Erkenntnis 83 (2017), pp. 1351-1365

DEONTIC MODALS: WHY ABANDON THE DEFAULT APPROACH

ANDRÉ FUHRMANN

Abstract. John Horty has proposed an approach to reasoning with ought-propositions which stands in contrast to the standard modal approach to deontic logic. Horty's approach is based on default theories as known from the framework of Default Logic. It is argued that the approach cannot be extended beyond the most simple kinds of default theories and that it fails in particular to account for conditional obligations. The most plausible ways of straightening out the defects of the approach conform to a simple theory of default reasoning in standard deontic language.

Keywords. Obligation, ought, conditional obligation, deontic logic, moral dilemmas, normative conflict, default theory, defeasible inference, nonmonotonic reasoning.

1. Introduction

"Why abandon the classical semantics?" asks John Horty [15] in the title of a paper in which he offers a novel approach to deontic reasoning on the basis of his brand of Default Logic and in which he presents a detailed comparison with the standard treatment of deontic logic as modal logic (SDL); see also [9, 10, 11, 12, 13, 14, 16]. The title must be taken with quite a grain of salt though in at least two respects. First, Horty does not propose to abandon the classical approach. In fact he starts by explaining why "the classical semantics fully deserves its status as the dominant approach" [15, p. 425]. Later he distinguishes two readings of ought-propositions, one of them being captured reasonably well by the modal semantics. He also observes in the course of his paper that the classical approach largely coincides with his own on the body of inferences distinguished as logical. Still, as one might expect, there are some and arguably significant differences concerning certain cases. More importantly, by looking at deontic reasoning in the setting of Default Logic, there is something to be learned about the nature of such reasoning understood as a process of balancing out possibly conflicting norms. As Horty presents the lesson, the notion of a reason figures prominently. But the core of Horty's default approach to oughts is independent of his theory of reasons—indeed in earlier work ([9, 10, 11]) this core had beed presented without the notion of a reason playing any systematic rôle. A second aspect in which expectations raised by the title need to be toned down is not explicitly mentioned by Horty: The title must not be understood as if an alternative *semantic* theory of deontic reasoning were in the offer. Default Logic in the way it is usually presented, including the way in which Horty presents it, is by most standards a syntactic logical theory. It sits on top of the classical relation of logical consequence and makes no essential use of semantic notions as, say, preferred models approaches to nonmonotonic inference do. It would be odd to view Default Logic as providing a semantics for anything, and Horty at no point suggests that it does.

Similarly to Horty's guarded use of the "Why abandon it?"-question, I do not wish to claim that there is something fundamentally wrong in expecting insights into deontic reasoning by studying it within the framework of Default Logic—that all such approaches are bound to fail and therefore ought to be abandoned. But I do argue that Horty's particular approach is not promising.

The plan of the paper is as follows. Section 2 rehearses the necessary background in Default Logic and introduces the kind of evaluation rules that play a pivotal role in Horty's approach. In Section 3 I look at simple default theories derived from sets of imperatives and observe that they do not make use of the resources of Default Logic in any substantial sense. In Section 4 I argue that Horty's approach to reasoning with oughts cannot be extended beyond simple imperative default theories, except for a limiting case of little interest. This is the main result of the paper. In Section 5 it is observed that the result continues to hold when Horty's framework is extended to treat conditional obligations. Independently of the criticism in Section 4 it is observed that Horty's account of conditional obligations overgenerates. In Section 6 it is argued that once the conceptual defect in Horty's approach is attended to, it is difficult to resist moving to a much simpler kind of theory. Such a theory, however, presents no rival to SDL.

2. Default Logic

As in logic in general, the subject of *Default Logic* (DL) is the study of certain inference relations. If we let FML denote the set of formulae of a given formal language, then an inference relation is a subset of $\wp(\mathrm{FML}) \times \mathrm{FML}$—or so we may assume for present purposes. (We thus eclipse here the topic of multiple conclusion inference as well as the option of collecting premisses in structures other than sets, such as multisets or sequences.) A default inference relation $\mid\!\sim_D$ depends on a fixed set D of default rules, "defaults" for short. The question as to whether a given pair (A, x) of a set of assumptions $A \subseteq \mathrm{FML}$ and a conclusion $x \in \mathrm{FML}$ is in $\mid\!\sim_D$ is settled by considering the so-called extensions of the default theory (D, A)—more about defaults and extensions in a moment.

We shall consider here propositional default theories only. Thus the assumption set A of a default theory (D, A) consists of formulae in a propositional language with a functionally complete set of Boolean connectives.

Defaults are pairs (a, b) of formulae. (Like Horty, I ignore here the more general notion of a default that has a third coordinate for a condition of application.) I shall write $a \Rightarrow b$ for defaults, calling a the antecedent (or premise) and b the consequent (or conclusion) of the default. As an informal gloss for $a \Rightarrow b$ I suggest the usual: if a, then normally b, or: from a normally infer that b. The latter expresses better than the former that it is the entire conditional that stands within the scope of a normality supposition, and that the conditional must be understood as a rule rather than an if-then sentence grammatically *en par* with, say, a conjunction. For the purpose of calculating extensions defaults function as rules of inference under which assumptions are to be closed if possible.

The study of default inference relations dates back to Reiter's work from the late seventies; see in particular [26]. The present paper contains no introduction to DL. For such an introduction the reader may want to consult Horty's [15] or the work of others, such as [2, 21, 23]. I shall assume that the reader is familiar with the general idea of DL and remembers in outline the standard way, due to Reiter, in which this idea can be formally executed. The central role played by the notion of an *extension* of a default theory (D, A) will thus be recalled. Actually, this is a loose way of speaking. It is not the default theory that is being extended but the assumption set of the theory; thus the Inclusion property stated below. Extension proceeds by trying to close A simultaneously, and as much as possible, under deductive consequence and the rules in D while maintaining consistency (in the principal case that the starting point A is consistent). Since there may be more than one way in which an assumption set can be extended, we may end up with a multi-membered set of extensions of a single default theory (relapsing here and below into the convenient loose talk).

We fix the inference relation $\mid\!\sim_D$, for a given set D of defaults, by quantifying over extensions. As always when quantifying over a domain, there are two salient options: $A \mid\!\sim_D x$ just in case x is in *every*, alternatively in *some* extension of (D, A). The first option is frequently called "sceptical", the second "brave" or "credulous" default inference. I shall here follow Horty in calling the sceptical mode of inference the *disjunctive* and the credulous the *conflict* mode of inference. I use $\mid\!\sim_D$ schematically as standing for either disjunctive or conflict inference. When I need to be specific, I use $\mid\!\sim^d_D$ for disjunctive and $\mid\!\sim^c_D$ for conflict inference. The rationale for these names is as follows. Consider a default theory with two extensions, the one containing x, the other y. Since extensions are closed under consequence (see below), both contain the disjunction $x \vee y$. Thus when we take the meet of conflicting extensions, as for disjunctive inference, conflicting formulae drop out of the conclusions but their disjunctions remain. By contrast, if we take the join of conflicting extensions, as for conflict inference, the formulae in conflict remain present as conclusions. In the literature on DL disjunctive (sceptical) inference is taken to be the by far more interesting relation. Why that should be so depends, of course, on the purposes one wishes DL to serve. It is one of the aims of Horty's approach to ought-reasoning to show that the join-definition of default inference—or the conflict-account, as Horty calls it, see Def. 3 in ch. 3 of [14]—serves a worth-while purpose when trying to capture ought-reasoning under conflicting norms. Horty develops the conflict account to one of the two logics investigated in ch. 3 of [14].

In Reiter [26] extensions are defined in two equivalent ways. In the subsequent literature on DL other definitions have been proposed, frequently combined with a proposal to add more structure to default theories. Horty too has his own way of defining extensions. But we need not go into details here. It will suffice to note some basic properties of extensions shared by all extant definitions of extensions.

If D is a set of defaults, let $Concl(D) = \{b : a \Rightarrow b \in D\}$ be the set of conclusions of D, and let Cn be the operation of taking the classical deductive consequences of a set of formulae. On any definition in DL of an extension E of a default theory (D, A) we have

$$E = \mathrm{Cn}(A \cup Concl(D'))$$

for some subset D' of D. Conveniently we may think of D' as containing the conclusions of those defaults in D that have been used—"triggered", as I shall also say—in the course of generating E. This property of extensions immediately yields two more specific ones. Let $Ext(D, A)$ be the set of extensions of a default theory (D, A). Then for all $E \in Ext(D, A)$ we have the following:

Inclusion	$A \subseteq E$;
Closure	$\mathrm{Cn}(E) \subseteq E$.

Inclusion and Closure are in turn mirrored by properties of the inference relations. As can be easily verified, for both disjunctive and conflict inference we have

Reflexivity	$A \mid\!\sim_D a,\ \forall a \in A$;
Supraclassicality	if $A \vdash x$, then $A \mid\!\sim_D x$.

Let me note in passing that what we characteristically do not have in general is Monotonicity and Substitutivity. From $A \mid\!\sim_D x$ we can neither infer that $A' \mid\!\sim_D x$ for any superset A' of A, nor can we infer that $\mid\!\sim_D$ holds between the results of uniform substitution in A and x. The first fails because adding new assumptions may inhibit the activation of defaults that contributed to generating conclusions in the context of the original set of assumptions. The second fails, briefly, because given Supraclassicality, it would render $\mid\!\sim_D$ trivial (by the Post-completeness of classical consequence, $\vdash$).

Suppose we introduce a predicate ϕ applying to propositions according to the following rule with no restriction on the range of the variable A for assumption sets:

For all A: A supports ϕx iff $A \vdash x$.

What is the property ϕ thus introduced? On salient readings of "support" it is arguably the property of being true. Certainly it is not the property of being tautologous, though it will be, if we restrict the range of A to the domain containing only the empty set. With a suitable restriction in place, ϕ may also distinguish the states of affairs Jack desires to be the case. Just let A exclusively range over collections of propositions Jack wishes to be the case. Then, setting aside reservations about the logical closure of Jack's desires, the optative interpretation of ϕ becomes feasible.

Likewise, and now moving closer to Horty's theory, we may consider the following unrestricted default evaluation rule for ϕ:

(DER1) For all D and A : (D, A) supports ϕx iff $A \mid\!\sim_D x$.

Again, truth offers itself as a candidate interpretation for ϕ. Poole has proposed the property of being predicted as another plausible interpretation of ϕ. As in the case above, by restricting the range over which the variables D and A run other interpretations become available. In general, i.e. when we assume nothing else than that D is a set of defaults and A is a set of assumptions, ϕ stands for what it stands for: see the right-hand-side of the evaluation rule. In particular cases, however, i.e. given particular assumptions about D and A, particular interpretations feeding on these assumptions become legitimate. This includes also deontic interpretations in which ϕ indicates what ought to be the case. Such interpretations will be discussed below. It will turn out that for deontic interpretations to become feasible, restrictions on D and A must be severe—too severe, as I shall argue.

Two Grammatical Remarks: 1. Read straightforwardly, (DER1) fixes the conditions of application (relative to a default theory) of a predicate ϕ to formulae of our formal language, where ϕ is not itself an expression of that language but functions as a schematic variable of our theory-language. Consequently, the relation of support does not hold between a default theory and a formula but between a default theory and the assertion that a formula x enjoys a certain property (indicated by ϕ). Horty sets up things differently. He imports ϕ into the formal language as a unary operator, and writes $\bigcirc$ instead of ϕ. But the operator can neither be nested nor can $\bigcirc$-formulae partake in formulating defaults or assumptions. (Moreover, Boolean combinations of $\bigcirc$-formulae need some extra footwork.) Horty makes it appear as if the truth-functional propositional language has been extended to a language of the type that underlies modal deontic logic. Of course, it hasn't really, and so whenever engaging in comparisons with SDL, Horty needs to repeat that the formation rules for $\bigcirc$-formulae are such that $\bigcirc$ is effectively formally indistinguishable from a predicate applying to names of Boolean formulae.

2. It is standard usage in the literature on DL that $\mid\!\sim$ denotes an inference relation between assumptions (premises) A and a single conclusion x, depending on a set of defaults D. The dependence on defaults is sometimes brought out by writing $(D, A) \mid\!\sim x$ instead of our preferred notation $A \mid\!\sim_D x$ (following Makinson [21]). In combination with what has been

explained in the previous remark, Horty uses $\mid\!\sim$ differently, namely for our relation of support. Thus, his evaluation rule reads:

$$(D, A) \mid\!\sim \bigcirc x \text{ iff } x \text{ is in all/some extensions of } (D, A).$$

On pain of fatally overloading notation, $\mid\!\sim$ cannot now also be used to denote default inference in the standard sense. To be sure, there is nothing incoherent in Horty's way of presenting things. But it may easily mislead readers to misidentify the point at which his theory connects with DL.

3. Simple imperative default theories

As does Horty, I follow here the simple Fregean idea [4] that an imperative is just a proposition (Frege: thought) that is not in the mode of assertion but in the mode of being commanded. If x is a proposition, then $!x$ will represent the imperative to see to it that x be the case. Given a set I of imperatives, we can transform I into a *simple imperative default theory* $(D(I), \emptyset)$ by putting

$$D(I) = \{\top \Rightarrow x : !x \in I\}.$$

Now consider the following default evaluation rule for ϕ:

$$(D(I), \emptyset) \text{ supports } \phi x \text{ iff } \emptyset \mid\!\sim_{D(I)} x.$$

Since all defaults have the trivial premise $\top$, they can be triggered without assumptions present. Two cases are to be considered. First, suppose that the conclusions of $D(I)$ are deductively consistent, i.e. it is not the case that $\{x : \top \Rightarrow x\} \vdash \bot$. Then all defaults will be triggered and $Ext(D(I), \emptyset)$ will contain a single extension, viz. (under normal assumptions) $E = \mathrm{Cn}(\{x : !x \in I\})$. In that case $\mid\!\sim^d_D$ and $\mid\!\sim^c_D$ coincide. Suppose, second, that the conclusions of the defaults are not mutually consistent. Then by triggering some defaults others will be inhibited by a consistency constraint on extensions. All depends on the sequencing of the defaults. Every maximal sequence of consistently triggered defaults will define an extension which is incompatible with any other extension. Thus $Ext(D(I), \emptyset)$ will be multiple and so $\mid\!\sim^d_D$ and $\mid\!\sim^c_D$ will diverge.

In both cases a deontic interpretation of ϕ is very plausible. In the first case ϕ simply records the deductive closure of the propositions that I presents in an imperative mode. So ϕ is the property of what ought to be the case according to I. In the second case the predicate ϕ as generated by the disjunctive relation $\mid\!\sim^d_{D(I)}$ represents a kind of consolidated ought (according to I): what ought to be the case according to I, once inconsistent commands are discounted. If we generate ϕ from the conflict relation $\mid\!\sim^c_{D(I)}$, then it fully mirrors the deductive closures of all maximally consistent subsets of the commanded propositions in I. So ϕ is the property of what ought to be the case according to I in an unconsolidated manner, i.e. allowing for conflicting obligations.

There are two observations here which are worth recording. First, if we restrict the variables D and A in the biconditional schema (DER1) such that D is generated as above from a set of imperatives and A is empty, then the predicate ϕ in the schema is deontically interpretable. Or, as we may say more succinctly, simple imperative default theories carry a deontic interpretation of ϕ. Second, when transforming a set of imperatives into a default theory, we can characterise an ought-predicate via the conflict relation of inference which treats conflicting obligations in a non-trivial way. This contrasts with the treatment of the ought-modality in Standard Deontic Logic (SDL). SDL is just the modal logic KD under a deontic interpretation, with the D-schema $\neg O\bot$ (or, equivalently, $O\neg x \to \neg Ox$) ruling out conflicting obligations.

Let us return to the first observation and add a further one. There is nothing genuinely "defaulty" going on here. We basically have detached the propositional content from a set of imperatives and have generated maximally consistent subsets of these contents. This continues a certain tradition in philosophical logic for dealing with inconsistencies, initiated by Jaśkowski [18] and followed in e.g. Rescher and Brandom [27] and Lewis [19]. The new techniques of the DL-framework were not really tapped. Horty [15, p. 438] writes:

> These [simple imperative default theories] are very simple, of course, but the normative interpretation can be generalised to richer theories as well—theories of the form $(D, A, <)$ in which the hard information from A may not be empty, the defaults from D might have nontrivial premises, and there might be real priority relations among them.

It will now be argued that the generalisation promised in the quote is not possible. Richer default theories typically cannot carry deontic interpretations.

4. Richer default theories

Richer default theories allow to express the idea that some defaults are in some sense better than others, so that the better ones can be identified as those that take precedence over others. To add the requisite expressive power to simple default theories we may either order the defaults in some way, thus treating of triples $(D, A, <)$. Such are default theories with *fixed priorities*. Or we may encode preference (priority) information in the underlying language. In the latter case we would want among our assumptions *formulae* like $d_1 < d_2$, with d_1 and d_2 naming defaults in D. These formulae may in turn serve as premises or conclusions of defaults.

The other way default theories typically are richer than simple imperative default theories is that their assumption sets are non-empty. Moreover, if assumptions are to interact with defaults, i.e. if they are to trigger defaults, defaults must have non-trivial premises. An immediate difficulty in considering non-empty assumption sets resides in the inclusion property of extensions and the resultant reflexivity of $\mid\!\sim_D$ in both the disjunctive and the conflict mode:

$$A \mid\!\sim_D x \; (\forall x \in A).$$

It follows by the default evaluation rule (DER1) for ϕ that ϕa for all $a \in A$ under both the disjunctive and the conflict reading of $\mid\!\sim_D$. If a is just a factual assumption, then a deontic interpretation of ϕ is thus plainly not available. Default theories with factual assumptions cannot carry a deontic interpretation of ϕ. Let us illustrate this with an example.[1]

Suppose Jack, without being obliged to do so, has invited Jill for lunch (p) and let q stand for the proposition that Jack has lunch with Jill. If Jack has invited Jill for lunch, then he normally has lunch with Jill. We may formulate this as the default $p \Rightarrow q$. So $p \mid\!\sim_{\{p \Rightarrow q\}} q$, whence $(\{p \Rightarrow q\}, \{p\})$ supports both ϕp and ϕq. (There are no conflicts here.) Perhaps Jack can be considered in some way to be under an obligation to have lunch with Jill (q). But that cannot be the content of the predicate ϕ, for, by hypothesis, Jack is not obliged to invite her (p). So this particular default theory does not fall into the range of those that carry a deontic interpretation of ϕ. (In Section 6 I shall look at an explanation for why one may be led to think that it does.)

Are there any richer default theories which do fall into that range? Yes. Consider, first, a pair (H, I) of sets of hard (H) and soft (I) imperatives. Let $(D(I), A(H)))$ be the default theory generated from (H, I) by putting

$$D(I) = \{\top \Rightarrow x : !x \in I\} \text{ and } A(H) = \{x : !x \in H\}.$$

This strategy succeeds in blunting the objection from the inclusion property of extensions to deontic interpretations of ϕ. The assumption sets of default theories can thus be populated while reading ϕ deontically. However, as default theories such structures are plainly uninteresting since the triggering of a default—still being of the form $\top \Rightarrow x$—can never depend on the assumptions available.

If non-trivial assumptions and defaults are to interact, then we need *genuine defaults*, i.e. defaults with non-trivial premisses. Suppose then, that $a \Rightarrow b$ is a genuine default in (D, A) and that this default can be triggered. Then there is an extension E of (D, A) such that $a, b \in E$. So by (DER1), (D, A) supports (at least in the conflict sense) both ϕa and ϕb. If (D, A) is to carry a deontic interpretation of ϕ, then neither a nor b can be purely factual. Put contrapositively: If the premise or the conclusion of a default is purely factual, then that default can never be triggered in a default theory carrying a deontic interpretation of ϕ. This finding describes a hard limit to the extent to which the approach introduced by way of simple imperative default theories can be extended to default theories with nontrivially-premised defaults.

[1] In what follows I shall proceed on the assumption that $\mid\!\sim$ must be reflexive. But, as an anonymous referee for this journal has observed, there is a way out of the difficulties raised here, if we give up Reflexivity. If we excluded factual assumptions from the extensions of a default theory, then, first, $\mid\!\sim$ would not be generally reflexive, and, second, plain facts could no longer occur on the right-hand-side of $\mid\!\sim$ and thus enjoy property ϕ.

In particular, a genuine default $a \Rightarrow b$ cannot in general be the result of transforming a conditional imperative into a default. By a 'conditional imperative' is meant an imperative on a condition, such as "Meet her, if you have invited her!". In a conditional imperative, considered on its own, the condition is always factual, since the condition does not itself stand in an imperative mode. However, the context may be such that the condition happens to be commanded ("Invite her!"). In that case we may say that in the conditional imperative the condition is not *purely* factual. Let us consider in the abstract a minimal setting that answers to the case just described. There is a single categorical imperative $!a$ and a conditional imperative $!b/a$. Because of $!a$, $!b/a$ is thus not purely factual. Suppose we transform this into a default theory $(\{a \Rightarrow b\}, \{a\})$. That default theory has the single extension $\mathrm{Cn}(\{a, b\})$ and the deontic interpretation of ϕ according to (DER1) is plausible. So here we have a default theory with a nonempty assumption set and a genuine default that can be triggered by interacting with a non-trivial assumption. But note that we have placed A and D under severe restrictions: All of A is rooted in imperatives and neither the premises nor the conclusions of the defaults in D can be purely factual. Default theories in which A and D satisfy these conditions do carry a deontic interpretation of ϕ. But the conditions limit the default approach to deontic reasoning to borderline cases.

Adding *fixed priorities* to a default theory has the effect of reducing the set of extensions of that theory. Thus, if an unprioritised default theory carries a deontic interpretation of ϕ, then so will every prioritised variant of that theory. Adding fixed priorities to simple imperative default theories cannot present a problem.

By a default theory with *dynamic priorities* we mean a default theory in a language including atomic formulae $d_1 \prec d_2$, where d_1 and d_2 are names of defaults syntactically closed under the binary predicate $\prec$. It is usually required that in assumption sets of dynamically prioritised (d.p.) default theories $\prec$ is irreflexive and transitive, thus making $\prec$ suitable for representing strict preference. But let us start with simple imperative default theories in which the assumption set is empty and ask how default priorities $d_1 \prec d_2$ may enter. Quite simply, they may enter as defaults $\top \Rightarrow (d_1 < d_2)$ by being derived from an imperative $!(d_1 < d_2)$. Without going into the technical details here, it suffices to note that the extensions of a d.p. simple imperative default theory are calculated by considering a family of simple imperative default theories with fixed priorities. These are, as we have seen, unproblematic w.r.t. their deontic interpretability.

As to be expected, problems arise when we start populating the assumption sets of prioritised default theories with priority-propositions and include priority defaults with non-trivial premises. We can now repeat for priority-propositions what we have observed a moment ago about propositions in general in the role of assumption or in the role of premises or conclusions of defaults. In particular, if a priority $d_1 < d_2$ is to trigger a default, then that priority must fall under the ϕ-predicate (at least under the conflict reading of $\mid\!\sim$ in DER1). Thus, the priority in question cannot be purely factual.

But wasn't the idea to use priorities (preferences) as a basis for determining what ought to be the case? Now it turns out that priorities can only serve that purpose, if they themselves ought to be the case. This is, to say the least, a surprising commitment.

5. Conditional obligation

Simple obligations are only a stepping stone towards the more general notion of a conditional obligation. Horty proceeds to the latter presupposing that his approach to simple obligations as generated by simple imperative default theories "can be generalised to richer theories as well" (witness the quote above). I have argued that this presupposition fails. Our findings continue to hold when moving to conditional obligations. In this section I shall point out specific problems with Horty's account of conditional obligation on the assumption that the presupposition holds. Here is Horty's default evaluation rule for what should turn out to be conditional oughts:

(DER2) (D, A) supports $\phi(b/a)$ iff $(D, A \cup \{a\})$ supports ϕb.

As in the case of (DER1) the schema has two instances, depending on whether support (which we now abbreviate to a colon, :) is meant in the disjunctive or in the conflict sense:

$$(D, A) :_d \phi(b/a) \text{ iff } A, a \mathrel{|\!\!\sim}^d_D b\,;$$
$$(D, A) :_c \phi(b/a) \text{ iff } A, a \mathrel{|\!\!\sim}^c_D b\,.$$

A grammatical note of caution (familiar from a similar issue in modal deontic logic): ϕ is not to be understood as a predicate applying to a formula b/a. There is no such formula and, in particular, there is no binary connective ./. applying to formulae. Rather $\phi(./.)$ is a binary predicate applying to formulae. (If we were to follow Horty in treating $\phi(b/a)$ as a formula of our formal language, $\phi(./.)$ would have to be a binary connective.)

The grammatical note brings us immediately to the question of the intended meaning of $\phi(a/b)$. The short paraphrase is: a is obligatory given that b. Now, as Horty observes, different theories of the concept of a conditional obligation need not be rival to each other; they may answer to different precisifications of the short paraphrase. Horty identifies two, the constrained optimality sense and the resultant sense of conditional oughts. (As in the DERs, I substitute ϕ for Horty's operator $\bigcirc$.)

- *Constrained optimality sense*: "The statement $\phi(a/b)$ is taken to mean, very roughly, that a holds in the best worlds in which b holds, even if these worlds are not among the best overall. It is this sense of the conditional ought that was originally explored by Bengt Hansson, in early papers by van Fraassen, Chellas and David Lewis [...]." [15, p. 448]
- *Resultant sense*: "The statement $\phi(a/b)$ is taken to mean, again very roughly, that the various prima facie norms—or imperatives, or reasons—at work under the condition b interact in a way that results in overall support for a. And it this sense that is captured by the framework of default logic." [15, p. 448]

Now we ask: Is (DER2) consonant with the resultant sense of conditional oughts? I shall argue that (DER2) overgenerates when interpreted in the resultant sense. (I use the unsubscribed colon if an assertion holds for support in the disjunctive as well as in the conflict sense.)

First observation: $(D, A) : \phi(a/a)$. This is immediate from the Reflexivity of $\mid\!\sim_D$. Let us try to gloss this in the resultant sense: The prima facie norms in effect under condition a favour a. That might be so in particular cases (let a be giving to charity) but it does not hold in general (let a be robbing a charity). Incidentally, where O stands for the ought-modality, the theorem $O(a/a)$, ubiquitous in conditional deontic logics, is a paradigm case for wheeling out the constrained optimality interpretation. Bengt Hansson [7, p. 144], e.g., comments: "One may think of obligations relative to circumstances a as obligations in a restricted universe: of all the a priori possible worlds only those in a are now achievable; therefore a plays the role of the universe and $O(a/a)$ only says that at least something is obligatory under a."

Second observation: If $a \to b \in A$, then $\forall D$, $(D, A) : \phi(b/a)$. For, assume that $a \to b \in A$. Then $A, a \vdash b$ and so by Supraclassicality $A, a \mid\!\sim_D b$ for any D. Thus, by (DER2), for any D, $(D, A) : \phi(b/a)$. Here is a counterexample: Assume that if Peter spits on the floor, then a bystander feels offended $(a \to b)$. So (?), whatever the moral defaults may be, the norms that are activated by Peter's spitting on the floor are such that a bystander ought to be offended. It is difficult to decide whether this is wrong or devoid of sense; in any case, it is untrue. The closest analogue of this in terms of a formula of deontic logic would be $(a \to b) \to O(b/a)$—which is not a theorem in any known deontic logic. This observation only prepares for worse to come.

Third observation: If $b \in A$, then $\forall D$, $(D, A) : \phi(b/a)$. This follows immediately from Reflexivity: $A, a \mid\!\sim_D b$ $(\forall b \in A)$. If b is assumed to be true, then b is obligatory whatever the circumstances (and the defaults).

Fourth observation: If $a \in A$, then $\forall D$, $(D, A) : \phi(b/\neg a)$. For, under the assumption we have $A, \neg a \vdash b$. So by Supraclassicality $A, \neg a \mid\!\sim_D b$, for any D. Anything becomes obligatory under circumstances that contradict the current assumptions (no matter what the defaults are). This echoes an observation made by Horty himself in [9] but left without comment there.

Again, the closest modal deontic analogues for the last two observations about conditional ϕ are $b \to O(b/a)$ and $a \to O(b/\neg a)$. Of course, none of these are theorems in any known conditional deontic logic.

To summarise: Even if the assumption could be discharged that rich default theories can be given a deontic interpretation in the way Horty suggests, the default evaluation rule for conditional ϕ overgenerates to an extent that calls into question any interpretation in terms of what ought to be the case; this is particularly so, if the intended resultant sense is to be distinct from the constrained optimality sense of conditional obligation.

6. Why not abandon the default approach

Why is it that Horty apparently so effortlessly succeeds in capturing example scenarios in default theories with deontic import? Let us consider one such example, the twin example, which occurs in various places in [14] and [15].

> I have inadvertently promised to have a private dinner tonight with each of two twins, Twin 1 and Twin 2. In that case, it is natural to say that I have an "obligation" to dine with Twin 1, due to my promise, and also with Twin 2. It can likewise be said that I have a "prima facie duty" to dine with Twin 1, and also with Twin 2, or that dining with Twin 1 is something I ought to do "other things being equal", as is dining with Twin 2. [14, p. 67]

There is nothing objectionable here. But let us see how the example is now transferred into a default theory. (It is supposed to be part of the above example, not mentioned in the quote above, that I will not dine with both twins. In the citation below some letters are changed so as to keep with our notation.)

> Take a_1 and a_2 as the propositions that I have arranged to dine with Twins 1 and 2, respectively, and d_1 and d_2 as the propositions that I do in fact dine with Twins 1 and 2, respectively. The example can then be encoded in the default theory $(D, A, <)$, where $A = \{a_1, a_2, \neg(d_1 \wedge d_2)\}$, where $D = \{a_1 \Rightarrow d_1, a_2 \Rightarrow d_2\}$, [...] and where the ordering $<$ is empty; neither default has a higher priority than the other. [14, p. 70]

Note that so far we have only captured the bare facts of the case in a default theory: That I have arranged, or promised, to dine with the twins, that I will not dine with both of them, and that if I have made such an arrangement with one of the twins, then I normally will meet that twin for dinner. There is no information here as to what I ought to do. And what is not "in" the specification of a default theory we cannot get "out" of it by calculating its extensions. Horty sees things differently. When we apply (DER1) in the disjunctive mode, then

> ... the disjunctive account yields $\bigcirc(d_1 \vee d_2)$ as an all considered ought. In the case of this example, then, where I have arranged to dine with each twin but cannot do so, rather than telling me that I ought to dine with Twin 1 and I also ought to dine with Twin 2, and so face a moral conflict, the disjunctive account tells me only that what I ought to do, all things considered, is dine with one twin or the other. [14, p. 73f.]

But the disjunctive use of (DER1) only yields the conclusion that $d_1 \vee d_2$ is distinguished by the property ϕ (Horty's suggestive $\bigcirc$) of being an element in all extensions of the example default theory. What part of the specification of that theory makes it the case that I *ought* to dine with one of the twins, rather than that I will, or that I wish, or—for that matter—that I ought *not* to?

For an answer one is naturally tempted to turn to the first quote above: to what is "natural to say". But when we look at the specification of the default theory in a straightforward way, we do not find what is natural to say among what has been said, i.e. among what has been recorded in that theory. We find the factual default that if I arranged to dine with Twin 1, then I normally dine with Twin 1. We do not find what may be natural to say, viz. the deontic default that if I arranged to dine with Twin 1, then I *ought* to dine with 1.

Perhaps, things are not meant to be as straightforward as they seem. Perhaps in the application at hand defaults $a \Rightarrow d$ are not meant to have a factual reading (as in "d is normally the case, given a") but a deontic one, something like: "d normally *ought* to be the case in circumstances a". If so, then the "ought" in the latter sentence obviously plays a systematic rôle in the theory. There are two ways in which it can play this rôle: either by d being of the form "it ought to be the case that b" (short: Ob), or by thinking of $a \Rightarrow d$ as a primitive deontic conditional (which we may alternatively write $O(d/a)$ to distinguish it from factual defaults). In the latter case, on the one hand, Horty's theory ends up with two notions of conditional oughts: the one used to encode deontic defaults, and the one delivered by (DER2). Whatever the relation between the two notions may be, the evaluation rôle (DER2) delivering $\bigcirc$-conditionals clearly depends on a prior understanding of O-conditionals.

If, on the other hand, we consider defaults of the form $a \Rightarrow Ob$ we are really entering a very different game, which may be best described as follows.[2] Consider an interpreted modal language with a unary operator O stipulated to express moral obligation. Then we represent the twins example in a default theory thus: a_1 and a_2 as well as d_1 and d_2 are as before. The initial assumptions are that I have separately arranged for dinner with the two twins. So a_1 and a_1 are in A as before. The defaults are that if a_i, then normally I ought to d_i ($i = 1$ or $i = 2$). So

$$D = \{a_1 \Rightarrow Od_1, a_2 \Rightarrow Od_2\}.$$

What I normally ought to do can be overridden. In our example it can be overridden by the assumption that, whatever I normally ought to do, I do not take myself obliged to have dinner with both twins, $\neg O(d_1 \wedge d_2)$. So, here is the complete set of assumptions:

$$A = \{a_1, a_2, \neg O(d_1 \wedge d_2)\}.$$

(In specifying the default theory we may ignore a priority ordering, since in the example it is empty.) Now we calculate the extensions of (D, A). In doing so we use a consequence operation Cn' based on whatever modal logic we assume to go with O. Then we have two extensions, depending on the order in which we use the defaults:

$$E_1 = \mathrm{Cn}'(A \cup \{Od_1\}) \text{ and } E_2 = \mathrm{Cn}'(A \cup \{Od_2\}).$$

[2] Prakken [24, p. 78] poses the question: "Why not use default logic as it is, with the only change that the language on which it is based, first-order predicate logic, is extended with modal deontic operators?" In response he develops a version of the theory sketched here.

So in the disjunctive mode of inference (where we take the meet of extensions) we have only $A \mid\!\sim^d_D Od_1 \vee Od_2$, whereas in the conflict mode (where we join the extensions) we have (*) both $A \mid\!\sim^c_D Od_1$ and $A \mid\!\sim^c_D Od_2$. Now there is a clear sense in which (*) is in conflict with the fact that (†) $A \mid\!\sim^c_D \neg O(d_1 \wedge d2)$. But $\mid\!\sim^c$ is not closed under Adjunction, i.e.

$$\frac{A \mid\!\sim^c_D x \qquad A \mid\!\sim^c_D y}{A \mid\!\sim^c_D x \wedge y}.$$

If it were, we could get from (*) that $A \mid\!\sim^c_D Od_1 \wedge Od_2$. Since $\mid\!\sim^c$ continues over $\vdash$ (Right Weakening in the literature), where $\vdash$ is the consequence relation that calculates the extensions, we would then have $A \mid\!\sim^c_D O(d_1 \wedge d_2)$. Together with (†) and, again, Adjunction that would give us $A \mid\!\sim^c_D \bot$. So, again by continuation over $\vdash$, $\mid\!\sim^c_D$-inference from A would be trivial. But all this is an empty threat, since our very example falsifies Adjunction: $Od_1 \wedge Od_2$ is neither in E_1 nor in E_2. So the conflict mode of inference registers the conflict of obligations in a non-trivial way.

This theory is not without problems of its own. Although the theory satisfies the principle of (factual) detachment, with $D = \{a \Rightarrow Ob\}$ we have $a \mid\!\sim_D Ob$, it fails deontic detachment: $Oa \not\mid\!\sim_D Ob$. Not surprisingly, deontic chaining will also fail: With $D = \{a \Rightarrow Ob, b \Rightarrow Oc\}$ we will not get $a \Rightarrow_D Oc$.[3] Note that deontic detachment is a rather suspicious principle. This shows particularly when considering contrary-to-duty obligations. Briefly (and schematically), the problem is as follows.[4] One may have (a) a duty to see to it that a; (b) the state of affairs a in turn creates an obligation that b, while (c) the contrary-to-duty state $\neg a$ creates an obligation that $\neg b$ be the case. This appears to be a set of logically independent rulings which should be compatible with (d) $\neg a$ being the case. Call the complete set, including the fact $\neg a$ a *Chisholm-set*. The challenge is to give a formal rendering of this set with the desired property that all elements are mutually independent and that the set does not generate inconsistent obligations.[5] We can encode a Chisholm-set in a default theory with $A = \{\neg a\}$ and $D = \{\top \Rightarrow Oa, a \Rightarrow Ob, \neg a \Rightarrow O\neg b\}$. We then have $A \mid\!\sim_D O\neg b$. In the presence of deontic detachment we would also have $A \mid\!\sim_D Ob$ — (A, D), thus failing to meet the challenge of consistency. But without deontic detachment no conflicting obligations can be generated. Are the items in $A \cup D$ mutually independent? On certain renderings of the conditional (b), $\neg a$ entails *if a then Ob*. But here we have encoded (b) by $a \Rightarrow Ob$, and there is no general connection between A and D that would logically force items into D on the basis of items present in A. So the challenge of independence can also be met in this way.[6]

[3] The two detachment principles are so difficult to combine that they may be used, as in [20], to subdivide dyadic deontic logics into two distinct families; see also the survey papers [1, 3]. Thanks to an anonymous referee for this journal for raising my attention to this point.

[4] For a far more complete exposition of this intricate topic see [3].

[5] See [3] for examples, a statement of the conditions for an adequate solution and a discussion of standard attempts at a solution.

[6] Note that we have represented the schematic Chisholm-set in a default theory with a

Whatever the merits of this or a similar theory may be, the sketch suffices to show what kind of theory will emerge, if we consider deontically loaded defaults to be of the form $a \Rightarrow Ob$. Deontic reasoning emerges as embedded within standard DL without a detour via a default evaluation rule such as (DER1). Moreover, the method is capable of delivering the kind of non-trivial conflict registration and resolution that Horty's approach is specifically designed for. Note that SDL is required as a background theory (for calculating extensions). It would thus make little sense to present this approach to reasoning with oughts as a rival to SDL. Whereas SDL is a theory of *deductive* reasoning with oughts, the topic of the approach sketched above is, not surprisingly, *defeasible* reasoning with oughts, thus keeping fully with the design idea of DL.

Acknowledgements. Part of this research was carried out at the Federal University of Rio de Janeiro in March and April 2015 while holding a grant from the FAPERJ (Brazil). I wish to thank Wilson Mendonça for stimulating discussions on the topic of reasons which soon drew my attention to Horty's work. Thanks are also due to Richard Dietz, Ali Esmi, Sven Ove Hansson, Dominik Kauß, and an anonymous referee for this journal for many helpful comments.

References

[1] Åqvist, L. Deontic logic. In Gabbay and Guenthner [5], pp. 147–264.

[2] Brewka, G., and Eiter, T. Prioritizing default logic. In *Intellectics and Computational Logic – Papers in Honour of Wolfgang Bibel*, S. Hölldobler, Ed., vol. 19 of *Applied Logic Series*. Kluwer, 2000, pp. 27–45.

[3] Carmo, J., and Jones, A. Deontic logic and contrary-to-duties. In Gabbay and Guenthner [5], pp. 265–343.

[4] Frege, G. Der Gedanke. *Beiträge zur Philosophie des deutschen Idealismus 2* (1918), 58–77.

[5] Gabbay, D. M., and Guenthner, F., Eds. *Handbook of Philosophical Logic (2nd Ed.)*. Springer, 2001ff.

[6] Ginsberg, M. L., Ed. *Readings in Nonmonotonic Reasoning*. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1987.

[7] Hansson, B. An analysis of some deontic logics. *Nous 3* (1969), 373–398. Reprinted in [8], pp. 121-147.

[8] Hilpinen, R., Ed. *Deontic Logic: Introductory and Systematic Readings*. Reidel, Dordrecht, 1981.

single extension. There is thus no reliance here on the hypothesis that the consistency of Chisholm-sets essentially depends on the defeasibility and indeed defeat of some of its members; see e.g. [22]. I take that hypothesis to be implausible for reasons explained in [25].

[9] Horty, J. F. Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic 23*, 1 (1994), 35–65.

[10] Horty, J. F. Nonmonotonic foundations for deontic logic. In *Defeasible Deontic Logic: Essays in Nonmonotonic Normative Reasoning*, D. Nute, Ed. Kluwer, 1997, pp. 17–44.

[11] Horty, J. F. Reasoning with moral conflicts. *Noûs 37*, 4 (2003), 557–605.

[12] Horty, J. F. Defaults with priorities. *Journal of Philosophical Logic 36* (2007), 367–413.

[13] Horty, J. F. Reasons as Defaults. *Philosophers' Imprint 7* (2007), 1–28.

[14] Horty, J. F. *Reasons as Defaults*. Oxford University Press, New York, 2012.

[15] Horty, J. F. Deontic modals. Why abandon the classical semantics? *Pacific Philosophical Quarterly 95* (2014), 424–460.

[16] Horty, J. Norm Change in the Common Law. In *David Makinson on Classical Methods for Non-Classical Problems*, S. O. Hansson, Ed. Springer Science+Business Media Dordrecht, 2014, ch. 15, pp. 335–355.

[17] Jaśkowski, S. Rachunek zdań dla systemów dedukcyjnych sprzecznych. *Studia Societatis Scientiarun Torunesis (Sectio A) 1* (1948), 55–77.

[18] Jaśkowski, S. Propositional Calculus for Contradictory Deductive Systems (Engl. tansl. of [17]). *Studia Logica 24* (1969), 143–157.

[19] Lewis, D. Logic for equivocators. *Nous 16* (1982), 431–441.

[20] Lower, B., and Belzer, M. Dyadic deontic detachment. *Synthese 54* (1983), 295–318.

[21] Makinson, D. *Bridges from Classical to Nonmonotonic Logic*. King's College Publications, London, 2005.

[22] McCarty, L. T. Defeasible deontic reasoning. *Fundamenta Informaticae 21* (1994), 125–148.

[23] Poole, D. Default logic. In *Handbook of Logic in Artificial Intelligence and Logic Programming*, D. M. Gabbay, J. Hogger, and J. Robinson, Eds., vol. 3: Nonmonotonic Reasoning and Uncertain Reasoning. Clarendon Press, Oxford, 1994, pp. 189–215.

[24] Prakken, H. Two approaches to the formalisation of defeasible deontic reasoning. *Studia Logica 57*, 1 (1996), 73–90.

[25] Prakken, H., and Sergot, M. Contrary-to-duty obligations. *Studia Logica 57*, 1 (1996), 91–115.

[26] Reiter, R. A logic for default reasoning. *Artificial Intelligence 13* (1980), 81–132; reprinted in [6, 68–93].

[27] Rescher, N., and Brandom, R. *The Logic of Inconsistency.* Blackwell, Oxford, 1980.

ANDRÉ FUHRMANN
INSTITUT FÜR PHILOSOPHIE
GOETHE-UNIVERSITÄT
60629 FRANKFURT a. M.

fuhrmann@em.uni-frankfurt.de