Contrary to Duty Obligations
A Study in Legal Ontology

Jaap Hage
Department of Metajuridica
Faculty of Law
Universiteit Maastricht
The Netherlands
jaap.hage@metajur.unimaas.nl

Abstract. In this paper the problems in deontic logic around contrary to duty obligations are used to conduct a study in basic normative ontology. Three causes of the problems around contrary to duty obligations are identified, that is 1) the attempt to analyze obligations in terms of what is ideally the case, 2) the application of deontic inheritance to the presuppositions of obligations, and 3) the failure to distinguish between what will be called 'inclusive' and 'exclusive ought-to-do'. These three causes are all attributed to insufficient distinctions on the ontological level.

1 The Relevance of Deontic Logic for Legal Ontology

A central issue in legal ontology is the study of normative positions. Such normative positions include being under an obligation, having a right, being entitled, etc.[1] The study of these positions and their mutual relations is a study of basic (legal) ontology, and includes the study of the relations between what ought to be done, what is permitted, and what is obligatory. In short, the central topics of deontic logic are part of the study of normative positions, which is in turn part of legal ontology.

In this paper I hope to make a contribution to the development of legal ontology by focusing on an issue that has plagued the development of deontic logic for quite a while. I refer to the issue of contrary-to-duty obligations (CTDO's). Since Chisholm [4] criticized the Standard System of Deontic Logic (SDL) for not being able to deal with duties which arise as a consequence of a norm violation, the issue of CTDO's has not disappeared from the stage of deontic logic.

My main thesis in this paper is that the complications that arise in connection with CTDO's have little to do with CTDO's themselves, but are consequences of deficient theories of the nature of obligations. Before starting my argument, I want to warn the reader that she should not expect a proper paper on deontic logic, let alone the umpteenth attempt to 'solve' the CTDO-paradoxes. My focus will be on analysis, not on the development of a logical system.

[1] See e.g. [1, 2, 3].

Jaap Hage, 'Contrary to Duty Obligations – A Study in Legal Ontology' in Bart Verheij, Arno R. Lodder, Ronald P. Loui and Antoinette J. Muntjewerff (eds.), *Legal Knowledge and Information Systems. Jurix 2001: The Fourteenth Annual Conference*. Amsterdam: IOS Press, 2001, pp. 89-102.

2 Ought-to-do and Ought-to-be

Prescriptive and prohibitory norms have as one of their main functions to guide human behavior. They indicate what people or organizations should or should not do. There seem to be two different ways in which this purpose can be achieved. A norm can indicate what kind of action is to be performed or refrained from, and for which actors this holds. Such norms are called ought-to-do (or Tun-sollen) norms.

A norm can also indicate how the world should look like, which states of affairs ought (not) to be the case. These norms are called ought-to-be (Sein-sollen) norms. Ought-to-be norms sometimes seem to occur in legal practice. In a sense they are problematic, however, because they do not specify what ought to be done, and which actors are to do what ought to be done.

It seems to me that ought-to-be norms can be given two different interpretations, which should be distinguished very carefully. On the one interpretation, ought-to-be norms specify how the world would ideally be. On this interpretation, ought-to-be norms themselves do not prescribe or prohibit behavior, because there is no logical relation between what is ideally the case and what ought to be done. Only if ought-to-be norms are complemented with the prescription to strive for the ideal, they guide behavior. But then the 'real' norm is the ought-to-do norm that prescribes to strive for the ideal. The ought-to-be norm merely provides content to this prescription.

On the other interpretation, ought-to-be norms prescribe to see to it that the obligatory state of affairs is achieved or maintained, depending on whether the ideal state in question already obtains. Briefly stated, the obligation is to see to it that the obligatory state of affairs obtains. On this interpretation, the ought-to-be norm is 'really' an ought-to-do norm in disguise. It is an incomplete ought-to-do norm, because it leaves the actor unspecified. This deficiency can be remedied, however, by saying that the actors are those who are responsible for seeing to it that the obligatory state of affairs is achieved or maintained.

If this analysis is correct, ought-to-be norms as such only seem to exist. In reality they are either no norms (in the sense of behavior guiding entities) at all, but the formulations of ideals, or they are ought-to-do norms with unspecified actors. On this analysis there would only be ought-to-do norms, and deontic logic should therefore be a logic of ought-to-do norms.

2.1 Attempts to Reduce Ought-to-do to Ought-to-be

It may seem, however, that there are real ought-to-be norms nevertheless. In this connection one might think of norms of the form 'It ought to be that X does (or: refrains from doing) A'. Moreover, one might even argue that ought-to-do norms are nothing else than such ought-to-be norms. 'X ought to do A' would on this view be analyzed as 'It ought to be the case that X does A'.

Such an analysis would in my opinion be defective. The sentence 'X ought to do A' means amongst others that X is responsible for doing A. This responsibility is lost on the analysis that it ought merely be the case that X does A. If it ought to be the case that X does A, it may very well also be the case that not X, but Y is responsible for X's doing A. In other words, if 'X ought to do A' is analyzed as 'It ought to be the case that X does A', something of the meaning of the original sentence is lost.

One might attempt to rescue the ought-to-be analysis by distinguishing between:

a It ought to be the case that X does A; and

b It ought to be the case that Y sees to it that X does A.

The former sentence might then be read as expressing that X is the responsible person for his own doing A, while the second sentence would be read as expressing that Y is responsible for X's doing A. This 'rescue' does not work, because either the second sentence leaves unspecified whether Y is responsible for his seeing to it that X does A, or it is not possible anymore to express that ideally Y sees to it that x does A, but that Y is not responsible for doing so.

Moreover, there is another objection against the ought-to-be analysis. Let us assume that sentences of the form 'It ought to be the case that ….' specify an ideal situation. On this assumption, the sentence 'It ought to be the case that X does A' means that in the ideal situation X does A. It may, however, be the case that X has promised to do A, while the world would be very slightly better if X would not do A. In this case, I feel that we can say that both X ought to do A, and that ideally X would not do A. Otherwise the obligation to keep one's promises would not exist if somebody promises to do what will lead to a sub-ideal outcome (the value of promise-keeping included). This is an understandable moral position (a kind of utilitarianism), but it should not be made the correct position on purely logical grounds.

This counter-intuitive result can be avoided by not assuming that 'It ought to be the case that …' indicates an ideal situation. However, if this assumption is not made, it is unclear how the phrase in question should be interpreted. The temptation is big to interpret 'It ought to be the case that X does A' as meaning the same as 'X ought to A', but then an ought-to-be is reduced to an ought-to-do, rather than the other way round.

In my opinion it is best to interpret ought-to-be sentences as specifying ideal situations, and on this interpretation ought-to-be and ought-to-do cannot be reduced to each other.

2.2 Standard Deontic Logic as a Logic of Ideals

SDL is an extension of propositional logic. Its language allows the formation of deontic sentences by pre-fixing a deontic operator, typically O, F, or P, to a sentence. Since the deontic operators operate on sentences, the resulting logic is a logic of the ought-to-be type. A sentence of the form O(P) would typically be read as 'It is obligatory that P', or 'It ought to be the case that P'.

SDL is best characterized by its semantics. A sentence of the form O(P) is interpreted as true iff the sentence P is true in all 'ideal' worlds. F(P) is true iff P is false in all ideal worlds, while P(p) is true iff P is true in at least one ideal world.

The 'official semantics' for SDL is formulated more precisely, by defining an accessibility relation R over the set of possible worlds. For each possible world w there is a non-empty set of possible worlds that stand in this relation to (are accessible from) w. The sentence is true in w iff the sentence is true in all the worlds that stand in the relation R to w. Although this official semantics is more precise, I also gave the intuitive reading, because it is, as we will see, more revealing about the nature of SDL.

The semantics of SDL validates a rule of inference to the effect that the logical consequences of what is obligatory are obligatory too. That this rule is validated is immediately clear if one realizes what is obligatory in a world is the case in its ideal alternatives. In these alternatives the facts that are logically implied by the facts that obtain must also obtain. Therefore the logical consequences of what is obligatory in a world w is the case in all ideal alternatives of w, which means that they are obligatory in w too. For instance, if O(P&Q) is true in w, P&Q is true in all its ideal alternatives. Therefore P is true in all its alternatives, which means that O(P) is true in w.

A special case of this general rule is the rule that allows arguments of the form *deontic detachment*:

> Obviously, if both p→q and p are true in all ideal worlds, q must be true in all ideal worlds too.

Given its semantics, SDL is a suitable logic for ideals. It specifies what logically follows from facts that ideally obtain. The facts that follow from what is ideal are ideal themselves. Apparently the characterization of the worlds that are accessible from a world w as ideal worlds relative to w is aptly chosen.

2.3 The Chisholm Paradox

We have seen that ought-to-be sentences can both be interpreted as sentences about what is ideally the case, and as ought-to-do sentences in disguise. An interesting question is whether SDL is also a suitable logic for this latter interpretation of ought-to-be sentences. Is SDL a suitable logic if O(P) is interpreted as 'The person who is responsible for situation P ought to see to it that P is the case'?

Briefly stated, the answer to this question is an unambiguous 'No'. SDL is not suitable for the ought-to-do interpretation of ought-to-be sentences. The first indication that this is so can be found if the rule for deontic detachment is inspected. If somebody ought to see to it that Q if P, and if he also ought to see to it that P, this does not logically imply that he ought to see to it that Q.

An example might make this clearer. Suppose that John ought to see to it that he helps his neighbors and also that if he helps them, he tells them so. Do these obligations together imply that John ought to see to it that he tells his neighbors that he will help them? Obviously not. John ought only tell his neighbors that he will help them if he will actually help them. The mere obligation to help them is not sufficient.[2] And yet this inference, a special case of deontic detachment, is validated by SDL. SDL leads to counterintuitive results if interpreted as an ought-to-do logic in disguise, because deontic detachment is not suitable for reasoning with ought-to-do obligations.

Chisholm [4] has employed this insight by showing how SDL leads to paradoxical results if interpreted as an ought-to-do logic in disguise. Chisholm's case ran as follows:

[2]The attentive reader will have noticed that I moved from ought-to-see-to-it terminology to ought-to-do terminology. This is on purpose. I think that in the present context the ought-to-see-to-it terminology is forced and only motivated by a misguided desire to stick as much as possible to ought-to-be like terminology.

- A certain man ought to go to assist his neighbors.
- He ought to tell them he is coming if he goes.
- If he does not go then he ought not to tell them he is coming.
- He does not go.

According to the usual model theoretic semantics of SDL this can be translated in terms of ideals as follows[3]:

- In the ideal situation a certain man goes to the assistance of his neighbors.
- In the ideal situation, if he does go he tells them he is coming.
- If he does not go then in the ideal situation he does not tell them he is coming.
- He does not go.

From the first two premises follows (by deontic detachment) that in the ideal situation, our man tells his neighbors he is coming. From the last two premises follows (by Modus Ponens) that in the ideal situation he does not tell that he is coming. In other words, framed as a theory about what is ideally the case, this set of premises is inconsistent. And yet, interpreted as a theory of what ought to be done, the premises are consistent.

The point of this paradox is that if ought-sentences are interpreted as sentences about what is the case in an ideal situation, they cannot be used to express what ought to be done in sub-ideal situations. Since there can obviously be obligations about what to do in sub-ideal situations (in particular CTDO's) SDL is unsuitable to deal with an important category of obligations. This point can even be generalized. If ought-sentences are interpreted as saying what is the case in all worlds that belong to a particular subset of possible worlds, whether these worlds are ideal or not, the resulting logic cannot be used to express CTDO's.

This is far from surprising, because SDL is a logic for ought-to-be interpreted as ideals, and not for ought-to-be interpreted as ought-to-do in disguise. The Chisholm paradox is a good way to point this out, but the validity of deontic detachment in SDL also sufficiently demonstrates that SDL should not be used as a logic for what ought to be done. If deontic logic is taken to be a logic for ought-to-do (and not for what is ideally the case), SDL or variants of it that allow deontic detachment should not be used for deontic logic. This conclusion is based on general considerations concerning ought-to-be 'obligations' and ought-to-do obligations, and not merely on difficulties in SDL to deal with CTDO's. Attempts to contest the conclusion by pointing out sophisticated logical theories that can 'deal with' CTDO's in a SDL-like logic are not sufficient, unless they are accompanied by a philosophical theory that explains how ought-to-do obligations can be reduced to ought-to-be's. As yet I have never seen such a philosophical theory.[4]

3 Deontic Inheritance

Some may think that I am too rash in limiting the validity of deontic detachment to a logic for ideals. Deontic detachment, it may be argued, is a special case of deontic inheritance, the

[3]There are all kinds of intricacies involved in the formalisation of this example, which are beyond the scope of this paper..

[4]The absence of satisfactory philosophical theories about the (non-)relation between ought-to-do and ought-to-be explains why this paper hardly pays any attention to the numerous attempts to deal with CTDO's within the boundaries of SDL.

validity of which is a desirable characteristic in deontic logic.

Informally stated, deontic inheritance is the phenomenon that if a state of affairs or an action is obligatory, its implications are also obligatory. For instance, if cars ought to be both locked and well-parked, it follows that cars ought to be locked. Or, if John ought to repay all his debts, it follows that John ought to repay his debt to Alice.

Both examples provide in my opinion illustrations of well-used deontic inheritance. But precisely what do they illustrate? The first example is probably best interpreted as dealing with deontic inheritance in a logic of ideals. If it ought to be the case that cars are both locked and well-parked, this means that in the ideal case, cars are both locked and well-parked. From this it follows, unproblematically as far as I can see, that in the ideal case cars are locked. This is precisely the kind of deontic inheritance that SDL allows, and in connection with what is ideally the case, this is also as it should be.

The second example leads to some complications. It may well be that if some type of action 'implies' some other type of action, the obligation to perform the first kind of action 'implies' the obligation to perform the second kind of action. The question, however, is what it means that an action implies another action. Actions are not propositions or sentences, and the implication can therefore not be logical entailment. The implication must here be some other kind of relation. Let us call this relation 'involvement', and let us hypothesize that if an action involves some other action, this latter action inherits the obligatory nature of the former action.

The question that still needs to be answered is when one action involves another action. A thorough investigation of this issue lies outside the scope of this paper, but in connection with CTDO's, one particular kind of involvement seems especially interesting. There is a kind of action that is closely related to propositions, namely seeing to it that some state of affairs obtains.[5] This close relation can be exploited by characterizing the involvement between two actions of the type 'seeing to it that ...' in terms of the entailment relation between the states of affairs that ought to be seen to. Some problems with CTDO's which were raised in connection with SDL then turn up again in connection with the logic for ought-to-do.

3.1 Some Linguistic Conventions

To avoid cumbersome sentences in discussing actions of the see-to-it type, I introduce some formalism. The two-place predicate`Od` is used to express ought-to-do obligations. The first parameter denotes the actor, the second the action type that is to be performed. For instance, `Od(john, pay_debts)` would mean that John ought to pay his debts.

A second convention is that if `P` is a sentence, `*p` denotes the state of affairs expressed by this sentence. If, for instance, the sentence `Closed(window)` expresses that the window is closed, `*closed(window)` denotes the state of affairs that the window is closed.

The function-expression `stit/1`is used to transform states of affairs into action types. `Stit` stands for 'see to it that'. For instance, `Od(john, stit(*closed_window))` would mean that John ought to see to it that the window is closed. This is an ought-to-do sentence, which should be distinguished from the sentence that it ought to be the case that John sees to it that the

[5]The notion of 'seeing to it that ...' has been developed in agency theory. A brief overview of this literature can be found in [5].

window is closed. In contrast to the latter, the former sentence expresses John's responsibility for seeing to it that the window is closed, and does not express that in the ideal case, John sees to it that the window is closed.

The one-place function expression ¬ is used to transform an action type into the action type which is refraining from the first action type. For instance, ¬cycle denotes the action type refrain from cycling.[6] In this connection I will assume that Od(actor, ¬action) entails ∼Od(actor, action).

Finally the two-place predicate Involves is used to express that performing the action type denoted by its first parameter involves performing the action type denoted by its second parameter. For instance, the sentence

Involves(fulfill_duties, pay_debts)

expresses that fulfilling one's duties involves paying one's debts.

3.2 Deontic Inheritance for Ought-to-do

Having the linguistic tools in place, we can look at the logical characterization of action types that involve each other. An attractive looking hypothesis is that if some sentence entails another sentence, seeing to it that the state of affairs expressed by the first sentence obtains involves seeing to it that the state of affairs expressed by the second sentence obtains. In our primitive formalism: if P entails Q, stit(*p) involves stit(*q). This hypothesis would have the desired consequence that seeing to it that the cars are locked and well-parked, involves seeing to it that the cars are locked:

Involves(stit(*locked(cars) & well_parked(cars)), stit(*locked(cars)))

Or, in general[7]:

$$N(P \rightarrow Q) \rightarrow \quad \forall x(\text{Involves}(\text{stit}(*p), \text{stit}(*q)))$$

This can be combined with the theory that if some action type involves another action type, the obligation to perform the first type of action 'entails' the obligation to perform the second type of action[8]:

$$\text{Involves}(a, b) \rightarrow (\text{Od}(x, a) \rightarrow \text{Od}(x, b))$$

[6] The operator '∼' is used for negation.

[7] The combination of the necessity-operator N in combination with the material conditional is used to express entailment.

[8] In my opinion this entailment is better represented as a rule, but to avoid additional complications I ignore that here.

3.3 The Paradox of the Good Samaritan

The construction proposed above to translate logical relations between sentences to involvement-relations between action types has the consequence that some paradoxes related to CTDO's re-appear. I will illustrate this by discussing the paradox of the Good Samaritan [6]. The biblical case of the Good Samaritan runs as follows: A Jew is robbed and the Samaritan ought to help the Jew who was robbed. This case might be formalized as:

$$O(\text{Helps}(\text{samaritan}, \text{robbed\_jew}))$$
$$N(\text{Helps}(\text{samaritan}, \text{robbed\_jew}) \rightarrow \text{Robbed}(\text{jew}))$$

These two sentences allow under (a somewhat extended version of) SDL the derivation of

$$O(\text{Robbed}(\text{jew}))$$

In other words, it is possible to derive from the sentence that it is obligatory that the Samaritan helps the robbed Jew, that the Jew ought to be robbed. A crucial assumption in this connection is that the sentence that the Samaritan helps the robbed Jew entails[9] the sentence that the Jew was robbed and that this entailment makes that the obligatoriness of the helping is inherited by the being robbed.

One way to solve this paradox is to make a scope distinction and leave the proposition that somebody was robbed outside the scope of the deontic operator[10]:

$$\exists x(\text{Robbed}(x)) \,\&\, \text{Od}(\text{samaritan}, \text{help}(x))$$

Notice, however, that in this formalization the deontic predicate occurs within the scope of the existential quantifier and that the resulting sentence is not of the ought-to-be type anymore (even disregarding the replacement of the O-operator by the Od-predicate). Apparently an adequate analysis of this type of case requires a deontic logic that is not of the ought-to-be type.

The question that is now to be addressed is whether the problem that occurs in an ought-to-be setting re-appears in an ought-to-do setting with stit action types. The following formalization nicely illustrates the point that I want to make:

$$\text{Od}(\text{samaritan}, \text{stit}(*\text{helps}(\text{samaritan}, \text{robbed\_jew})))$$
$$N(\text{Helps}(\text{samaritan}, \text{robbed\_jew}) \rightarrow \text{Robbed}(\text{jew}))$$

Under the hypothesis formulated in the previous section, the entailment of $\text{Robbed}(\text{jew}))$ by $\text{Helps}(\text{samaritan}, \text{robbed_jew})$would, through deontic inheritance, allow the derivation of $\text{Od}(\text{stit}(\text{samaritan}, \text{robbed\_jew}))$. The problem that occurred in connection with SDL seems to

[9] Later in this section I will raise doubts concerning this 'entailment'.

[10] The scope distinction proposed here was inspired by a similar distinction proposed by Smullyan [7] in connection with alethic modal logic. A strongly related approach was proposed by Castañeda [8]. He distinguished between actions (practitions) and factual circumstances implied by those actions. For instance, the action of helping the robbed Jew implies the factual circumstance that the Jew was robbed. According to Castañeda, the obligation has a deontic focus which is only directed on the helping, and not on the fact that the Jew was robbed.

re-appear under the present version of ought-to-do- logic. Apparently the transition from ought-to-be to ought-to-do logic is not sufficient, and a deeper solution is required.

This deeper solution is in my opinion that the relation between the state of affairs that the Samaritan helps the robbed Jew and the state of affairs that the Jew was robbed must be analyzed more thoroughly. The first state of affairs does not entail the second state of affairs, but *presupposes* it. That is a different relation than that between the state of affairs that the cars are both locked and well-parked on the one hand and that the cars are locked on the other hand. This latter relation would in my opinion be one of *entailment* proper.[11] The first state of affairs so to speak 'comprises' the second one. In such a case, deontic inheritance seems unobjectionable. An obligation to bring about a particular state of affairs extends to what this state of affairs comprises. If one state of affairs presupposes another one, however, deontic inheritance seems in general not to be acceptable. The obligation to bring about a particular state of affairs does not extend to this state of affairs' presuppositions.

Therefore I propose to limit the notion of entailment that allows deontic inheritance, in order to exclude presupposition. This rescues the hypothesis of the previous section that in connection with stit action types, deontic inheritance can be based on entailment relations between sentences. The rescue goes at the cost, however, of having an entailment relation that is in need of further elucidation.

It turns out that the paradox of the Good Samaritan can easily be avoided by means of a scope distinction in connection with an analysis in ought-to-do terminology. The paradox threatens to re-occur if the action types in the case are analyzed as stit-type actions. This re-occurrence can in turn be avoided by limiting deontic inheritance to entailment relations between sentences that do not include presupposition.[12]

4 Doing Something in a Particular Way

A distinction that is superficially related to that between a state of affairs or action and its presuppositions is that between performing an action in a particular way and performing it as such. The latter distinction plays a role in, amongst others, the case of the cycling children.

The case of the cycling children [10] is more difficult to handle than that of the Good Samaritan, because it involves proper entailment between action types. The case runs as follows: In London, children, including Alice, ought not cycle on the streets. But if they cycle, they should do so on the left hand side of the road. Suppose that Alice happens to cycle. She should therefore cycle on the left. Cycling on the left involves cycling, and therefore, it seems, if Alice cycles, she ought to cycle, which is in contradiction with the prohibition to cycle.

I will confine my discussion of this case to an analysis in terms of non-stit action types.[13] Given this starting point, the case can be formalized as:

$$\mathrm{Od}(\mathrm{alice}, \neg \mathrm{cycle})$$

[11] In the appendix to this paper, I attempt to analyze the distinction between presupposition and entailment.

[12] The solutions of the Good Samaritan paradox can, with minor adaptations, also be used for the paradox of the Knower [9], which also results from treating presupposition as entailment.

[13] If the case is formalised in terms of stit-actions, it strongly resembles the stit-analysis of the case of the Good Samaritan and should be treated analogously.

Cycles(alice) → Od(alice, cycle_left))
Involves(cycle_left, cycle)
Cycles(alice)

These sentences allow the derivation of

Od(alice, cycle)

which is inconsistent with the first premise.

The problem with this case is the step from ‘Alice ought to cycle on the left’ to ‘Alice ought to cycle’. It seems that not always when some action involves another action, the obligation to perform the first action extends to the second action.

To solve this problem, we must pay attention to an ambiguity in the sentence ‘Alice ought to cycle on the left’. Suppose that Alice is not a child, but a soldier who has received the command to cycle on the left hand side of the road to some building. In this case Alice’s obligation is twofold. Not only should she cycle to the building, but she should do so on the left hand side of the road. In this case there is no problem to derive the obligation to cycle from the obligation to cycle on the left.

The original case is different, because there the obligation to cycle on the left is an obligation to do the cycling on the left. This obligation does not imply the obligation to cycle. Apparently the obligation to do something in a particular way can regard both the main action and the way it is performed, or only the way in which the main action is to be performed.

Let me introduce two new conventions to explicate this distinction. The first convention is that the term denoting an action type can be extended to indicate the way in which the main action is to be performed. This extension is done by indicating the modus operandi between square brackets, as follows: cycle[on_the_left]

The second convention is the introduction of an additional predicate for ought-to-do Od_m, which expresses that the obligation is not directed at the main action, but only at the mode in which the main action is to be performed:

Od_m(alice, cycle[on_the_left])

Given these conventions, the sentence

Od(actor, action[mode])

entails the sentence

Od_m(actor, action[mode])

but not the other way round. The rule that if some action type involves another action type, the obligation to perform the first action type extends to the second action type, only holds for the original Od, and not for the relation Od_m.

The original cycle-case should then be formalized as follows:

Od(alice, ¬cycle)
Od_m(alice, cycle[on_the_left]))
Involves(cycle[on_the_left], cycle)
Cycles(alice)

On this formalization it is not possible anymore to derive

Od(alice, cycle)

which is as it should be.

The second cycle-case (where Alice is a soldier) should be formalized as:

Od(alice, ¬cycle)
Od(alice, cycle[on_the_left]))
Involves(cycle[on_the_left], cycle)

Barring defeasibility this set is inconsistent, because by means of deontic inheritance it is possible to derive

Od(alice, cycle))

The approach to the case of the cycling children proposed here is also available for other cases that seem problematic because of the distinction between doing *something* in a particular way, and doing something *in a particular way*. These other cases include that of the Gentle Murder [11].

5 Ontological Conclusions

It becomes time to look back and draw some conclusions. The first conclusion is that there is no single problem with CTDO's.

The first problem, that arises in connection with the Chisholm-case, is the consequence of treating sentences that express what ought to be done as if they were saying what would be the case in an ideal world. The Chisholm-case brings out clearly that such an analysis of ought-to-do makes it impossible to express what ought to be done in a non-ideal world. However, this is not a problem of CTDO's, but the problem of a wrong analysis of ought-to-do sentences.

The second problem has to do with the role of presuppositions. This problem is exemplified by the case of the Good Samaritan. It arises through deontic inheritance. If the truth of a sentence has certain presuppositions, these presuppositions will be true in all possible worlds in which this sentence is true. In combination with a logic that treats obligations as states of affairs that ought to be the case, this leads to the problem that if a state of affairs ought to be the case, its presuppositions also ought to be the case.

The philosophical solution to this problem is that deontic inheritance should not be combined with presuppositions. If an action type ought to be performed, this ought does not extend to the presuppositions of this action.[14] An ought-to-do of an action type only extends to

[14]Most often these presuppositions will not even be actions, but states of affairs, and the ought-to-do of the action type cannot extend to the presuppositions for that very reason. This explains the attractiveness of Castañeda's distinction between practitions and factual circumstances.

other action types which the obligatory action type involves. An action type does not involve its presuppositions.

Technically this solution can be translated into formalization in two ways. One is the non-occurrence of the presupposition as, or within, a parameter of the ought-to-do-predicate, as is shown by the following analysis of the case of the Good Samaritan:

$$\exists x(\mathsf{Robbed}(x) \rightarrow \mathsf{Od}(\mathsf{samaritan}, \mathsf{help}(x)))$$

This solution is only available if the deontic expressions are not treated as operators on sentences, as is done in SDL.

The second way is not to treat presupposition as a kind of entailment. This solution does not work within SDL, because the semantics of SDL cannot distinguish between presuppositions and entailment.[15] It is available in the ought-to-do approach proposed above, however, because this approach makes use of a dedicated predicate 'involves' that can be defined as to exclude presuppositions, as is shown in the appendix.

Notice that this second problem has again nothing to do with CTDO's. Philosophically it is the problem how to deal with presuppositions, and technically it is a problem of scope or the definition of involvement. CTDO's merely provide a suitable occasion to highlight this problem that already existed independently.

The third problem arises because of an ambiguity in sentences of the form 'X ought to do A in way W'. The obligation on X may concern both doing A and doing it in way W, or it may only concern doing A in way W, without committing X to doing A. This problem underlies the cases of the cycling children and the gentle murder.

On a superficial analysis, this problem can be seen as a case of presuppositions, where doing a in way W presupposes doing A. A more transparent solution is to distinguish between two versions of ought-to-do, that reflect the ambiguity of these sentences. There is an inclusive ought-to-do that concerns both the action type and the way in which it is performed, and an exclusive ought-to-do, which only concerns the way an action type is performed.

Deontic inheritance based on the involvement of action types is only allowed in the case of the inclusive ought-to-do. Cycling on the left involves cycling, but the obligation to cycle on the left only implies the obligation to cycle if the first obligation is inclusive.

Notice that also in this third case, CTDO's merely function as an occasion to bring out a problem that in itself has nothing to do with CTDO's. Actually, I think that there is no intrinsic problem with CTDO's at all. I see no good reason why obligations that arise as a consequence of a violation of another obligation should be any different from other obligations. Their only special characteristic seems to be that they are useful to highlight deficiencies in logical theories that exist independently of them. For this characteristic we should be grateful, because it helps us to improve our logical theories.

Which improvements should we make as a consequence of these seeming problems? First and foremost we should sharply distinguish between sentences that express what is ideally the case and sentences that express what ought to be done. In my opinion it is principally impossible to express what ought to be done in terms of what would ideally be the case, or

[15] On the analysis given in the appendix, both entailed and presupposed states of affairs obtain in all possible worlds in which the entailing or presupposing state of affairs obtains. This suffices for deontic inheritance to work in SDL.

- in general – what would be the case in a particular set of worlds. The fundamental point is that there is an irreducible difference between what ought to be done and what is the case. Therefore, in my opinion, logics that analyze what ought to be done in terms of what is the case in any world, ideal or sub-ideal, are bound to fail. If this failure is not brought to light by CTDO's, there will turn up some other problem.

The second improvement that should be made is that a clear theory should be developed about the legitimacy of deontic inheritance. It seems to me to be beyond doubt that deontic inheritance is acceptable in a large number of case. It seems also beyond doubt that there are cases, such as those which involve presuppositions, in which deontic inheritance is not acceptable.

As yet, I know of no good theory that explains when deontic inheritance is acceptable, and when not. As a first hypothesis I would say that 'deontic' inheritance is acceptable in a logic of what is ideally the case, in the sense that if something is ideally the case, its logical consequences (excluding presuppositions) are also ideally the case. Deontic inheritance is acceptable in a logic for ought-to-do in the case of action types that involve each other. This merely shifts the problem, however, to the question when one action type involves another action type.

The third improvement is to distinguish between an inclusive and an exclusive ought-to-do, and to connect this distinction to the allowance of deontic inheritance. In section 4 I have indicated how this might be done.

Appendix: Entailment and Presupposition

A familiar theory of presupposition runs that if the truth of sentence A presupposes the truth of sentence B, the falsity of A also presupposes the truth of B [12, p. 167f.]. For instance, both 'The king of France is bold' and 'The king of France is not bold' presuppose the truth of the sentence 'France has a king'. Let me use this theory to give an analysis that distinguishes between entailment and presupposition.

I will treat both presupposition and entailment as relations between states of affairs and will exploit the convention about the relation between states of affairs and the sentences expressing them exposed in section 3.1. Given these assumptions, the difference between entailment and presupposition can be characterized as follows[16]:

Presupposition
Presupposes(*a, *b) is true, iff $N(A \rightarrow B) \,\&\, N(\sim A \rightarrow B) \,\&\, \sim N(B)$ is true.
Entailment
Entails (*a, *b) is true, iff $(N(A \rightarrow B) \,\&\, \sim N(\sim A \rightarrow B)) \vee N(B)$ is true.

Intuitively this boils down to that a state of affairs *a presupposes some other state of affairs *b if *b does not obtain in all possible worlds, but does obtain in all possible worlds in which either *a or *$\sim$a obtains.[17] State of affairs *a entails state of affairs *b if either *b obtains in all possible worlds, or if it obtains in all possible worlds in which *a obtains, but not in all

[16] Here the N-operator should be interpreted as 'in all possible worlds it is the case that …'

[17] This presupposes a semantics (interpretation function on sentences) according to which *$\sim$s does not by definition obtain in all possible worlds in which *s does not obtain.

possible worlds in which *$\sim$a obtains. Clearly, it is not possible that a state of affairs both entails and presupposes some other state of affairs

Given the above definitions of entailment and presupposition, it is possible to define an involves-relation between action types that excludes presupposition:

> Involvement
> Involves(action_1, action_2) is true, iff N($\forall x$(Entails(*performs(x, action_1), *performs(x, action_2))) is true.

Given this definition of involvement, cycling on the left involves cycling, but helping a robbed person does not involve anything about robbing.

References

[1] Hohfeld, W.N. Fundamental Legal Conceptions as Applied in Judicial Reasoning, 23 *Yale Law Journal*, 16 (1913).
[2] Lindahl, L. *Position and Change; A Study in Law and Logic*. Reidel, Dordrecht 1977.
[3] Allen, L.E. and C.S. Saxon. A-Hohfeld: A Language for Robust Structural Representation of Knowledge in the Legal Domain to Build Interpretation-Assistance Expert Systems. J-J. Ch. Meyer and R.J. Wieringa (eds.), *Deontic Logic in Computer Science. Normative System Specification*. John Wiley & Sons, London e.a. 1993, p. 205-224.
[4] Chisholm, R.M. Contrary-to-duty Imperatives and Deontic Logic, *Analysis* 24 (1963), p. 33-36.
[5] Bartha, P. Moral Preference, Contrary-to-Duty Obligation and Defeasible Oughts. P. McNamara and H. Prakken (eds.), *Norms, Logics, and Information Systems*, IOS Press, Amsterdam e.a.(1999), p. 93-108.
[6] Lewis, D. Semantic analyses for dyadic deontic logic. S.Stendlund (ed.), *Logical Theory and Semantic Analysis*, Reidel, Dordrecht 1974, p. 1-14.
[7] Smullyan, A.F. Modality and description. *The Journal of Symbolic Logic* vol. 13 (1948) nr. 1, p. 113-127.
[8] Castañeda, H.N., The Paradoxes of Deontic Logic: The Simplest Solution to All of Them in One Fell Swoop. R. Hilpinen (ed.), *New Studies in Deontic Logic*, Reidel Dordrecht 1981, p. 37-86.
[9] Åqvist, L.. Good Samaritans, Contrary-to-Duty Imperatives, and Epistemic Obligations. *Noûs* 1 (1967), p. 361-379.
[10] Prakken, H. and M. Sergot, Contrary-to-duty obligations and defeasible deontic reasoning, *Studia Logica* 57(1) (1996), p. 91-115.
[11] Forrester, J.W. Gentle murder, or the adverbial Samaritan. *Journal of Philosophy* 19 (1984), p. 193-197.
[12] Levinson S.C. *Pragmatics*. Cambridge University Press 1983.