NEW FOUNDATIONS FOR DEONTIC LOGIC: A PRELIMINARY SKETCH

Peter B. M. Vranas
vranas@iastate.edu
Iowa State University
13 May 2002

Abstract. I outline six components of a comprehensive proposal for overhauling the foundations of deontic logic. (1) Actions and prescriptions are temporally indexed; more precisely, they attach to *nodes* of a *tree* in a branching time structure. (2) *Actions* are (modeled as) sets of branches and can be *coarse*- or *fine*-grained depending on whether or not they have proper subsets which are also actions. (3) *Prescriptions* have *satisfaction* and *violation* sets; these are sets of branches which may—but need not—be or include actions. (4) *Prescriptive propositions*, which state that an action is obligatory/permitted/forbidden *according* to a given prescription, are defined by relating the action with the satisfaction and violation sets of the prescription. (5) *Conditional* prescriptions can—but need not—be *derived* from unconditional or even from other conditional ones. (6) *Thick* prescriptions, in contrast to *thin* ones, prescribe or proscribe actions with varying *intensities*, and can have embedded *subprescriptions* (some of which are *negative*, namely "contrary-to-duty"). Most of the above components are inspired by the literature, but their combination is novel.

0. Introduction

Deontic logic is a thriving research field. Nevertheless, it remains true that "there is virtually no single issue in [the field] upon which a settled consensus has been reached" (Åqvist 1984: 607, 1987: 8). This paper has the ambitious aim to provide the beginning of a remedy for this undesirable state of affairs.

I should make clear at the outset two general limitations of the paper. (1) As the title indicates, the paper provides only a preliminary sketch of my comprehensive proposal for overhauling the foundations of deontic logic; it is intended as a prelude to a more thorough and careful treatment. (2) The paper has a severely circumscribed scope. It deals with ought to do but not with ought to be;[1] with actions but not with strategies;[2] with prescriptions that attach to a particular moment but not with those that attach to time periods;[3] and it ignores

[1] On the relation between ought to do and ought to be see: Brown 1999: 115-6; Chisholm 1964: 150; Geach 1982; Harman 1986: 131-2; Horty 2001: chap. 3; Krogh & Herrestad 1996; von Wright 1981a: 9, 1981b: 409-11.

[2] On strategies see: Belnap, Perloff, & Xu 2001: chap. 13; Horty 2001: chap. 7; von Kutschera 1986: 274-9.

[3] On the temporal dimensions of prescriptions see: Åqvist 1997: 60; Brown 1996, 1999, 2000; Hansson 1969/1971: 141; Kamp 1979: 285 n. 3.

probabilities, multiple agents,[4] and conflicts of prescriptions.[5] I will drop some hints, however, on how my proposal can overcome some of these restrictions.

My proposal integrates in a novel way a number of ideas most of which are more or less well known in the literature. The starting point is a branching time structure (Belnap et al. 2001) or *tree framework* (§1). Actions are indexed to nodes of the tree and correspond to sets of histories (or branches). *Coarse*-grained actions, unlike *fine*-grained ones, have proper subsets that are also actions (§2). Prescriptions are also indexed to nodes of the tree and are characterized by a *satisfaction set*, namely the set of histories that satisfy the prescription, and a *violation set* (§3). This simple characterization of prescriptions is the heart of my proposal. It enables me to generalize Anderson's (1956/1967, 1958) reduction of deontic to alethic modal logic: a violation set is a *prescription-specific* version of Anderson's "sanction". It turns out that prescriptive propositions, which state that an action is obligatory/permitted/forbidden *according* to a given prescription, can be reduced to set-theoretic propositions relating the action with the satisfaction and violation sets of the prescription (§4). (The reduction captures both von Wright's (1968) distinction between weak and strong permission propositions and Brown's (1996) distinction between type 1 and type 2 obligation propositions.) My characterization of prescriptions enables me also to give a unified treatment of conditional and unconditional prescriptions: a prescription is *unconditional* exactly if its violation set is the complement of its satisfaction set, and is *conditional* (on the union of its satisfaction and violation sets) otherwise (§5). Finally, a generalization of my characterization of prescriptions allows of embedded *subprescriptions*, some of which correspond to "contrary-to-duty" cases (§6). I conclude in §7.

1. The tree framework

The intuitive picture is that of a decision-theoretic tree: at each *decision* node an agent is faced with a number of choices, each choice corresponding to a branch that leads to a further node, while at each *chance* node a random process has a number of possible outcomes, each outcome corresponding to a branch (associated with a certain probability) that leads to a further node. This picture, unfortunately, is based on some more or less questionable assumptions. Namely, that time is discrete; that branches have finite lengths; that nodes correspond to finitely many branches; and that a single agent is making choices. Belnap, fortunately, has elaborated a theory of "branching time structures" which avoids

[4] On multiple agents see: Bailhache 1991: chaps. 5, 8; Belnap et al. 2001: chap. 10; Feldman 1986: chap. 7; Horty 2001: chap. 6; Tännsjö 1989; Tuomela 1989a, 1989b.
[5] On conflicts of prescriptions see especially: Brink 1994; Horty unpublished.

these assumptions without substituting other ones.[6] In Belnap's theory it is left open whether time is discrete, continuous, or mixed; whether branches have finite or infinite lengths; whether nodes correspond to finitely or to infinitely many branches; and whether one or more agents are making choices (Belnap et al. 2001: 342-3). In the next paragraph I give a very brief overview of Belnap's theory; readers uninterested in technical details may just skim that paragraph.

A *tree* is a nonempty set of *moments*, partially ordered by a reflexive, transitive, and antisymmetric relation ≤ (not later than). Two moments m_1 and m_2 are *comparable* exactly if $m_1 \leq m_2$ or $m_2 \leq m_1$. A *history* is a maximal set of comparable moments; i.e., a set of pairwise comparable moments no proper superset of which is a set of pairwise comparable moments. It is assumed that (1) every two moments m_1 and m_2 have a lower bound (i.e., a moment m_3 with $m_3 \leq m_1$ and $m_3 \leq m_2$; so every two histories intersect), and that (2) no two incomparable moments m_1 and m_2 have an upper bound (i.e., a moment m_3 with $m_1 \leq m_3$ and $m_2 \leq m_3$; so there is no backward branching of histories). These two assumptions make the tree "treelike". Two histories *split* at a moment exactly if the moment is the least upper bound of their intersection. A *node* is a moment at which histories split.[7]

2. Actions

I understand an *action* as a (possible) action *token* ("individual action"), like the murder of Caesar by Brutus, not as an action *type* ("generic action") like murder (von Wright 1951: 2, 1963: 36, 1981a: 15). So an action is always indexed to a node[8] and to an agent. (Or to a group of agents. In this paper to simplify I consider a single agent and I omit reference to agents.) An action at node n (an *n-action*) corresponds to (or, as I will say, *is*) a non-empty set of histories each of which contains n (*n-histories*). Not every non-empty set of n-histories is an n-action, however. For example, the set H_n of *all* n-histories is not an n-action: it contains histories at which the agent falls asleep or dies (and thus performs no

[6] Belnap 1991, 1996; Belnap et al. 2001: chap. 7; Prior 1967; Thomason 1970, 1981a, 1981b, 1984.

[7] Zanardo (1985, 1991, 1996; cf. Brown 1999: 110-3) has developed an alternative branching time theory which takes branches rather than moments to be primitive. Zanardo's theory is more general than Belnap's but may be less adequate as a description of our world (Belnap et al. 2001: 199-201). I think that the differences between the two theories don't matter for my purposes.

[8] A *strategy*, by contrast, is indexed to a *set* of nodes: it is a function from nodes in that set to actions indexed to those nodes. I don't examine strategies in this paper.

action).[9] It follows that tautologies, which correspond to H_n, correspond to no *n*-action.[10]

I assume that the union of any *n*-actions is an *n*-action: if one can perform *n*-actions A_{n1} and A_{n2}, then one can perform A_{n1} *or* A_{n2}. It follows that the union of *all* *n*-actions is itself an *n*-action; call it U_n, the *universal* *n*-action. U_n is distinct from H_n, the set of all *n*-histories: the latter, as I said, is not an *n*-action. So although U_n is an action that the agent performs no matter what she *does*, it is not an action that she performs no matter what *happens*. I also assume that the complement with respect to U_n (henceforth simply *complement*) of any *n*-action other than U_n is an *n*-action. The complement of an *n*-action A_n is the *n*-action of *refraining* from A_n. It is well known that refraining from A_n, which is an action, should be distinguished from *not performing* A_n, which is not an action.[11] This distinction is captured in my theory by noting that not performing A_n corresponds to the complement of A_n with respect to H_n (not U_n), and is thus not an action because contains histories at which the agent falls asleep or dies. The complement of U_n, namely the empty set, is not an action; the *n*-action of "doing nothing" is—like every *n*-action—a *non*-empty set of *n*-histories. It follows from the above assumptions that the non-empty intersection of any *n*-actions is an *n*-action.

A distinction crucial for what follows and for deontic logic in general is that between coarse- and fine-grained actions. An *n*-action is *coarse-grained* exactly if some proper subset of it is an *n*-action (i.e., exactly if the action has a proper *subaction*), and is *fine-grained* otherwise (i.e., exactly if it has no proper subaction).[12] A paradigmatic example of a coarse-grained action is a disjunctive action: turning right *or* left. The disjuncts, however, are normally also coarse-grained: one can turn left slowly or speedily. Indeed, almost *every* action specified in ordinary discourse is coarse-grained: it has proper subactions. If I pay you ten dollars, I can pay by cash, check, or money order; if I pay by cash, I can give you a ten, fives, or ones; if I give you a ten, I can give you this one, that one, or the other one; if I give you this one, I can hand it to you face up, down, right, or left; and so on (Chisholm 1963: 4; cf. Brown 1996: 50). This

---

[9] As another example, Belnap et al.'s postulate of "no choice between undivided histories" (2001: 216-7) corresponds to the claim that no set is an *n*-action if it contains only one of two *n*-histories that split at a node properly *later* than *n*.

[10] So the claim that H_n is not an *n*-action is a simple analog of stit theory's "negative requirement" (Belnap et al. 2001: 37).

[11] Cf. Belnap et al. 2001: 40-3; Brown 2000: 91-2; Horty 2001: 25-9; von Wright 1963: 45, 1988: 45.

[12] Belnap et al. never mention this distinction with respect to actions (they talk in effect only about fine-grained actions), but they vaguely realize its necessity with respect to strategies (2001: 368-9).

"decomposition" process, however, must stop before the "bottom": before the level of single histories. This is because "no agent's action at a given moment will determine the whole future course of the universe in all its details. To say [otherwise] would be to deny the role of other agents, of chance, and indeed of the agent's capacity for further choice in the future" (Brown 1999: 114). So fine-grained actions correspond to the limit of the agent's ability to make a difference to the state of the world and are not *maximally specific*: they don't contain just one history. (A maximally specific action, on the other hand, if such a thing existed, would be fine-grained: by definition it would contain only one history, so it would have no non-empty proper subset and thus no proper subaction.)

Contrary to what is standardly assumed,[13] the set of all fine-grained n-actions does not *partition* the set H_n of all n-histories. One reason is that some n-histories, like those at which the agent falls asleep or dies, belong to no n-action. Another reason is that fine-grained n-actions can overlap. Suppose, for example, that an agent (or a robot) can raise her right arm with a precision of 2^o. If she tries to raise it at an angle of 45^o, she raises it at an angle between 44^o and 46^o; if she tries to raise it at an angle of 44^o, she raises it at an angle between 43^o and 45^o; so histories at which she raises it at an angle of 44.5^o belong to at least *two* fine-grained actions. This example suggests that in general it makes no sense to talk of *the* fine-grained n-action that corresponds to a given n-history. More generally, it makes no sense to talk of *the* action that the agent performs: if she performs A_n, then she performs every action that is a superset of A_n. For example: if she makes a left turn, then she turns.[14] It does not follow, however, that she turns *or* makes the sun shine:[15] not *every* superset of an action is an action.

---

[13] E.g., by Belnap et al. (2001), Brown (1996: 54, 1999: 117, 2000: 98), Horty (2001: 12), and von Kutschera (1986: 267). Belnap et al. (2001: 214-5), however, admit that the assumption may be mistaken and encourage the development of alternative theories.

[14] Contrary to what Brown (2000: 90) suggests, the relation between these two actions is not mysterious: the latter is a proper superset of the former. So I don't see a need for what Brown calls an *outcome* operator, which would "give one of the *consequences*" of performing an action, in addition to what he calls an *action* operator, which would "give a full and precise expression" of the action. On the other hand, I strongly agree with Brown that we need two *obligation* operators (see §4.3).

[15] Stit theory in general cannot avoid this unwelcome implication (though the difference between the deliberative and the achievement stit matters here: Belnap et al. 2001: 47), so its "negative requirement" only partially solves the problem it is intended to address. My theory is both simpler and more effective. One might argue that my theory has less expressive power, since it does not allow for iterations. But even actions that correspond to iterations of stit operators are sets of histories and thus can be *expressed* in my theory. Admittedly, however, the *structure* of such actions is more explicit in stit theory.

3. Prescriptions

We now get closer to deontic logic. First, some metaphysical remarks. I understand a *prescription* as an abstract entity, analogous to a proposition. Declarative sentences (linguistic entities) and declarative utterances (concrete events) *express* propositions, but propositions exist regardless of whether they are ever expressed; similarly, imperative sentences and imperative utterances *express* prescriptions,[16] but prescriptions exist regardless of whether they are ever expressed. ('Statement' is ambiguous between 'declarative utterance' and 'proposition'; similarly, 'command' and 'order' are ambiguous between 'imperative utterance' and 'prescription'. I disambiguate by choosing the latter senses.) Propositions *track* facts if they are *true* (every proposition "purports" to be true), but don't *generate* facts (declarative utterances do so); similarly, prescriptions *track* reasons for action if they are *binding* (every prescription "purports" to be binding), but don't *generate* reasons (imperative utterances may do so, if the utterers are qualified authorities). So prescriptions are not true or false: they are rather binding or non-binding, depending on whether they have or lack normative force, they do or do not track reasons for action. (If there are distinct kinds of reasons—e.g., moral, prudential, legal—then there are corresponding kinds of "bindingness".) I reserve the terms 'norm' and 'duty' for binding prescriptions.[17]

A prescription is always indexed to a node[18] and to an agent (the *prescriptee*);[19] I omit reference to agents, and also sometimes to nodes. A prescription at node n (an *n-prescription*), C_n, has a *satisfaction set*, S_{Cn}, and a *violation set*, V_{Cn}; these are nonoverlapping subsets of U_n (the universal n-action). An n-history *satisfies* C_n exactly if it is a member of S_{Cn}, *violates* C_n exactly if it is a member of V_{Cn}, and neither satisfies nor violates C_n exactly if it is a member of neither; similarly

---

[16] I am not denying that declarative sentences (e.g., "you may not park here" or even "you will open the door") can be *used* imperatively (Alchourrón & Bulygin 1981: 97).

[17] Two remarks for the cognoscenti. (1) Is my conception of prescriptions *expressive* or *hyletic*? (See Alchourrón & Bulygin 1981.) Neither. Against the expressive conception (as I understand it), I hold that prescriptions exist even if they are never expressed; against the hyletic conception, I hold that prescriptions are *not* the result of an operation on propositions. (2) What von Wright (1969: 89) calls *idealist* and *realist* "traditions in the ontology of norms" are better understood as views on the *source* of bindingness (or *normativity*). I am not taking a stand on this issue.

[18] A *guideline*, by contrast, is indexed to a *set* of nodes: it is a function from nodes in that set to prescriptions indexed to those nodes. (Cf. the distinction between actions and strategies in footnote 8.) I take imperative sentences like "*never* kill" to express guidelines rather than prescriptions. I don't examine guidelines in this paper.

[19] Or to a group of agents. I think that talk of "impersonal" prescriptions, which have no prescriptee(s) and correspond to ought to be, is a misnomer; I won't say more about this in the paper.

for *n*-actions (rather than *n*-histories), by replacing 'member' with 'subset'. These definitions validate the plausible principle that an action satisfies a prescription exactly if *every* history in the action satisfies the prescription. Now some distinctions.

- A prescription is *unconditional* exactly if its violation set is the complement of its satisfaction set, and is *conditional* (on the union of its satisfaction and violation sets) otherwise. (More on this in §5.)
- A prescription is *vacuous* exactly if its violation set is empty, and is *non-vacuous* otherwise. "Do whatever you like", for example, (is an imperative sentence that) expresses a vacuous prescription.
- A prescription is *impossible* exactly if its satisfaction set is empty, and is *possible* otherwise. "Raise and don't raise your right arm", for example, expresses an impossible prescription.
- A prescription is *infeasible* exactly if its satisfaction set includes no action, and is *feasible* otherwise. Clearly, every impossible prescription is infeasible. Not the other way round, however. "Raise your right arm at an odd angle (i.e., 1^{o}, 3^{o}, 5^{o}, …)", for example, expresses a possible but infeasible prescription (satisfied by many *histories* but by no *action*). "Toss this fair coin so that it comes up heads" gives another example. Some infeasible prescriptions correspond to cases in which a prescriptor—i.e., someone who expresses a prescription—is mistaken about the prescriptee's abilities or does not abide by the "ought implies can" principle (which can be rigorously formulated as the principle that *no infeasible prescription is binding*).
- A prescription is *normal* exactly if its satisfaction set is an action, and is *non-normal* otherwise. Clearly, every infeasible (and thus every impossible) prescription is non-normal. Not the other way round, however. "Raise your right arm at an odd angle or don't raise it at all", for example, expresses a feasible but non-normal prescription. "Don't get killed while driving" gives another example: the prescription it expresses is feasible because it is within my power to avoid driving at all, but is non-normal because, given that I drive, not being killed is not within my power and is thus not an action (so its union with the action of not driving at all is not an action either). Some normal prescriptions correspond to cases in which a prescriptor specifies what the prescriptee is to *do*, whereas some non-normal prescriptions correspond to cases in which a prescriptor specifies what the prescriptee is to *achieve* (cf. von Wright's (1981a: 10-1) process/achievement distinction).

I haven't yet said what a prescription *is*; more on this in §6. For the moment just note that *every* prescription *has* a satisfaction and a violation set, but by

definition a *thin* *n*-prescription just *is* an ordered pair $<S_{Cn}, V_{Cn}>$ of nonoverlapping subsets of U_n.

4. The logic of prescriptive propositions

4.1. Prescriptive propositions and the structure of deontic logic

At last we get to deontic logic proper. A *prescriptive* (i.e., obligation, permission, or prohibition) proposition, in the context of single-node/single-agent deontic logic (which is the target of this paper), is a proposition stating that a given *n*-action is *obligatory* (required), *permitted* (allowed), or *forbidden* (prohibited) *according* to a given *n*-prescription. (A *deontic* proposition is a truth-functional combination of prescriptive propositions.) It can hardly be overemphasized that a prescriptive proposition is always indexed to a given prescription: the claim that an *n*-action is (e.g.) obligatory *simpliciter* is shorthand for the claim that the action is obligatory *according* to a very special prescription, the *all-things-considered n*-prescription (i.e., the combination of *all binding n*-prescriptions). Moreover, a prescriptive proposition is always indexed to a *specific* prescription, not to a *system* of prescriptions. Although several authors distinguish "norms" from "normative propositions", they typically index the latter to "normative systems": they understand the claim that an action is (e.g.) obligatory according to a given normative system as the claim that *there is* in the system a prescription according to which the action is obligatory.[20] But what if there is *also* in the system a prescription according to which the action is forbidden? Of course I also need to address the issue of conflicting prescriptions, but I do so (in another paper) by finding ways to combine prescriptions into a *single* prescription, so I don't need to index prescriptive propositions to *systems* of prescriptions.

Those who focus on *systems* of prescriptions typically understand (e.g.) "$\sim(Op \,\&\, P\sim p)$" as the claim that "there are no two norms in [a given system] such that the one norm pronounces a state obligatory and the other its contradictory permitted" (von Wright 2000: 176). A focus on *single* prescriptions motivates a different understanding: no prescription is such that according to it an action is obligatory but the complement of (i.e., refraining from) the action is permitted. A failure to clearly distinguish these two understandings is, I submit, responsible for part of the confusion that plagues deontic logic. The former understanding implicitly promotes an impoverished conception of prescriptions, according to which a given prescription is associated with a *single* action. In a

[20] Alchourrón 1969: 245, 1993: 44; Bulygin 2000: 133; von Wright 1963: 106, 1969: 102, 1981a: 26, 1981b: 405, 2000: 173. Cf. Belzer & Loewer 1994: 400; Castañeda 1981; Hansson 1969/1971: 123.

sense this is not false: a (normal, unconditional, thin) prescription is fully determined by a single action, namely its satisfaction set. But this should not obscure the fact that in another sense a given prescription is associated with a multitude of actions: those which are obligatory, permitted, or forbidden according to the prescription. A single-prescription understanding of deontic propositions makes this richer conception of prescriptions explicit.

The above considerations suggest an answer to a question that has been "tormenting" von Wright (1999a: 20, 1999b: 32) for decades: is deontic logic a logic of prescriptions or a logic of prescriptive propositions? The answer is that deontic logic has two parts. The first part is a logic of prescriptive (and deontic) propositions indexed to a *single* prescription. It addresses three main questions: which *n*-actions are (1) obligatory, (2) permitted, and (3) forbidden according to a given *n*-prescription? (It will be seen that the answers do take into account the structure of prescriptions in terms of satisfaction and violation sets.) The second part is a logic of prescriptions. It addresses mainly the question: how to *combine* the prescriptions in a set into a single prescription? (It turns out that the answer depends on whether the prescriptions in the set are binding.) To find what is (e.g.) obligatory according to a system of prescriptions, one needs to follow a two-step procedure: apply first the logic of prescriptions to find the combined prescription, and apply next the logic of prescriptive propositions to find what is obligatory according to the combined prescription. So the two parts of deontic logic are intimately related.

In this section I deal with the logic of prescriptive propositions indexed to an *unconditional* prescription (so the violation set is the complement of the satisfaction set). (Conditional prescriptions I address in §5; the logic of prescriptions I address in another paper.) It will be seen that a lot of things fall into place once we put together the three main ideas introduced so far: that actions are normally coarse-grained, that prescriptions have satisfaction and violation sets, and that prescriptive propositions are indexed to single prescriptions.

4.2. Weak and strong permission and prohibition propositions

Start with what should be an easy case: if an action is included in the satisfaction set of a given prescription, then the action is clearly permitted according to the prescription. For example: the action of closing the door with my left hand is included in the action of closing the door and is thus clearly permitted according to a prescription whose satisfaction set is the latter action. The standard objection is that not *every* way of closing the door will do: I am

surely not permitted to close the door so violently that it breaks your nose.[21] I reply that the objection hinges on a subtle shift from an indexed to an unindexed permission proposition. As far as the *specific* prescription under consideration is concerned, it doesn't matter whether I break your nose or not: all that matters is whether I close the door or not. So the *indexed* permission statement *is* true: *according* to that prescription, the violent action is permitted. Of course it does not follow that the violent action is permitted *simpliciter*: presumably some *other* (binding) prescription forbids it, and its combination with the prescription under consideration also does. In response one might argue that, when you utter "close the door", you do not thereby grant me permission to break your nose; so the violent action is not permitted *even* according to the specific prescription expressed by your utterance. But then, I reply, the satisfaction set of that prescription is *not* the action of closing the door; when you utter "close the door" you express instead a prescription with a more restricted satisfaction set, something like the action of closing the door not too violently. How exactly to specify that satisfaction set is a vexed and widely discussed problem;[22] but it's a problem in the pragmatics of deontic speech, not in deontic logic proper. The important point is that we have no counterexample to the principle that any action included in the satisfaction set of a given prescription is clearly permitted according to the prescription. Similarly for prohibition propositions: any action included in the violation set of a given prescription is clearly forbidden according to the prescription.[23]

Having thus disposed of the "easy" case, let us turn to the hard one. Consider a *mixed* action: a (coarse-grained) action which has both a clearly permitted subaction (included in the satisfaction set) and a clearly forbidden one (included in the violation set). Is such an action permitted, forbidden, both, or neither? I answer that, although we don't want to say the action is forbidden, sometimes we want to say it is permitted and sometimes we want to say it is not. We vacillate. Such vacillation, I submit, is responsible for part of the confusion that plagues deontic logic.

---

[21] Cf. Hilpinen 1981: 322; Kamp 1974: 62-3. Equivalently: I am not permitted to close the door *and* (e.g.) break your nose—or so the objection goes.

[22] More precisely, it's a related problem about *permissive* (rather than *imperative*) utterances that has been widely discussed. Cf. Kamp 1974, 1979; Lewis 1979a/1983: 234-6, 1979b/2000b; Nute 1985: 183-7; Rohrbaugh 1997.

[23] Doesn't the "Penitent's" (or "Robber's") paradox follow, namely that "[i]f one is forbidden to commit a crime, one is forbidden to commit a crime and do penitence for it"? (Meyer & Wieringa 1993: 6; cf. al-Hibri 1978: 24; Nowell-Smith & Lemmon 1960: 294.) No: the actions of committing a crime and of doing penitence are indexed to *different* nodes.

To feel the force of the vacillation, consider an example. Suppose I ask my dentist to pull out a certain tooth—but *gently*. It seems then that the prescription I expressed permits the mixed action of pulling out the tooth (which includes both the clearly permitted subaction of pulling out the tooth gently and the clearly forbidden subaction of pulling out the tooth non-gently). Indeed, it would sound absurd to say to the dentist: "What I said does not allow you to pull out the tooth; it only allows you to pull out the tooth *gently*." The dentist could properly reply: "How can I pull out the tooth gently without pulling out the tooth? What you said does not allow me to pull out the tooth *non*-gently; it does allow me to pull out the tooth." Here is then a case where a mixed action seems permitted. Suppose, however, that we (equivalently) redescribe the mixed action as the action of pulling out the tooth gently *or* non-gently. This is analogous to the case in which I ask you to post a certain letter and we wonder whether what I said allows you to post *or burn* the letter (Ross 1941: 62). The answer seems clearly negative. To continue a modified version of the above conversation, after the dentist says "How can I pull out the tooth gently without pulling it out gently *or* non-gently?", I could respond: "You are not *forbidden* to pull out the tooth gently or non-gently; this would imply that you are forbidden to pull out the tooth gently, and you are not. But you are not *permitted* either to pull out the tooth gently or non-gently; this would imply that you are permitted to pull out the tooth *non*-gently, and you are not." It seems then natural to say that the mixed action is not permitted after all. Do you now feel the force of the vacillation?

We are thus naturally led to von Wright's distinction between weak and strong permission propositions: an action is *weakly* permitted according to a given prescription exactly if it *overlaps* (is *partly* in) the satisfaction set of the prescription, and is *strongly* permitted exactly if it is *included* in (is *wholly* in) the satisfaction set.[24] So pulling out the tooth (like posting or burning the letter) is weakly but not strongly permitted. A similar distinction can be made for prohibition propositions: an action is *weakly* forbidden according to a given prescription exactly if it *overlaps* the violation set of the prescription, and is *strongly* forbidden exactly if it is *included* in the violation set. So pulling out the tooth (like posting or burning the letter) is weakly but not strongly forbidden. It follows that a *mixed* action is both weakly permitted and weakly forbidden, and neither strongly permitted nor strongly forbidden. Note that the proposition that an action is weakly permitted is just the negation of the proposition that the action is strongly forbidden. (Indeed, given that for unconditional prescriptions the violation set is the complement of the satisfaction set, the proposition that an

[24] Cf. von Wright 1968: 26, 1969: 94-5, 1971: 160, 166, 172-3, 1981a: 25, 32, 1981b: 414; Åqvist 1984: 638, 1987: 54; Hilpinen 1981: 320; Nute 1985: 173-83.

action overlaps the satisfaction set is just the negation of the proposition that the action is included in the violation set.) Similarly, the proposition that an action is strongly permitted is the negation of the proposition that the action is weakly forbidden.

The vacillation can be now explained as follows. There are four possible positions on the status of a mixed action: two symmetric and two asymmetric ones. A first symmetric position is that the action is both permitted and forbidden; i.e., both weakly permitted and weakly forbidden. Apparently no one wants to say this. The other symmetric position is that the action is neither permitted nor forbidden; i.e., neither strongly permitted nor strongly forbidden. This position is supported by two considerations. First, it is the only plausible symmetric position, and we do want a symmetric position: it seems hard to justify an asymmetry given that the description of the situation is symmetric (a mixed action overlaps both the satisfaction and the violation set). Second, just as it is plausible to claim that any subaction of a forbidden action is itself forbidden, it is also plausible to claim that any subaction of a permitted action is itself permitted. The latter claim corresponds to the widely discussed "principle of free choice permission".[25] (This is stronger than the claim, defended four paragraphs ago, that any action included in the satisfaction set is permitted; but it can be defended against objections in a similar way.) Now a first asymmetric position is that the action is permitted but not forbidden; i.e., weakly permitted but not strongly forbidden. This is the position of "standard deontic logic" (e.g., Nute & Yu 1997: 3). It is supported by a reluctance to say that the action is forbidden (the only plausible symmetric position agrees on this) in conjunction with a tendency to say that whatever is not forbidden is permitted (as we saw, strong prohibition is just the negation of weak permission). Indeed, '*im*permissible' is often used synonymously with 'forbidden'. Finally, the other asymmetric position is that the action is forbidden but not permitted; i.e., weakly forbidden but not strongly permitted. Apparently no one wants to say this, because for some as yet mysterious reason everyone agrees that the action is not forbidden. (I propose an explanation for this agreement later on.) So we are left with two plausible positions, a symmetric and an asymmetric one. They both agree that the action is not forbidden, so they use strong rather than weak prohibition. The symmetric position is motivated by symmetry and by the principle of free choice permission; the asymmetric position is motivated by the complementarity between permission and prohibition. So the conflict between (1) symmetry plus free choice permission and (2) complementarity seems to

[25] Cf. von Wright 1968: 22, 1981a: 7-8, 1981b: 417, 1999a: 17, 1999b: 29; Åqvist 1987: 47-54; Føllesdal & Hilpinen 1971: 22-3; Hilpinen 1981: 320; Kamp 1974; Makinson 1984; Meyer & Wieringa 1993: 6; Nute 1985: 183-7; Stenius 1982.

explain the vacillation. Whatever its explanation, the vacillation is here to stay. The important point is that, by talking about coarse-grained actions in relation to satisfaction and violation sets, the situation has been perfectly clarified. Whether we call a mixed action neither permitted nor forbidden or permitted but not forbidden seems now inconsequential, a mere verbal choice. What matters is that such an action overlaps both the satisfaction and the violation set.

4.3. Weak and strong obligation propositions

Which actions are obligatory according to a given prescription? If the prescription is normal (i.e., its satisfaction set is an action), then one clearly obligatory action is the satisfaction set itself. For example: if I ask my dentist to pull out a certain tooth gently, then according to the prescription I expressed the dentist is clearly required to pull out the tooth gently. Is that all? In one sense yes: what is obligatory is to choose one of the strongly permitted actions, so the one and only obligatory action is the union of all strongly permitted actions—namely the satisfaction set (or, if the prescription is not normal, the maximal action included in the satisfaction set). But in another sense there are many more obligatory actions. After all, if the dentist is required to pull out my tooth gently, then the dentist is required to pull out my tooth. (This inference corresponds to the widely discussed "principle of obligation execution" or "deontic inheritance rule".[26]) But how can the action of pulling out my tooth be obligatory when it includes the strongly forbidden subaction of pulling out my tooth *non*-gently? Here is then another vacillation.

We are thus naturally led to Brown's (1996: 50, 1999: 114, 2000: 85-6) distinction between two kinds of obligation propositions, which he calls 'type 1' and 'type 2' but I call 'weak' and 'strong': an action is *strongly* obligatory according to a given prescription exactly if it *coincides* with (the maximal action included in) the satisfaction set of the prescription, and is *weakly* obligatory exactly if it *includes* (is a *superset* of) the satisfaction set.[27] So pulling out the

---

[26] This is the principle that, if one is required to do A and one cannot do A without doing B (i.e., $A \subseteq B$), then one is required to do B; in other words, every action which is a *superset* of an obligatory action is itself obligatory. For endorsements of or arguments supporting the principle see: al-Hibri 1978: 16-7; Asher & Bonevac 1997: 163; Belzer & Loewer 1997: 51; Brink 1994: 234-5; Knuutila 1981: 240; Nute 1997: 290; Nute && Yu 1997: 5-6; Prakken & Sergot 1997: 224; Sinnott-Armstrong 1985: 164. For rejections of or arguments against the principle see: Belzer & Loewer 1997: 51; Brink 1994: 234-5; Forrester 1984: 197; Knuutila 1981: 241; Loewer & Belzer 1986: 125-6.

[27] A similar distinction is that between "primary" and "derived" obligation propositions (Belzer & Loewer 1997: 51; cf. 1993: 65-6; Loewer & Belzer 1986: 125).

tooth (like posting or burning the letter) is weakly but not strongly obligatory.[28] Note that an action is *strongly* obligatory exactly if it is both a superset and a subset of (the maximal action included in) the satisfaction set; in other words, exactly if it is both *weakly* obligatory and *strongly* permitted. So strong obligation entails strong permission. (Similarly, weak obligation entails weak permission if the prescription is possible.) Note also that a mixed action can be both weakly obligatory and weakly forbidden; given that apparently no one wants to say that an action is both obligatory and forbidden, the attractiveness of weak obligation may explain the previously noted agreement that a mixed action is not forbidden (and thus the unattractiveness of weak prohibition).

The vacillation between weak and strong obligation can be now explained as follows. Weak obligation is supported by the seemingly uncontroversial claim that (the maximal action included in) the satisfaction set is obligatory, in conjunction with the principle of obligation execution. Strong obligation is supported by the very plausible claim that every obligatory action is permitted (i.e., "ought implies may"),[29] in conjunction with the principle of free choice permission (which motivates strong permission). So the conflict between (1) obligation execution and (2) ought implies may plus free choice permission seems to explain the vacillation. In any case, the important point is again that, by talking about coarse-grained actions in relation to satisfaction and violation sets, the situation has been perfectly clarified. Whether we call a mixed action that includes the satisfaction set obligatory or not seems now inconsequential, a mere verbal choice. What matters is that we know what we are required to do: choose any action included in the satisfaction set.

---

[28] One could make a similar distinction for prohibition propositions, but such a distinction would not be very useful. (1) On the one hand, it is implausible to say that there is exactly *one* forbidden action, namely (the maximal action included in) the violation set: apparently everyone agrees that *every* action included in the violation set is forbidden. So no interesting prohibition concept corresponds to *strong* obligation. (2) On the other hand, the principle of obligation execution (footnote 26) has the following analog for prohibition: if one is forbidden to do A and one cannot do B without doing A (i.e., $B \subseteq A$), then one is forbidden to do B. In other words, every action which is a *subset* of a forbidden action is itself forbidden; this leads us back to *strong* prohibition, so no *new* prohibition concept corresponds to *weak* obligation.

[29] On the "ought implies may" principle see: Anderson 1969: 109; Åqvist 1984: 616, 1987: 19; Moritz 1963; Stenius 1963: 254; Nute 1997: 312. Some people deny this principle because they think it has the consequence that an action and its complement cannot be both obligatory, and they take this consequence to conflict with the possibility of moral dilemmas. But the consequence in question is that an action and its complement cannot be both obligatory according to the *same* prescription, whereas in moral dilemmas an action and its complement are obligatory according to *different* prescriptions. This is an example of the importance of distinguishing single-prescription from multiple-prescription readings of deontic propositions.

4.4. A taxonomy of deontic logic systems

Let A be an n-action and C an n-prescription (with satisfaction set S_C and violation set V_C). The main results of §4.2 and §4.3 can be summarized in symbols as follows.

- Weak obligation (A is weakly C-obligatory): $O_{-C}(A) \equiv A \supseteq S_C$.
- Strong obligation (A is strongly C-obligatory): $O_{+C}(A) \equiv A = S_C$.[30]
- Weak permission (A is weakly C-permitted): $P_{-C}(A) \equiv A \cap S_C \neq \varnothing$.
- Strong permission (A is strongly C-permitted): $P_{+C}(A) \equiv A \subseteq S_C$.
- Weak prohibition (A is weakly C-forbidden): $F_{-C}(A) \equiv A \cap V_C \neq \varnothing$.
- Strong prohibition (A is strongly C-forbidden): $F_{+C}(A) \equiv A \subseteq V_C$.

In theory, eight possible systems of deontic logic (by which in this subsection I mean "logic of single-prescription deontic propositions) correspond to the eight possible ways of choosing a triple consisting of an obligation, a permission, and a prohibition concept. The "standard" system of deontic logic corresponds to the triple <weak obligation, weak permission, strong prohibition> (though note that the standard system is for ought to be rather than ought to do). This is most easily seen by comparing the above definitions with Anderson's (1956/1967, 1958, 1959, 1969) reductive definitions, his "sanction" S corresponding to the violation set and the negation Q of the sanction corresponding to the satisfaction set. According to Anderson, A is obligatory exactly if $\Box(Q \rightarrow A)$, is permitted exactly if $\Diamond(Q \wedge A)$, and is forbidden exactly if $\Box(A \rightarrow S)$; these are isomorphic to $S_C \subseteq A$, $S_C \cap A \neq \varnothing$, and $A \subseteq V_C$ respectively. My definitions, like Anderson's, have the advantage of dispensing with the need to choose a set of *axioms*: once a triple of *concepts* is chosen, every derivable proposition of the corresponding system of deontic logic follows from the *definitions* of the concepts. As an example consider a version of the principle of obligation execution: if A is weakly C-obligatory and one cannot do A without doing B, then B is weakly C-obligatory. This is valid because obviously $S_C \subseteq A$, in conjunction with $A \subseteq B$, entails $S_C \subseteq B$.

What about the remaining seven systems of deontic logic? As we already saw, not all of them are equally interesting. Four of them are eliminated if weak prohibition is considered uninteresting. Two more (namely those that combine weak obligation with strong permission) are eliminated if ought implies may. So only two of the seven systems stand out: <strong obligation, strong permission, strong prohibition>, and <strong obligation, weak permission, strong prohibition>. This is not to say that the other five systems are useless; their

[30] This is if C is normal. For non-normal prescriptions, replace throughout S_C with the maximal action (i.e., the union of all actions) included in S_C.

usefulness is a question for future research. (For example, a system proposed by von Wright (1981b: 418-20) corresponds to <strong obligation, strong permission, weak prohibition>.) The important point for now is the existence of the taxonomy.[31] Note finally that some deontic propositions use (e.g.) *both* weak and strong obligation. For example, the following version of obligation execution is easily seen to be valid: if *A* is *strongly C*-obligatory and one cannot do *A* without doing *B*, then *B* is *weakly C*-obligatory. So *mixed* systems of deontic logic may also be useful.

5. Conditional and derived prescriptions

A *conditional* prescription, as I said in §3, is a prescription whose satisfaction and violation sets are (nonoverlapping but) not complementary: they do not exhaust the universal action. (In this section and the next for simplicity I drop reference to nodes, but everything should be understood as indexed to a given node *n*.) Call the union of the satisfaction and violation sets of a given prescription the *context* (cf. Prakken & Sergot 1997: 236) of the prescription. If the prescription is conditional, then call its context its *condition*. (So an unconditional prescription has a context—namely the universal action—but has no condition.) The three questions of the logic of prescriptive propositions arise again: which actions are (1) obligatory, (2) permitted, and (3) forbidden according to a given conditional prescription? The complication now is that some actions lie partly or wholly in the complement of the condition. To appreciate the issues, take an example.

Suppose I tell my dentist: "*if* you pull out the tooth, do it gently". I expressed a conditional prescription whose satisfaction set is the (dentist's) action of pulling out the tooth gently, whose violation set is the action of pulling out the tooth non-gently, and whose condition is the action of pulling out the tooth. Consider three kinds of actions.

- Actions wholly in the complement of the condition. For example, the action of *cleaning* (and thus not pulling out) the tooth. It seems natural to say that such actions are neither obligatory nor permitted nor forbidden according to the prescription: they are "outside" the context of the prescription.
- Coarse-grained actions which, in addition to overlapping both the satisfaction and the violation set, overlap the complement of the condition. For example, the action of using a certain instrument, *I*, which can be used either for pulling out—gently or non-gently—the tooth or for cleaning it. It seems that the discussion of §4.2 applies here as well: such actions are

[31] Compare Kripke's (1963) taxonomy of modal logic systems in terms of the properties of the accessibility relation between possible worlds.

weakly but not strongly permitted and weakly but not strongly forbidden, and this follows indeed from the definitions in §4.4.

- Coarse-grained actions which overlap the satisfaction set and the complement of the condition but not the violation set. For example, the action of using instrument *I* gently (and thus either pulling out the tooth gently or cleaning it gently). According to the definitions in §4.4, such actions are not strongly permitted: they are not *wholly* in the satisfaction set. It seems, however, that they *are* clearly permitted: *given* that the dentist pulls out the tooth, she is permitted to use instrument *I* gently. The fact that the action of using instrument *I* gently is not wholly in the satisfaction set (because it includes the subaction of cleaning—and thus not pulling out—the tooth gently) is irrelevant: what matters is that the *intersection* of this action with the condition *is* wholly in the satisfaction set. So we need to modify the definition of strong permission: an action is strongly permitted according to a conditional prescription exactly if the *intersection* of the action with the condition of the prescription is included in the satisfaction set of the prescription (provided the intersection is non-empty). It can be seen that such a simple modification gives the desired results in all cases.

To summarize in symbols, let B be the context of a prescription C: $B = S_C \cup V_C$. In the definitions of §4.4, replace A with $A \cap B$. For explicitness I will write (e.g.) $O_{-C}(A|B)$ whenever C is a prescription conditional on $B = S_C \cup V_C \neq U$. (The $O_{-C}(A|B)$ notation does not exclude the case in which C is unconditional, namely $B = U$. So I will use $O_{-C}(A|B)$ *if*, not *exactly if*, C is conditional. I will reserve $O_{-C}(A)$ for unconditional prescriptions.) So we have, for a conditional prescription C (if $A \cap B \neq \varnothing$, otherwise all propositions are false):

- Weak conditional obligation (A is weakly C-obligatory given B):
  $O_{-C}(A|B) \equiv A \cap B \supseteq S_C$.
- Strong conditional obligation (A is strongly C-obligatory given B):
  $O_{+C}(A|B) \equiv A \cap B = S_C$.
- Weak conditional permission (A is weakly C-permitted given B):
  $P_{-C}(A|B) \equiv A \cap B \cap S_C \neq \varnothing$.
- Strong conditional permission (A is strongly C-permitted given B):
  $P_{+C}(A|B) \equiv A \cap B \subseteq S_C$.
- Weak conditional prohibition (A is weakly C-forbidden given B):
  $F_{-C}(A|B) \equiv A \cap B \cap V_C \neq \varnothing$.
- Strong conditional prohibition (A is strongly C-forbidden given B):
  $F_{+C}(A|B) \equiv A \cap B \subseteq V_C$.

The above definitions, however, can be considerably simplified by taking advantage of the fact that B is $S_C \cup V_C$. With some work it can be seen that the definitions are equivalent to the following:

- Weak (conditional) obligation: $O_{-C}(A|B) \equiv A \supseteq S_C$.
- Strong (conditional) obligation: $O_{+C}(A|B) \equiv A \supseteq S_C \wedge A \cap V_C = \varnothing$.
- Weak (conditional) permission: $P_{-C}(A|B) \equiv A \cap S_C \neq \varnothing$.
- Strong (conditional) permission: $P_{+C}(A|B) \equiv A \cap V_C = \varnothing$.
- Weak (conditional) prohibition: $F_{-C}(A|B) \equiv A \cap V_C \neq \varnothing$.
- Strong (conditional) prohibition: $F_{+C}(A|B) \equiv A \cap S_C = \varnothing$.

Note that in the simplified *definientia* (right-hand sides of definitions) B does not appear. Moreover, if C is unconditional, then the simplified *definientia* are equivalent to those in §4.4: in three cases they are identical, and in the other three cases the equivalence follows from the fact that, for unconditional prescriptions, V_C is the complement of S_C (so, for example, for strong permission $A \subseteq S_C$ is equivalent to $A \cap V_C = \varnothing$).[32] We have thus reached a unified treatment of conditional and unconditional prescriptive propositions: in the above, final definitions C may be conditional or unconditional. As Brown (2000: 93) notes, some authors take unconditional prescriptive propositions as primitive and define conditional ones in terms of them, whereas (more typically) other authors take conditional prescriptive propositions as primitive and define unconditional ones as limiting cases. In my theory there is no need to choose.

A limitation of the above definitions might be taken to be that, at least when the satisfaction and violation sets are actions, the condition must be an action. Aren't we also interested in prescriptions like the one expressed by "if it rains, cancel the picnic", where the antecedent corresponds to a proposition rather than an action? (See Belzer & Loewer 1994: 405-6.) We are, but such a prescription is not conditional. The function of the antecedent is rather to *attach* an *unconditional* prescription ("cancel the picnic") to a given node; namely, to a node at which it rains (or to all such nodes in a set of nodes). So I don't think we need conditional prescriptions whose conditions are not actions.

The above definitions can be seen to satisfy, e.g.: $P_{-C}(A|B) \equiv P_{-C}(A \cap B|B)$. Note an analogy with conditional probability: $P(A|B) = P(A \cap B|B)$. This analogy is also useful in another respect: similarly to the way in which from a probability measure P' we can derive a new probability measure P by conditioning on B

[32] Since the two expressions are not equivalent for conditional prescriptions, strictly speaking one could distinguish between *strong* permission, amounting to $A \cap V_C = \varnothing$, and *strongest* permission, amounting to $A \subseteq S_C$; similarly for prohibition and obligation. But I argued above (with the dentist example) that strongest permission is not useful.

($P(A) = P'(A \cap B)/P'(B)$), from a prescription C' we can derive a new prescription C by conditioning on an action B: $S_C = S_{C'} \cap B$, $V_C = V_{C'} \cap B$. For explicitness, when C is derived from C' by conditioning on B, I may write $O_{C'|B}(A|B)$. (This is equivalent to $A \supseteq S_{C'} \cap B$, whereas $O_{C'}(A|B)$ is equivalent to $A \supseteq S_{C'}$.)

Conditional probability runs into trouble when the probability of the condition is zero. Similarly, derived prescriptions run into trouble when the condition is strongly forbidden. Indeed, in such a case $B \subseteq V_{C'}$, so $B \cap S_{C'} = \varnothing$, so $S_C = \varnothing$: the derived prescription is impossible (§3). How, then, can "contrary-to-duty" prescriptions be possible? I address this question in the next section.

6. Thick prescriptions and subprescriptions

Recall (from §3) that, by definition, a *thin* prescription is fully characterized by a satisfaction and a violation set. For some prescriptions, however, such a characterization is inadequate. Suppose there are three buttons in front of you: a red, a blue, and a green one. I tell you: "don't push any button, *especially* the red one". The prescription I expressed is violated if you push the blue button, but is *more strongly* violated if you push the red one. Now suppose I tell you: "push any button, *preferably* the red one". The prescription I expressed is satisfied if you push the blue button, but is *more strongly* (i.e., *better*) satisfied if you push the red one. Less artificial examples are also easy to find. I can tell you a secret and ask you: "don't tell anyone, *especially* my mother". Or I can tell you: "open a window, *preferably* the left one". These examples suggest that some prescriptions are characterized by *intensities*, namely *degrees* to which they are satisfied or violated by various actions. Define then (to a first approximation) a *thick* prescription C as an ordered pair of (1) an *intensity function*, I_C, which assigns a number to every action, and (2) a *primary threshold*, T_C, namely a number such that an action A (i) satisfies C exactly if its intensity, $I_C(A)$, exceeds the threshold and (ii) violates C exactly if its intensity is below the threshold (and thus (iii) neither satisfies nor violates C exactly if its intensity is equal to the threshold). The satisfaction set of C is then the union of all actions that satisfy C, and the violation set of C is the union of all actions that violate C.[33]

We need to impose some conditions on the intensity function, however, to ensure that every action included in the satisfaction set has intensity above the threshold (and similarly for the violation set). It turns out that the following two conditions will do the trick. For any actions A_1 and A_2:

[33] This definition has the consequence that C is normal. I omit the more general definition.

(1) If $I_C(A_1) > T_C$ and $I_C(A_2) > T_C$, then $I_C(A_1 \cup A_2) > T_C$.
(2) If $I_C(A_1) < T_C$ and $I_C(A_2) < T_C$, then $I_C(A_1 \cup A_2) < T_C$.

It can be shown that (1) and (2) entail:

(3) If $I_C(A_1) > T_C$ and $A_2 \subseteq A_1$, then $I_C(A_2) > T_C$.
(4) If $I_C(A_1) < T_C$ and $A_2 \subseteq A_1$, then $I_C(A_2) < T_C$.
(5) If $I_C(A_1) > T_C$ and $I_C(A_2) < T_C$, then $I_C(A_1 \cup A_2) = T_C$.
(6) If $I_C(A_1) = T_C$, then $I_C(A_1 \cup A_2) = T_C$.

(3) guarantees that every subaction of an action that satisfies C itself satisfies C, and similarly for (4). (5) says that a mixed action which consists of a subaction that satisfies C and a subaction that violates C neither satisfies nor violates C.

Although the above characterization of thick prescriptions is quite general, for some prescriptions it is still inadequate. Suppose I tell you: "don't push any button, but if nevertheless you push a button then push the blue or the green one". The prescription I expressed is violated if you push the blue button, and is more strongly violated if you push the red one; so far the case is analogous to that of a previously examined prescription, namely the one expressed by "don't push any button, especially the red one". But in contrast to the previous prescription, as far as the new prescription is concerned you do satisfy something if you push the blue button: *embedded* in the new prescription is a *subprescription*, namely the conditional prescription expressed by "if you push a button, then push the blue or the green one". This subprescription is a "contrary-to-duty" prescription. But talk of "contrary-to-*duty*" prescriptions, namely what can be called *negative* subprescriptions, obscures the fact that there are *positive* subprescriptions as well, whose context is a subset of the *satisfaction* (rather than the violation) set of the prescription in which they are embedded. Suppose I tell you: "push any button you like, but if you don't push the green button then push the red one". The prescription I expressed is satisfied if you push the blue button, and is more strongly violated if you push the red one; so far the case is analogous to that of the prescription expressed by "push any button, preferably the red one". But in contrast to the previous prescription, as far as the new prescription is concerned you do violate something if you push the blue button: embedded in the new prescription is a positive subprescription, namely the conditional prescription expressed by "if you don't push the green button, then push the red one". This subprescription is *not* a "contrary-to-duty" prescription.

Fortunately, a simple modification of the above definition of thick prescriptions suffices to capture the idea of embedded subprescriptions. It suffices to understand T_C (not as a number, but rather) as a *threshold function* which

assigns a number to those actions to which embedded subprescriptions correspond. The idea is simple: if B is an action (in the domain of T_C) that violates C, then every action included in B has intensity lower than $T_C(U)$, the primary threshold of C; but some of these actions may have intensities higher than $T_C(B)$, the threshold of the B-subprescription, and the union of these actions is the satisfaction set of the B-subprescription. Note that the B-subprescription can itself have embedded subprescriptions if some action included in B is in the domain of T_C. More generally, for any action B in the domain of T_C, the B-subprescription of C is simply the ordered pair of the *restrictions* to subactions of B of the intensity and threshold functions of C.

7. Conclusion

This paper is only a preliminary report of a small part of a larger project. Unfortunately I did not have the time to talk about conflicts of prescriptions or to explain how my proposal illuminates the standard paradoxes of deontic logic. But I hope that what I said makes a case for the fruitfulness of my approach.

REFERENCES

Alchourrón, Carlos E. (1969). Logic of norms and logic of normative propositions. *Logique et Analyse* 12: 242-268.

Alchourrón, Carlos E. (1993). Philosophical foundations of deontic logic and the logic of defeasible conditionals. In J.-J. Ch. Meyer & R. J. Wieringa (Eds.), *Deontic logic in computer science: Normative system specification*, 43-84. New York: Wiley.

Alchourrón, Carlos E., & Eugenio Bulygin. (1981). The expressive conception of norms. In R. Hilpinen (Ed.), *New studies in deontic logic*, 95-124. Dordrecht: Reidel.

al-Hibri, Azizah. (1978). *Deontic logic: A comprehensive appraisal and a new proposal.* Washington, D.C.: University Press of America.

Anderson, Alan Ross. (1958). A reduction of deontic logic to alethic modal logic. *Mind* 67: 100-103.

Anderson, Alan Ross. (1959). On the logic of commitment. *Philosophical Studies* 10: 23-27.

Anderson, Alan Ross. (1967). The formal analysis of normative systems. In N. Rescher (Ed.), *The logic of decision and action*, 147-213. Pittsburgh: University of Pittsburgh Press. (Originally published 1956.)

Anderson, Alan Ross. (1969). Comments on von Wright's "Logic and ontology of norms". In J. W. Davis, D. J. Hockney, & W. K. Wilson (Eds.), *Philosophical logic*, 108-113. Dordrecht: Reidel.

Åqvist, Lennart. (1984). Deontic logic. In D. Gabbay & F. Guenthner (Eds.), *Handbook of philosophical logic, vol. II*, 605-714. Dordrecht: Reidel.

Åqvist, Lennart. (1987). *Introduction to deontic logic and the theory of normative systems.* Napoli: Bibliopolis.

Åqvist, Lennart. (1997). Systematic frame constants in defeasible deontic logic. In D. Nute (Ed.), *Defeasible deontic logic*, 59-77. Dordrecht: Kluwer.

Asher, Nicholas, & Daniel Bonevac. (1997). Common sense obligation. In D. Nute (Ed.), *Defeasible deontic logic*, 159-203. Dordrecht: Kluwer.

Bailhache, Patrice. (1991). *Essai de logigue déontique.* Paris: Librairie Philosophique J. Vrin.

Belnap, Nuel. (1991). Before refraining: Concepts for agency. *Erkenntnis* 34: 137-169.
Belnap, Nuel. (1996). Agents in branching time. In B. J. Copeland (Ed.), *Logic and reality: Essays in pure and applied logic, in memory of Arthur Prior*, 239-271. New York: Oxford University Press.
Belnap, Nuel, Michael Perloff, & Ming Xu. (2001). *Facing the future: Agents and choices in our indeterminist world.* New York: Oxford University Press.
Belzer, Marvin, & Barry Loewer. (1993). Absolute obligations and ordered worlds. *Philosophical Studies* 72: 47-70.
Belzer, Marvin, & Barry Loewer. (1994). Hector meets 3-D: A diaphilosophical epic. *Philosophical Perspectives* 8: 389-414.
Belzer, Marvin, & Barry Loewer. (1997). Deontic logics of defeasibility. In D. Nute (Ed.), *Defeasible deontic logic*, 45-57. Dordrecht: Kluwer.
Brink, David O. (1994). Moral conflict and its structure. *The Philosophical Review* 103: 215-247.
Brown, Mark A. (1996). Doing as we ought: Towards a logic of simply dischargeable obligations. In M. A. Brown & J. Carmo (Eds.), *Deontic logic, agency, and normative systems. ΔEON '96: Third international workshop on deontic logic in computer science, Sesimbra, Portugal, 11-13 January 1996*, 47-65. Berlin: Springer.
Brown, Mark A. (1999). Agents with changing and conflicting commitments: A preliminary study. In P. McNamara & H. Prakken (Eds.), *Norms, logics and information systems: New studies in deontic logic and computer science*, 109-125. Amsterdam: IOS Press.
Brown, Mark A. (2000). Conditional obligation and positive permission for agents in time. *Nordic Journal of Philosophical Logic* 5: 83-112.
Bulygin, Eugenio. (2000). On norm-propositions. In W. Krawietz, R. S. Summers, O. Weinberger, & G. H. von Wright (Eds.), *The reasonable as rational? On legal argumentation and justification. Festschrift for Aulis Aarnio*, 129-134. Berlin: Duncker & Humblot.
Castañeda, Hector-Neri. (1981). The paradoxes of deontic logic: The simplest solution to all of them in one fell swoop. In R. Hilpinen (Ed.), *New studies in deontic logic*, 37-85. Dordrecht: Reidel.
Chisholm, Roderick M. (1963). Supererogation and offense: A conceptual scheme for ethics. *Ratio* 5: 1-14.
Chisholm, Roderick M. (1964). The ethics of requirement. *American Philosophical Quarterly* 1: 147-153.
Feldman, Fred. (1986). *Doing the best we can: An essay in informal deontic logic.* Dordrecht: Reidel.
Føllesdan, Dagfinn, & Risto Hilpinen. (1971). Deontic logic: An introduction. In R. Hilpinen (Ed.), *Deontic logic: Introductory and systematic readings*, 1-35. Dordrecht: Reidel.
Forrester, James W. (1984). Gentle murder, or the adverbial Samaritan. *The Journal of Philosophy* 81: 193-197.
Geach, Peter. (1982). Whatever happened to deontic logic? *Philosophia* 11: 1-12.
Hansson, Bengt. (1971). An analysis of some deontic logics. In R. Hilpinen (Ed.), *Deontic logic: Introductory and systematic readings*, 121-147. Dordrecht: Reidel. (Originally published 1969 in *Noûs* 3: 373-398.)
Harman, Gilbert. (1986). *Change in view: Principles of reasoning.* Cambridge, MA: MIT Press.
Hilpinen, Risto. (1981). Conditionals and possible worlds. In G. Fløistad (Ed.), *Contemporary philosophy: A new survey. Vol. 1: Philosophy of language, philosophical logic*, 299-335. The Hague: Martinus Nijhoff.
Horty, John F. (2001). *Agency and deontic logic.* New York: Oxford University Press.

Horty, John F. (Unpublished). Reasoning with moral conflicts. Draft of October 24, 2001.
Kamp, Hans. (1974). Free choice permission. *Proceedings of the Aristotelian Society* 74: 57-74.
Kamp, Hans. (1979). Semantics versus pragmatics. In F. Guenthner (Ed.), *Formal semantics and pragmatics for natural languages*, 255-287. Dordrecht: Reidel.
Knuutila, Simo. (1981). The emergence of deontic logic in the fourteenth century. In R. Hilpinen (Ed.), *New studies in deontic logic*, 225-248. Dordrecht: Reidel.
Kripke, Saul. (1963). Semantical considerations on modal logic. *Acta Philosophica Fennica* 16: 83-94.
Krogh, Christen, & Henning Herrestad. (1996). Getting personal: Some notes on the relationship between personal and impersonal obligation. In M. A. Brown & J. Carmo (Eds.), *Deontic logic, agency, and normative systems. ΔEON '96: Third international workshop on deontic logic in computer science, Sesimbra, Portugal, 11-13 January 1996*, 134-153. Berlin: Springer.
Lewis, David K. (1983). Scorekeeping in a language game. In D. K. Lewis, *Philosophical papers, vol. 1*, 233-249. New York: Oxford University Press. (Originally published 1979 in *Journal of Philosophical Logic* 8: 339-359.)
Lewis, David K. (2000a). Semantic analyses for dyadic deontic logic. In D. K. Lewis, *Papers in ethics and social philosophy*, 5-19. New York: Cambridge University Press. (Originally published 1974 in S. Stenlund (Ed.), *Logical theory and semantic analysis*, 1-14. Dordrecht: Reidel.)
Lewis, David K. (2000b). A problem about permission. In D. K. Lewis, *Papers in ethics and social philosophy*, 20-33. New York: Cambridge University Press. (Originally published 1979 in E. Saarinen, R. Hilpinen, I. Niiniluoto, & M. Provence Hintikka (Eds.), *Essays in honour of Jaakko Hintikka*, 163-175. Dordrecht: Reidel.)
Loewer, Barry, & Marvin Belzer. (1986). Help for the good Samaritan paradox. *Philosophical Studies* 50: 117-127.
Makinson, David. (1984). Stenius' approach to disjunctive permission. *Theoria* 50: 138-147.
Meyer, John-Jules Ch., & Roel J. Wieringa. (1993). Deontic logic: A concise overview. In J.-J. Ch. Meyer & R. J. Wieringa (Eds.), *Deontic logic in computer science: Normative system specification*, 3-16. New York: Wiley.
Moritz, Manfred. (1963). Permissive Sätze, Erlaubnissätze, und deontische Logik. In H. Bratt, S. Dunér, M. Moritz, & H. Regnéll (Eds.), *Philosophical essays dedicated to Gunnar Aspelin on the occasion of his sixty-fifth birthday the 23rd of September 1963*, 108-121. Lund: CWK Gleerup.
Nowell-Smith, P. H., & E. J. Lemmon. (1960). Escapism: The logical basis of ethics. *Mind* 69: 289-300.
Nute, Donald. (1985). Permission. *Journal of Philosophical Logic* 14: 169-190.
Nute, Donald. (1997). Apparent obligation. In D. Nute (Ed.), *Defeasible deontic logic*, 287-315. Dordrecht: Kluwer.
Nute, Donald, & Xiaochang Yu. (1997). Introduction. In D. Nute (Ed.), *Defeasible deontic logic*, 1-16. Dordrecht: Kluwer.
Prakken, Henry, & Marek Sergot. (1997). Dyadic deontic logic and contrary-to-duty obligations. In D. Nute (Ed.), *Defeasible deontic logic*, 223-262. Dordrecht: Kluwer.
Prior, Arthur N. (1967). *Past, present, and future*. New York: Oxford University Press.
Rohrbaugh, Eugene. (1997). *Scalar interpretation in deontic speech acts*. New York: Garland.
Ross, Alf. (1941). Imperatives and logic. *Theoria* 7: 53-71.
Sinnott-Armstrong, Walter. (1985). A solution to Forrester's paradox of gentle murder. *The Journal of Philosophy* 82: 162-168.

Stenius, Erik. (1963). The principles of a logic of normative systems. *Acta Philosophica Fennica* 16: 247-260.
Stenius, Erik. (1982). Ross' paradox and well-formed codices. *Theoria* 47: 49-77.
Tännsjö, Torbjörn. (1989). The morality of collective actions. *The Philosophical Quarterly* 39: 221-228.
Thomason, Richmond H. (1970). Indeterminist time and truth-value gaps. *Theoria* 36: 264-281.
Thomason, Richmond H. (1981a). Deontic logic as founded on tense logic. In R. Hilpinen (Ed.), *New studies in deontic logic*, 165-176. Dordrecht: Reidel.
Thomason, Richmond H. (1981b). Deontic logic and the role of freedom in moral deliberation. In R. Hilpinen (Ed.), *New studies in deontic logic*, 177-186. Dordrecht: Reidel.
Thomason, Richmond H. (1984). Combinations of tense and modality. In D. Gabbay & F. Guenthner (Eds.), *Handbook of philosophical logic, vol. II*, 135-165. Dordrecht: Reidel.
Tuomela, Raimo. (1989a). Actions by collectives. *Philosophical Perspectives* 3: 471-496.
Tuomela, Raimo. (1989b). Collective, action, supervenience, and constitution. *Synthese* 80: 243-266.
von Kutschera, Franz. (1986). Bewirken. *Erkenntnis* 24: 253-281.
von Wright, Georg Henrik. (1951). Deontic logic. *Mind* 60: 1-15.
von Wright, Georg Henrik. (1963). *Norm and action: A logical enquiry.* New York: Humanities Press.
von Wright, Georg Henrik. (1968). *An essay in deontic logic and the general theory of action. Acta Philosophica Fennica* Fasc. 21. Amsterdam: North Holland.
von Wright, Georg Henrik. (1969). On the logic and ontology of norms. In J. W. Davis, D. J. Hockney, & W. K. Wilson (Eds.), *Philosophical logic*, 89-107. Dordrecht: Reidel.
von Wright, Georg Henrik. (1971). Deontic logic and the theory of conditions. In R. Hilpinen (Ed.), *Deontic logic: Introductory and systematic readings*, 159-177. Dordrecht: Reidel.
von Wright, Georg Henrik. (1981a). On the logic of norms and actions. In R. Hilpinen (Ed.), *New studies in deontic logic*, 3-35. Dordrecht: Reidel.
von Wright, Georg Henrik. (1981b). Problems and prospects of deontic logic: A survey. In E. Agazzi (Ed.), *Modern logic - A survey: Historical, philosophical, and mathematical aspects of modern logic and its applications*, 399-423. Dordrecht: Reidel.
von Wright, Georg Henrik. (1988). Action logic as a basis for deontic logic. In G. di Bernardo (Ed.), *Normative structures of the social world*, 39-63. Amsterdam: Rodopi.
von Wright, Georg Henrik. (1999a). Deontic logic—as I see it. In P. McNamara & H. Prakken (Eds.), *Norms, logics and information systems: New studies in deontic logic and computer science*, 15-25. Amsterdam: IOS Press.
von Wright, Georg Henrik. (1999b). Deontic logic: A personal view. *Ratio Juris* 12: 26-38.
von Wright, Georg Henrik. (2000). On norms and norm-propositions. In W. Krawietz, R. S. Summers, O. Weinberger, & G. H. von Wright (Eds.), *The reasonable as rational? On legal argumentation and justification. Festschrift for Aulis Aarnio*, 173-178. Berlin: Duncker & Humblot.
Zanardo, A. (1985). A finite axiomatization of the set of strongly valid Ockhamist formulas. *Journal of Philosophical Logic* 14: 447-468.
Zanardo, A. (1991). A complete deductive system for since-until branching-time logic. *Journal of Philosophical Logic* 20: 131-148.
Zanardo, A. (1996). Branching-time logic with quantification over branches: The point of view of modal logic. *Journal of Symbolic Logic* 61: 1-39.