as long as it is not or does not include that action.[10]

Thirdly, PDL contains the expressions $[\alpha]\varphi$ and $\langle\alpha\rangle\varphi$.[11] $[\alpha]$ extends the semantics of the necessity operator □ from modal logic: $[\alpha]\varphi$ is true in *S* – formally $S \vDash [\alpha]\varphi$ – iff φ is true in all states *S´* which are the result of performing α in state *S*, or α cannot be performed in *S*. In the same way, $\langle\alpha\rangle$ extends the possibility operator ◊. $\langle\alpha\rangle\varphi$ is true in *S* – formally $S \vDash \langle\alpha\rangle\varphi$ – iff, as Meyer (1988, p.114) phrases it, "there is some way by doing action α to achieve φ", that is, α is performable in *S,* and its performance does not guarantee the truth of φ in the resulting state *S'.* As in traditional modal logic, $[\alpha]\varphi$ is equivalent to $\neg\langle\alpha\rangle\neg\varphi$.

The Figure 1 illustrates the framework. We adopt the convention that single arrows represent *atomic actions which* change a state of the world into another one so that the truth values of some sentences change. Different arrows starting from the same state of the world then indicate a nondeterministic choice between atomic actions, that is, one or several of these actions are actually performed but it is left open which. A complex action α, which produces several subsequent state changes, is a recursive construction of sequential compositions, joint actions and (non)deterministic choices of atomic actions and corresponds to a complex directed graph of arrows. [12]

---

[10] A good discussion of action negation is Broersen (2004).

[11] PDL harbors still more operators, but these do not play a role in our analysis.

[12] When Meyer, van de Hoek and van Linder apply PDL to formalize the dynamics of beliefs, desires, intentions, commitments, (Meyer et all 1999) they use a deterministic version of PDL and even in a deterministic context this fine-tuned psychological dynamics turns out to be fairly complicated. To paraphrase engineering knowledge we think as a first approximation we may dispense with these agent specific notions. Engineering means-end knowledge strives to be as least as possible subject-dependent or subjective. We cannot however avoid the nondeterministic version of PDL since engineers often need to choose between different in principle possible actions to achieve their goal.

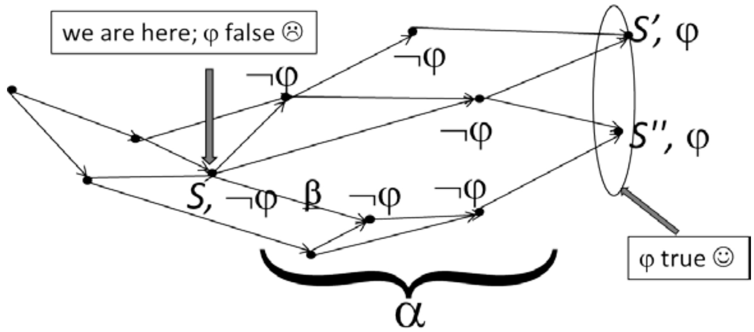

Figure 1: Complex action α in *S* to achieve φ in *S′* or *S″*

PDL takes the weakest form of a normal modal logic. Besides modus ponens and the rule of necessitation ($\vdash \varphi$ implies $\vdash \Box\varphi$, or in other words, logical truths are necessary) it only uses the Kripke (distribution) axiom:

(Kα) $[\alpha](\varphi \rightarrow \psi) \rightarrow ([\alpha]\varphi \rightarrow [\alpha]\psi)$.

(Kα) says that if action α necessary leads to states where φ cannot be true without ψ being true, then if the performance of α will always lead to φ it will also always lead to ψ. So if α makes φ always sufficient for ψ, then we only need perform α to achieve φ in order to achieve ψ, which is not insignificant for engineers. We do not impose any other property on the action operator.

Sufficient means

In order to apply PDL to engineering means-end knowledge, we follow (Hughes, Kroes, & Zwart, 2007) and paraphrase the *means* to achieve something with the *actions* that agents can perform to realize a particular state of the world or, in the language of the adopted framework, make a possible state of the world the actual state. Accordingly, we take *ends* to be the *possible states* in which certain desired states of affairs are instantiated.[13] With these definitions, a *weakly sufficient means* can be defined as follows:

Definition 1: Action α in state *S* is a *weakly sufficient means* for realizing φ iff there is some way by doing action α to achieve φ: $S \vDash \langle\alpha\rangle\varphi$.

[13] We recognize that in natural language means are often identified with objects or instruments and not actions, but do not think this poses an important objection to our approach. We may always translate the object-as-means language into to the actions the means-object performs. If an ignition coil is the means to cause an ignition in the combustion chamber at exactly the right time, we may rephrase this means as the action of igniting.

Weakly sufficient means play an important role in forensic and reversed engineering. It may be helpful for formalizing reasoning about which processes led to an accident, or for functionally decomposing an artifact. In contrast, $[\alpha]\varphi$ cannot be used to define a strongly sufficient means in the same straightforward way. We could have done so if $[\alpha]\varphi$ were false when it is impossible to perform α in state S, which corresponds to $M_{\alpha,S}$ being empty[14]. However, it is a consequence of the adopted definition of the $[\alpha]$ operator that $S \vDash [\alpha]\varphi$ is (trivially) satisfied if performing α is impossible in that state of the world. Consequently, the definition of (strongly) sufficient additionally requires that α can be performed or equivalently, that $M_{\alpha,S}$ is non-empty. A condition that serves this purpose is $S \vDash \langle\alpha\rangle\top$, with $\top$ being the tautology.

Definition 2: Action α in state S is a *(strongly) sufficient means* for realizing φ iff the performance of α in state S always leads to a state in which φ is true and it is possible to perform α: $S \vDash [\alpha]\varphi \wedge \langle\alpha\rangle\top$.[15]

Necessary means

In principle an action β can be defined as *a necessary means* to achieve φ if all weakly sufficient means α for φ involve β. However, this requires a formal definition of the intuitive notion of involvement, which complicates things considerably. The possible complexity of compound actions enables the involved actions to be performed 'in chunks along the line' of other complex actions. If α is an action that is sufficient for achieving one's overall goal, which is being a writer, then writing a book, action β, certainly will be a necessary part of α. But obviously action β is not an atomic action. A complex action α may therefore subsume many other actions in the form of α-parts that together amount to β.

To accommodate the idea of the performance of an action along the way of another, in Hughes et al (2007) the involvement relation has been introduced and axiomatized. The intuitive idea of an action α involving an action β is that one cannot do α without doing β 'along the way'.

---

[14] $M_{\alpha,S}$ denotes the set of states of the world that is the result of performing α in state S.

[15] Although this definition seems intuitively plausible, it lacks important relevance conditions. Typically we do not call an action a sufficient means to some end if this end is inevitable anyway, such as, for example, the setting of the sun. According to definition 2 all actions which start before sunset and end after sunset are sufficient means for the occurrence of the sunset. There are several attempts to repair this, but for the purposes of the present paper we can ignore this problem (cf. Hughes et al. 2007, p. 215/6).

If, for instance, α consists of performing the sequence of sub-actions $\alpha_1 ; \beta_1 ; \alpha_2 ; \beta_2 ; \alpha_3$, (in this order), and $\beta := \beta_1 ; \beta_2$, then α involves β, although α does not involve the action $\beta_2 ; \beta_1$ because the order of the sub-actions counts. We write $\alpha \trianglerighteq_S \beta$ to indicate that α involves β. The subscript index S is required because whether or not an action involves another action may depend on the state of the world. Note that the involvement relation between actions α and β comes close to the idea that doing β is necessary for doing α.

With the involvement relation in place we are able to construct our definition of β being a necessary means for achieving φ in state S. We take this to correspond to the situation where there is a path σ such that (a) the initial state of σ is S; (b) the final state of σ is a φ-state; and (c) σ does β along the way, and where additionally every path σ that satisfies (1a) and (1b) also satisfies (1c).

The first clause guarantees that some part of actions from the current state to a desired state exists and that arriving at the latter requires doing α along the line. The second clause makes α a necessary condition since all paths require that α is being done along the way. The formal definition reads as follows (Hughes et al, 2007, definition 3.2):

Definition 3: An action β is a *necessary means* to φ in state S iff

1. For some action α: $S \vDash \langle\alpha\rangle\varphi$ and $\alpha \trianglerighteq_S \beta$
2. For every $\underline{\cup}$-free action α: if $S \vDash \langle\alpha\rangle\varphi$ then $\alpha \trianglerighteq_S \beta$

To see that the second clause, which its restriction to $\underline{\cup}$-free actions, is strong enough, consider an action α that achieves φ but is not $\underline{\cup}$-free. Then we can rewrite α into $\alpha_1 \underline{\cup} \ldots \alpha_i \ldots \underline{\cup} \alpha_n$ with all α_i being $\underline{\cup}$-free. If all α_i involve β, β is necessary.

Technical Norms

With these three definitions in place, we are now able to sketch the semantics of technical norms of the type

(TN'_N) 'If you want *A* and you are in situation *B*, then you ought to do X.'

As stated above, Niiniluoto's claim is that a norm of this type is true if and only if it is true that 'Doing X is a necessary cause of *A* in situation *B*'.

First, however, we must deal with a problem that this truth condition for technical norms has. The 'and only if' part implies that not wanting *A* in situation *B* is sufficient for X being a necessary cause of *A* in situation *B*. This is a consequence we do not accept. We only subscribe to the right-to-left direction, i.e., *A* being necessary for *B* is sufficient for the truth of (*) but not necessary. This still has the consequence that if (*) is the case, then you do not want *A* or you are not in situation *B* or you ought to do X, which is perhaps somewhat odd. More adequately one could define: if you are in situation *B* and you want *A*, then you ought to do X iff doing X is a necessary cause of *A* in situation *B*.

With Niiniluoto's proposal modified in this way, we come up with the following semantics for it, which basically just amounts to reading 'necessary cause of ...' as 'necessary means to ...':

(*Necessary reading*) 'If you want φ, and you are in situation S, then you should do β' is true in S if

$$\exists\alpha : S \vDash \langle\alpha\rangle\varphi \text{ and for every action } \underline{\cup}\text{-free action } \alpha \text{ in } S\text{: if } S \vDash \langle\alpha\rangle\varphi \text{ then } \alpha \rhd_S \beta$$

Of course in the practice of engineering and design technical norms with this definition of their truth-value are almost always false. Many roads lead to Rome and not all pass the Rubicon. But as stated in the introduction, empirical adequacy is not our primary goal here; we want to show how technical norms, or means-end statements, can get truth-values, by using dynamic logic.

As mentioned above in footnote 3, Niiniluoto also considered weaker variants of (TN'_N). Our semantics generalizes naturally to the variant of the norm that states that 'it is sufficient for you to do α'. This variant we take to be true if doing α is a (strongly) sufficient means to φ in situation S.

(*Sufficient reading*) 'If you want φ, and you are in situation S, then it suffices to do β' is true in S if

$$S \vDash [\beta]\varphi \wedge \langle\beta\rangle\top.$$

This definition achieves our goal of showing how technical norms can get a truth-value, but again one may wonder to what extent this definition of the second reading is empirically adequate. We will briefly discuss the complexities involved in this issue in the Discussion. We complete this section by addressing the remaining issue of the formal validity of practical inference.

3 Deontic PDL and Practical Inferences

Just as the semantics of propositional dynamic logic (PDL) suffices to formulate the truth conditions for technical norms, it suffices also to define the truth condition of (PI)'s second premise, the anankastic statement, which expresses that something is (or is not) a necessary condition for something else. We assume this truth condition to be the same as the sufficient truth condition of (TN'_N), or rather we assume the second premise of a practical inference simply to mean that given the state of the world a particular action is a necessary means to a given end:

(*Anankastic premise)* 'Unless ε is performed, it will not be the case that φ' is true iff

$\exists\alpha : S \vDash \langle\alpha\rangle\varphi$ and for every $\underline{\cup}$-free action α in S: if $S \vDash \langle\alpha\rangle\varphi$ then $\alpha \rhd_S \varepsilon$

Thus, if all paths from S in M that make the hut habitable involve heating the hut, then 'unless the hut is heated, it will not become habitable' is true in S.

This is how far we can proceed using just dynamic logic, which mainly describes how actions may change states of the world. To formulate the semantics of the first premise of (PI), which contains the notion of 'wanting', and its conclusion, which contains the 'must' of practical necessity, we need a language that is minimally equipped with the deontic 'ought to' operator. To that end we turn to Dynamic Deontic Logic as formulated by John Jules Meyer in his seminal (1988). In this paper, Meyer sets out to circumvent the well-known paradoxes of traditional deontic logic by letting the 'Forbidden', 'Ought to' and 'Permissible' operators have actions as their arguments instead of states of the world. To make this work, he introduces a special deontic proposition V, which "is a sentential constant denoting the so-called 'undesirable state-of-affairs', e.g., sanction, (liable to) punishment, trouble (with conscience, for example)". He then defines the notion that some action α is forbidden (F) in a state S by:

Definition F: $S \vDash F\alpha$ iff $S \vDash [\alpha]V$.

This means that α in state S is *forbidden* iff performing α in S always leads to the undesirable state-of-affairs. Likewise an action α is *obligatory* in S iff failing to perform α in S always brings one to the undesirable state:

Definition O: $S \vDash O\alpha$ iff $S \vDash [\sim\alpha]V$

Finally, an action α is permitted in S if it is not forbidden:

Definition P: $S \vDash P\alpha$ iff $S \vDash \neg F\alpha$

These definitions seem to enable us to formalize the conclusion of the Action version of (PI) – 'Therefore one must heat the hut' – in DDeL as $S \vDash O\varepsilon$, where ε is the action of heating the hut. We can define this statement to be true in S iff all performances of ε in S lead to states S' in which V is false, provided ε can be performed in S.[16]

When formalizing the deontic aspect of practical inference in this way, we in a sense ignore the issue of individuality. The conclusion would state that the action of heating the hut must be performed, without specifying by whom it must be performed, although supposedly all actions are performed by someone. We cannot likewise ignore this issue for the first premise, however. Its 'want' must be formalized as the wanting of someone if the formal treatment is to make sense. Our semantics must therefore be able to manage obligations for specific individuals.[17] For this purpose, and thus extending DDeL to give DDeL*, we introduce a personalized deontic operator $O_A(\alpha)$ which expresses that person A has an obligation to perform α. Individual obligations, however, need not be imposed. Agents can also voluntarily take obligations upon themselves. Thus when an agent A has obliged herself to perform action α, we also write this as $O_A(\alpha)$.[18]

With individualized 'oughts' we are now fully equipped not only to formalize (PI)'s conclusion in its 'true' form 'A ought to heat the hut' as $S \vDash O_A(\varepsilon)$, where ε is the action 'heating the hut'.[19] For the first premise, we take seriously Von Wright's (1972: p. 55) remark that "action cannot possibly follow logically from premises about intentions and epistemic attitudes" and

---

[16] Note that ε need not be atomic but may consist for instance of two subsequent subactions: $\varepsilon := \varepsilon_1; \varepsilon_2$. In such a case $O[\varepsilon_1; \varepsilon_2] \equiv O\varepsilon_1 \wedge [\varepsilon_1]O\varepsilon_2$. (Note that in contrast $F[\varepsilon_1; \varepsilon_2] \equiv [\varepsilon_1]F\varepsilon_2$.) This means that V is true in all states along the paths towards the state where the hut is finally heated; only in that final state V is false.

[17] Individual deontic operators may also be found in for instance work of Wieringa and Meyer, (1993), Herzig and Lorini, (2010), and other publications in STIT logic.

[18] There are ways for DDeL* to bring out the distinction between voluntary obligations and imposed obligations. Voluntary obligations result from a previous action of the person who has the obligation, namely the action of taking an obligation upon oneself, whereas for imposed obligations the action of imposing the obligation is performed by another person than the person who has the obligation.

[19] It is here that the step from the *Dynamic-States version* to the *Action version* of PI is actually made.

interpret the statement '*A* wants α' to mean or stand for '*A* has obliged herself to perform α', or formally $S \vDash O_A(\alpha)$. Of course one may object that wanting is not always a form of obliging oneself. However, we take the wanting in the first premise not to be some casual wanting like in 'I want it to be nice weather'. From such wanting indeed little could follow in the strict deductive sense. Instead we look upon the wanting as signifying a top priority in agreement with Von Wright who explains the first premise as follows: "we want to attain the end as a result or consequence of something which we do" (PI, p.160). Such wanting involves a form of commitment that we see as at least close to obliging oneself to bring about the stated end.[20] We return to this issue in the Discussion section.

Finally we are ready to investigate whether practical inference, as introduced by Von Wright's and made more explicate by us, is logically valid in DDeL*. In our adapted, individualized Action version the inference is:

(PI´) A has obliged herself to make the hut habitable.
Unless the hut is heated, it will not become habitable.

A is under the obligation to heat the hut.

With φ being 'the hut is habitable', α being 'making the hut habitable' and ε being 'heating the hut' the semantical counterpart of this inference reads:

(PI´) $S \vDash O_A(\alpha) \wedge [\alpha]\varphi \wedge \neg\varphi$

$S \vDash \langle\alpha\rangle\top \wedge \forall\alpha\, (\langle\alpha\rangle\varphi \rightarrow \alpha \trianglerighteq_S \varepsilon\,)$

$S \vDash O_A\,(\varepsilon)$

To show that the truth of the premises implies the truth of the conclusion, let us first assume that 'making the hut habitable' is $\underline{\cup}$-free, which means that it is a linear order of actions; that it is possible for *A* to perform it, and that it involves 'heating the hut'. As mentioned before, ε,

[20] Meyer and others similarly claim that 'being committed to α' intuitively corresponds to 'having promised (to oneself) to perform α next (or at least a.s.a.p.)' (Meyer et al. 1999, p16). Note, further, that although one may want unattainable things, here we do assume that the things wanted are at least empirically attainable. This is relevant for engineering and will be taken up again in the final Discussion section.

heating the hut, need not be an atomic action. We let ε refer to that string of actions starting from *S* to *S´* in which the hut is finally heated. We let β refer to the remainder of the string of actions, between *S´* and *S˝* in which α is completed and the hut is habitable. So we have $\alpha = \varepsilon; \beta$ (and of course β may be non-existent). Consequently, more syntactically, the DDeL* representation of (PI) is:

(PI´) $O_A(\varepsilon; \beta) \wedge [\varepsilon;\beta]\varphi \wedge \neg\varphi$

$\langle\varepsilon; \beta\rangle\top$

$O_A(\varepsilon)$

That (PI´) is valid follows very easily from one of the theorems of DDeL which states that $O(\varepsilon; \beta) \equiv O\varepsilon \wedge [\varepsilon]O\beta$. In case α is not $\underline{\cup}$-free, it may be rewritten in a disjunctive normal form $\alpha := \alpha_1 \underline{\cup} \ldots \alpha_i \ldots \underline{\cup}\, \alpha_n$ in which all the α_i are $\underline{\cup}$-free. Because α is finite, all its branches end up in states where φ is true (since $[\varepsilon;\beta]\varphi$), and all these branches involve the activity of heating the hut. Thus we can identify ε with a structure of branches where each branch starts in state *S* and ends in a state *S´* in which the heating has finished, and take β to be the rest of α, and then we can again apply $O(\varepsilon; \beta) \equiv O\varepsilon \wedge [\varepsilon]O\beta$ to show that that (PI´) is valid.

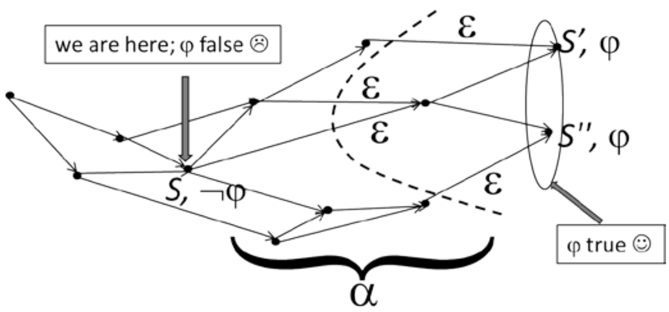

Figure 2: To achieve φ in *S′* or *S″* one cannot escape to perform ε somewhere along the line.

Let us summarize, then, why (PI´) is valid. An agent *A* in some state of the world ‘really’ wants something, which is not yet the case in that state and which is her goal φ, and therefore obliges herself to achieve φ by performing an action α (the first premise). By empirical necessity, by the laws of nature jointly, so to speak, all action paths that bring her from the initial state of the world to a state where her goal is realized involve the performance of another action β (second premise). Then we may validly conclude that *A* is obliged to perform β.

4 Discussion

To recap the starting point of our paper, in a practical inference (PI) as defined by Von Wright, an obligation to act is deduced from a want or an intention to realize an end and from an empirical constraint on how this end can be achieved. Technical norms (TN), considered as 'contracted' PIs, are conditional statements with the setting of an end as the antecedent and an obligation to act as the consequent. In this paper we used dynamic logic to determine formal truth conditions for technical norms and to define validity for practical inferences. But by doing so, we adjusted them considerably to make them fit into our DDeL* framework. The main question that remains to be discussed here is to what extent the results regarding our adapted PI and TN versions still say something substantial about the original ones. To what extent do our paraphrased versions of TN and PI do justice to the intuitions that underlie our uses of these statements and forms of reasoning?

A reason to doubt that they do them justice is the observation that, in practice, technical norms and practical inferences are always defeasible or non-monotonic. Every engineer would agree that for hardly any design problem, a necessary condition exists that is part of all solutions. No engineering task seems to be absolutely determined in this way. Thus in real-life situations a technical norm will have the form: 'If you want φ, and you are in situation *S*, then *normally* you should (or: it suffices to) do α'. This is so because in real-life situations the underlying anankastic statement is seldom deterministic and almost always has the form: 'unless α is performed, *normally* it will not be the case that φ'. Accordingly, the pertaining conclusion is that someone who wants φ and is in situation *S* *normally* should perform α.

Formally, our semantics can cope with non-monotonicity by introducing a relation *r* that for every state determines the set of normal states S_n that are, up to a certain extent, similar to *S*. This *r* enables us to define a defeasible conditional $\psi \Rightarrow \varphi$, which is already true in *S* if only the *S-normal* ψ-states satisfy φ (but not necessarily all ψ-states). Then, we can use ψ as a sufficient precondition for the necessary means-end relation involving α and φ by letting the conditions 1 and 2 of Definition 3 not apply to the unique state *S* but to all normal states S_n and in this way arrive at a defeasible TN version. With the same mechanism we may render the formal version of PI non-monotonic.[21]

Where would this leave us with regard to the adequacy of our analysis? On the one hand we have engineers and other professional practitioners who, on a daily basis, draw defeasible conclusions about what should be done, and on the other we have constructed a non-monotonic framework for drawing indefeasible conclusions from self-obligations and anankastic statements. The proposed introduction of defeasibility into the framework, although opening a way to bring our approach more in line with reasoning as it actually occurs, requires a similarity- between-states-of-the-world notion, which is barely recognizable in real-life means-end reasoning. More in-depth research should find out how in practice professionals cope with the defeasibility of their conclusions and whether our similarity-based mechanism can accommodate this reasoning in practice.

A second, seemingly even more profound departure from real-life practical reasoning of our formal semantics is the assumption that an agent can only oblige herself to perform *successful* actions. In every state S and for every action α, the semantics determines whether $S \vDash \langle\alpha\rangle\top$ holds, that is, whether α can be performed in S. Moreover, for all S and performable α the semantics fixes the resulting world states. The framework therefore does not allow for actions with undetermined end results.[22] In real life, agents often initiate actions without being sure whether they will result in the intended end. In the course of their work, engineers regularly change their course of action to adapt to changing circumstances or changing requirements. In addition, research engineers, who are engaged in extending the limits of what is technically possible, often start on projects of which they are not even sure that they are realizable. Take for instance, the first transatlantic communications cable or keeping heavier-than-air contraptions airborne.

Whether the success assumption of α poses a profound or only a superficial obstacle to the adequacy of our model, and how much it takes to repair this, depends heavily on fundamental philosophical questions. Our model comes close to a form of technical realism according to

---

[21] See for more technical details Hughes et al. (2007), pp. 225-227.

[22] This differs from the claim that for $\langle\alpha\rangle$ and $[\beta]$, α can and β possibly cannot be performed. The reader interested in the logic of effectiveness may consult Harz (2007) (in German).

which all technical possibilities are fixed once and for all. From this point of view one may argue that taking up an unattainable engineering task is not performing an action at all and therefore does not need to be represented in the framework. Technical constructivism, according to which the technically possible is a realm of human construction, is much more difficult to conciliate with our framework, especially for engineering tasks that have never been performed before. The quest for pushing technological boundaries will often involve the performance of actions that fail to achieve the projected ends. Technical realism is closely related to scientific realism according to which the laws and theories of physics lie hidden 'out there', to be discovered by science. Constructivism, in contrast, is related to scientific empiricism or instrumentalism, which acknowledges a good deal of human construction in the development of scientific theories. Furthermore, the two points of views differ markedly in that technical realism is primarily backward-looking and technical constructivism primarily forward-looking. It is not our task here to decide between the two. We only note that our framework fits the former perspective better than it fits the latter.

The last and perhaps most salient difference between our formal account and the intuitive notions of practical inference and technical norm is the interpretation of the 'want' or intention-statement. In our formal version of PI we paraphrase '*A* wants φ' as '*A* obliges herself to (bring about) φ'. Even if we restrict 'wants' to a person's top priorities, for which it is plausible that they involve some form of self-obligation, we have to acknowledge important differences between the two. Wants are akin to desires (Irvine, 2006). The wants in practical inference appear to be much more similar to desires than to self-obligations, and as the distinction between desires and self-obligations is obvious, so must the wants be different from self-obligations. To oblige oneself, is the performance of an action, whereas ending up wanting something does not appear to be related to the performance of a prior explicit action. Wants are typically generated beyond the grasp of one's own will; they happen to you and may fluctuate, whereas self-obligations seem rather to be decisions taken because one wills them, at a precise moment, and do not fluctuate.

Whereas we acknowledge that in general wants do not imply self-obligations, in our formal analysis we have interpreted wanting in its strongest form, which we take to be close enough to self-obligation to be allowed to treat them as equivalent. In doing so, we are convinced that we have excavated and modeled all the logical structure and the validity there is to practical inference. This is a strong claim, which may be underpinned by the following robustness argument. None of the authors who so far have seriously studied the validity of practical inference have found a valid way to deduce an (obligation to) an action from a want or intention and an empirical necessity. All of them, when establishing a logically valid connection between the premises and the conclusion of a practical inference, have, more or less independently, used forms of the same strategy, which we have applied, and which can be likened to a reduction-to-the-same-denominator. Similar to the process of reducing fractions to the same denominator, this strategy puts the 'want'-premise and the 'must'-conclusion on the same footing. For instance, the form in which Von Wright ultimately considered practical inference to be valid is (1972, p.47):

(PI[I]) *X* intends to make it true that E

He thinks that, unless he does *A*, he will not achieve this

X intends to do *A*.

In Von Wright's own words this form of practical inference exhibits a *transmission of intentions* from premises to the conclusion. It is exactly this form that Meyer et al (1999, p. 15) paraphrase in their framework. In the same sense we may call our valid version of the practical inference *transmission of obligations*. Also Broome, in his analysis of practical reasoning (2003), applies the same reduction-to-the-same-denominator strategy. He paraphrases the 'want'-statement and the 'must'-conclusion in the future tense and infers 'Chris will borrow money' from 'Chris will buy a boat' and 'for Chris to buy a boat a necessary means is for Chris to borrow money'. Like Von Wright before him, however, Broome (ibid, section 2) claims the inference to be valid without further ado. Finally we may even turn to Immanuel Kant, who claims in his *Groundwork* that 'whoever wills the end, wills also (according to the dictate of reason necessarily) the indispensable means thereto which are in his power'. Again, this, according to Kant, analytically true statement may be seen as an instance of the same strategy. Answering Segerberg's (1980) call to apply dynamic logic to philosophical problems, we also came to the same result, and which we now in retrospective consider to be a sound philosophical (and not technical) norm: If you want a practical inference to be logically valid, you should apply something similar to the reduction-to-the-same-denominator strategy.

5 References

Bratman, M. E., Israel, D. J., & Pollack, M. E. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, *4*(3), 349–355.

Broersen, J. (2004). Action negation and alternative reductions for dynamic deontic logics. *Journal of Applied Logic*, *2*(1), 153–168.

Broome, J. (2002). Practical reasoning. In J. Bermúdez and A. Millar, eds., *Reason and nature: Essays in the theory of rationality*. Oxford University Press, 85–111.

Harz, M. (2007). *Zur Logik der technologischen Effektivität*. Brandenburgischen Technischen Universität, Cottbus. Retrieved from https://opus4.kobv.de/opus4-UBICO/frontdoor/index/index/docId/3738

Herzig, A., & Lorini, E. (2010). A dynamic logic of agency I: STIT, capabilities and powers. *Journal of Logic, Language and Information*, *19*(1), 89-121. https://doi.org/10.1007/s10849-009-9105-x

Hughes, J., Kroes, P., & Zwart, S. (2007). A semantics for means-end relations. *Synthese*, *158*(2), 207–231.

Irvine, W. B. (2006). *On desire: Why we want what we want*. Oxford University Press.

Meyer, J.-J. (1988). A different approach to deontic logic: Deontic logic viewed as a variant of dynamic logic. *Notre Dame Journal of Formal Logic*, *29*(1), 109–136.

Meyer, J.-J. C., van der Hoek, W., & van Linder, B. (1999). A logical approach to the dynamics of commitments. *Artificial Intelligence*, *113*(1), 1–40.

Niiniluoto, I. (1993). The aim and structure of applied research. *Erkenntnis*, *38*(1), 1–21. https://doi.org/10.1007/BF01129020

Segerberg, K. (1980). Applying modal logic. *Studia Logica*, *39*(2–3), 275–295.

Wieringa, R. J., & Meyer, J.-J. C. (1993). Actors, actions, and initiative in normative system specification. *Annals of Mathematics and Artificial Intelligence*, *7*(1–4), 289–346.

Von Wright, G. H. (1963PI). Practical inference. *The Philosophical Review*, *72*(2), 159–179.

Von Wright, G. H. (1963NA). *Norm and Action: A logical enquiry*. London: Routledge & Kegan Paul.

Von Wright, G. H. (1963VG). *The varieties of goodness*. London: Routledge & Kegan Paul.

Von Wright, G. H. (1972). On so-called practical inference. *Acta Sociologica*, *15*(1), 39–53.

Von Wright, G. H. (1999). Deontic logic-as I see it. In P. McNamara & H. Prakken, *Norms, logics and information systems: New studies in deontic logic and computer science* (Artificial Intelligence and Applications, vol. 49). Amsterdam: IOS Press, 15–25.

Figure 1

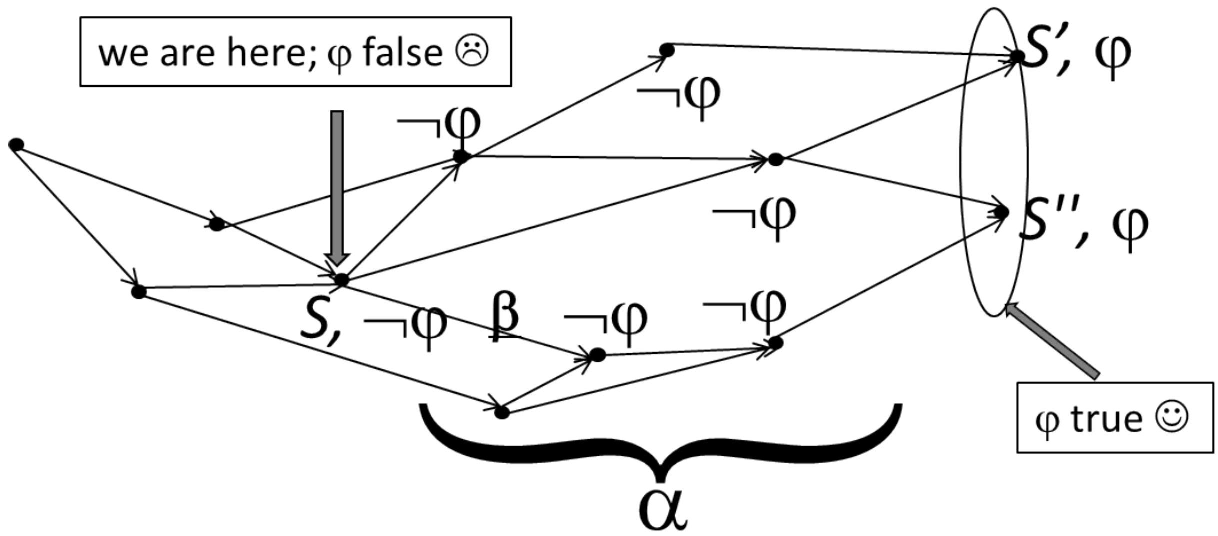

Figure 2