CONDITIONAL OBLIGATION AND POSITIVE PERMISSION FOR AGENTS IN TIME

MARK A. BROWN

This paper investigates the semantic treatment of conditional obligation, explicit permission (often called positive permission), and prohibition based on models with agents and branched time. In such models branches (rather than moments) are taken as basic, and the branching provides a way to represent the indeterminism which is normally presupposed by talk of free will, responsibility, action and ability. Careful treatment of the relation between ability and responsibility avoids many common problems with accounts of conditional obligation. Recognition of the generality often involved in conditional obligations makes possible a sensitive way of expressing some kinds of general prohibitions, which in turn makes it possible to account for the special role of explicit permission.

1. Introduction

Deontic logicians have been understandably eager to provide a simple core theory which might later, it is hoped, be adjusted and elaborated. This is a reasonable—even indispensable—strategy for dealing with any complex subject matter. Some simplifying assumptions in deontic logic are made consciously, with conscious acceptance of some of the risks involved. For example, we commonly choose to leave tacit the agent whose obligations are under discussion, though we will then be unable to discuss ways in which one agent's obligations might conflict with those of another. Other assumptions may be less conscious. For example, we normally give no thought to the possibility that in a sentence like "I ought to go home" the phrase "to go home" might not fully and precisely characterize the obligation involved, but might simply indicate one salient consequence of fulfilling one of my obligations.

Simplifying assumptions—both those we make consciously and those we fall into all unawares—can have powerful side effects, however, and as deontic logic matures, we must constantly reexamine our work to identify the assumptions we use and to consider their effects. From the first, treatments of deontic logic have been beset by a variety of "paradoxes"—or rather, as Makinson (1999) has more accurately put it, have

Nordic Journal of Philosophical Logic, Vol. 5, No. 2, pp. 83–112.

failed to deal satisfactorily with various benchmark examples. These failures, we must suspect, are often symptomatic of the acceptance of unduly simple assumptions.

In this paper I relinquish some common assumptions and embrace some complications that are often left aside. I do this because of the growing conviction that ignoring these details invites subtle mistakes in our thinking about deontic logic in general and about conditional obligation, prohibition and permission in particular.

I believe we need to keep the following considerations, among others, in or near the foreground:

- different agents often have different obligations;
- an agent's obligations can conflict with one another;
- agents are (presumed to be) free to make choices;
- agents can perform different actions at different times;
- our obligations change over time, partly as the result of our actions and the actions of others;
- our expressions of obligations are of various sorts:
  —some are, but some are not, intended to capture the exact content of the indicated obligation;
  —some are, but some are not, intended to indicate general, or standing, obligations.
- some, though not all, prohibitions are general in character;
- explicit permission often works against a background of general prohibition.

I argue that, partly through consideration of such details, we find that a number of simple (and initially plausible) assumptions commonly made—whether consciously or not—in the formalization of deontic logic are in fact highly suspect. Among these are the following:

- that a single modal operator can serve virtually all deontic purposes;
- that deontic logic can be treated apart from the logic of action;
- that personal obligation and personal permission are duals of one another;
- that unconditional obligation can be defined in terms of conditional obligation;
- that the conditions involved in conditional obligations are normally made fully explicit in our expression of such obligations;
- that there is basically just one kind of permission.

The first three of these points have been discussed in Brown 1999 and somewhat more fully, along with the fourth and fifth, in Brown (forthcoming). These will be reviewed in Sections 2–5 below. The sixth assumption will be discussed in Section 6 below.

In Brown 1996 I introduced discussion of a particular species of obligation: simply dischargeable obligations. These are *dischargeable* in the sense that they can be fulfilled and then cease to be further in effect. They are thus unlike standing obligations, such as the obligation to honor our parents, which continue in force no matter how much we do to fulfill them. Dischargeable obligations are also presumably distinct from prudential, or strategic, *ought*s. While these seem circumstantial, like dischargeable obligations, they are in a sense relentless, like standing obligations: I must always and in all circumstances be prudent, and choose the best course of action, if such there be. Having made a good choice at this moment brings no relief from such general prudential obligations.

Simply dischargeable obligations are *simply* dischargeable in the sense that they are not cumulative, i.e. they do not merge with other obligations the way two monetary debts to the same person might be thought to merge into a single larger debt, thus losing their individual identities. In this paper, I continue to focus on simply dischargeable obligations, giving an account of conditional obligation in Section 8 in terms of (unconditional) simply dischargeable obligations. This account of conditional obligation will help make it possible to give an account of standing obligations. Cumulative obligations present special problems, and will have to wait for another occasion.

2. Two Types of Deontic Operators

At least two distinct types of modal operator (called Type 1 and Type 2 in the terminology adopted in Brown 1996, Brown 1999, Brown forthcoming) are needed to express ordinary unconditional claims of obligation, because we make (and *need* to be able to make) such claims sometimes with lesser, sometimes with greater, specificity as to the exact nature of the obligation. A deontic operator $\mathcal{O}^1$ of Type 1 will indicate an obligation by citing one salient consequence of its fulfillment, with $\mathcal{O}^1 p$ interpreted to mean that the (tacit) agent has some (not fully specified) obligation or other whose fulfillment will necessitate the truth of p. However the truth of p will in general be no guarantee of the fulfillment of the obligation which underlies the claim that $\mathcal{O}^1 p$. In short, the truth of p will be a necessary, but not a sufficient, condition for the fulfillment of some otherwise unspecified obligation. Operators of Type 1 will be subject to the rule RM:

RM $\quad$ *From* $\vdash p \rightarrow q$, *infer* $\vdash \mathcal{O}^1 p \rightarrow \mathcal{O}^1 q$.

If p expresses a consequence of fulfilling one of my obligations, and p

entails q, then q expresses another consequence of fulfilling that same obligation. To track chains of such consequences, acknowledging the deontic status of each, we will want to use a Type 1 operator.

An operator $\mathcal{O}^2$ of Type 2 will be used to state (up to logical equivalence) precisely what the obligation itself is, with $\mathcal{O}^2 p$ interpreted to mean that the (tacit) agent has an obligation for whose fulfillment the truth of p is not only a *necessary*, but also a *sufficient* condition. Operators of Type 2 will not satisfy the rule RM, but will instead satisfy the weaker rule RE:

RE *From* $\vdash p \leftrightarrow q$, *infer* $\vdash \mathcal{O}^2 p \leftrightarrow \mathcal{O}^2 q$.

Operators of Type 2 are therefore not normal modal operators, but are instead the weaker sort of operator which Chellas (1980) calls *classical* operators, first introduced in Segerberg 1971.

Many systems of monadic deontic logic employ a Type 1 modal operator, including all systems based on normal systems (Kripke systems) of modal logic. Because such operators obey the rule RM, they also validate the schema known as Ross's "Paradox", first discussed in Ross 1941:

$$\vdash \mathcal{O}^1 p \rightarrow \mathcal{O}^1 (p \vee q).$$

When p expresses the claim that I mail your letter and q expresses the claim that I burn your letter, this entailment seems to have the effect of saying that if I have promised to mail your letter I can fulfill that obligation, or at least I can fulfill *some* obligation, by burning your letter and thus making $p \vee q$ true. Careful, consistent interpretation of the Type 1 operator will prevent such a reading, of course: all that is really said here is that if I ever fulfill my promise to mail your letter, then it will at that point be true that I either mail it or burn it.

The psychological force of the paradox derives from the fact that it is natural to construe the sentence "I ought to mail your letter" as expressing (very nearly) precisely the nature of the obligation under consideration, and thus as giving necessary *and sufficient* conditions for its fulfillment, not merely one necessary consequence of its fulfillment; hence it is humanly tempting to fall carelessly into a Type 2 reading of whatever operator is used in the antecedent, and then follow through with a Type 2 reading of the formula in the consequent. If we yield to this temptation, we end up taking $p \vee q$ as expressing the sufficient condition for fulfillment of this (or at least of some) obligation.

The naturalness of this mistake argues for the need for a Type 2 operator. We do sometimes wish to be understood as expressing the precise content of an obligation. This will, ideally, be the case whenever we agree on some contractual arrangement, for example. Whenever there is something we clearly can mean to express, there ought to be some way to express it clearly. In this case, a Type 2 operator is what is needed to make that possible.

On the other hand, we also need to be able to identify and track the consequences of, and/or the means of fulfilling, our obligations; and it seems natural to express the character of such conditions using *ought*. When I have an obligation to mail your letter next Monday, merely mailing your letter next week will not by itself fulfill any obligation. Nonetheless, we find it natural and appropriate to say that (because I ought to mail your letter next Monday) I ought to mail your letter next week. To express this less specific *ought* without implying that merely mailing the letter sometime next week will by itself fulfill any obligation, we need a Type 1 operator.

Some of our obligations may well be too subtle and complex for us to be able to give any simple sentences which capture necessary and sufficient conditions for their fulfillment. It is in any case certainly true that spelling an obligation out fully would often be more than is desired for the purposes at hand. When I say that I ought to go home, it is not normally the case that merely going home will fulfill the obligation that lies behind my announcement. Perhaps I ought to fix dinner for my family, who are at home, and I can only do this if I first go home myself. Or perhaps I ought to fix the broken window before nightfall. But for present purposes the exact nature of the obligation doesn't matter. What matters is that, to fulfill my obligations, I must now go home. In such cases, we are accustomed to directing attention to the existence of an obligation using sentences which merely indicate salient consequences of fulfilling the obligation, i.e. we have recourse to Type 1 operators.

In short, both types are needed for a full treatment of ordinary normative discourse. Unfortunately, although there is a close semantic connection between the two, neither is syntactically definable from the other in any comfortable way. In particular, expressing Type 1 operators using an operator of Type 2 would require something like explicit quantification over propositions or over obligations. We would prefer not to make such quantification explicit in the syntax, but to leave it "hidden" within the underlying semantics. It could be argued that part of the point of modal logic is its ability to keep certain forms of quantification (over possible worlds, for example) tacit in the syntax by reserving them to the semantic metalanguage.

Some clarification of the distinction between Type 1 and Type 2 operators is appropriate here. Type 2 operators are clearly not normal modal operators, since they defy RM. One might easily assume that Type 1 operators will differ partly by being normal. But although normal modal operators will be of Type 1, it is also possible to have non-normal operators of this type. Indeed, if deontic operators are to be adequate to the task of expressing potentially conflicting obligations, then even Type 1 deontic operators will not be normal modal operators, since they will not obey the rule RR of agglomerativity:

RR $\quad$ *From* $\vdash(p \wedge q) \rightarrow r$, *infer* $\vdash(\mathcal{O}^1 p \wedge \mathcal{O}^1 q) \rightarrow \mathcal{O}^1 r$.

The notion of a Type 2 operator, as first introduced in Brown (1996), can in retrospect be seen to be a bit vague. The terminology was introduced in a way that would suggest it should apply only to operators which obey the rule RE but defy the rule RM, i.e. to operators generating systems which, in the terminology of Chellas (1980), are classical but non-monotonic. But the discussion also suggested that the terminology might be applied still more narrowly, to operators which in addition obey an unusual modal cancellation rule, which we might see as an inverse of RE:

RE_i $\quad$ *From* $\vdash \mathcal{O}^2 p \leftrightarrow \mathcal{O}^2 q$, *infer* $\vdash p \leftrightarrow q$.

Then again, and confusingly, the notion of a Type 2 operator has sometimes been described as if it encompasses any operator which defies the rule RM—i.e. any operator generating a system which, in the terminology of Chellas (1980), is non-monotonic—without regard for whether it is even classical. To make things still worse: even in normal systems, not all modalities (as contrasted with primitive operators) satisfy the rule RM—negative modalities such as non-necessity instead obey its converse RM_c. For an operator or modality $\circledast$ the rule would be:

RM_c $\quad$ *From* $\vdash p \rightarrow q$, *infer* $\vdash \circledast q \rightarrow \circledast p$.

In this paper, I will count any operator or modality which obeys either RM or RM_c as being of Type 1, and any operator or modality which obeys RE but defies both RM and RM_c as being of Type 2. However, the critical basic impersonal deontic modality of Type 2 will in fact also obey RE_i, and it is important for its deontic role that it should do so.

3. Deontic Logic and the Logic of Action

Suppose we take the seemingly trivial step of making explicit the agent that is normally tacit in deontic logic. Consider first how this will affect our treatment of Type 2 operators. Instead of writing that $\mathcal{O}^2 p$, let us for the moment write that $\mathcal{O}^2 p(\alpha)$, where α is the agent in question.

Once we bring the agent explicitly into the picture, we must recognize that the only method of fulfilling an agent's obligation is through that agent's actions, not through mere happenstance or accident, or the agency of others. My obligation to help the accident victim will not be fulfilled unless *I* act to aid the victim. The obligation might cease in some other way (someone more able takes over, relieving me of my obligation, perhaps), but it will not be *fulfilled* unless *I* act appropriately. With that in mind, it becomes evident that the exact expression of what must occur for the fulfillment of my obligation must be of the agentive form "I do such-and-such". Expressing this adequately, i.e. in a way that differentiates it from a non-agentive expression of the form "such-and-such happens", will require that we employ a logic of action.

Knowing that we must be prepared to have agentive action operators in place anyway, we can now recognize the distinct possibility that a personal obligation can be treated as an impersonal obligation to have a personal agentive proposition be true. I.e. we can expand $\mathcal{O}^2 p(\alpha)$ as $\mathcal{O}^2 \Delta_\alpha p$, where $\Delta_\alpha p$ expresses the (non-deontic) agentive claim that agent α "does" p, i.e. brings it about that p, sees to it that p, or acts in a way that assures that p. Versions of this compositional strategy, as we may call it, have been considered by many authors, including Chisholm (1964), Kanger (1971), Makinson (1986), Horty (1996) and Sergot (1999). García (1986) discusses an early proposal by Meinong to the same effect. Objections to various versions of the compositional strategy have also been put forward from time to time. The most sustained and detailed objections, presented in Krogh and Herrestad 1996, have been shown in Brown 1999 to be inapplicable to the version offered here, since all the arguments in that work are directed against the inadequacies of normal modal operators for such deontic purposes, whereas the operators used here are all non-normal. Objections to another version are discussed in Horty 1996, but these, too, are inapplicable to our account of obligation. Horty is concerned with a utilitarian account of *ought* and treats it as a strategic (or perhaps prudential) notion, rather than as one addressed to simply dischargeable obligations. The operator he considers for this purpose is a normal modal operator, and his arguments against the compositional strategy apply to that normal operator, not to the non-normal operators considered here.

No doubt it is technically possible to have operators of each of the four types that result from cross-classifyingagentive/non-agentive with normative/non-normative operators, but surely the most appealing treatment of operators that are both agentive and normative would be to treat them as compound modalities. We would then have the following simple and satisfying pattern:

	non-agentive	agentive
non-normative	p	$\Delta_\alpha p$
normative	$\mathcal{O}p$	$\mathcal{O}\Delta_\alpha p$

Incorporating the agency directly into the deontic operator, in a single normative agentive operator instead of a compound modality, would leave us with no ready and natural way to relate non-normative statements concerning our actions to normative statements of our obligations, so as to determine whether the obligation had or had not been fulfilled by such actions.

Similar considerations apply to Type 1 operators: we can and perhaps should treat personal obligations, understood in the Type 1 manner, as impersonal obligations of Type 1 for the performance of personal actions. However we will not want, in this case, to attach the Type 1 impersonal obligation operator to an expression involving an action operator which purports to give a full and precise expression of the action in question. An action operator of such a sort would be a Type 2 action operator, and combined with the Type 1 obligation operator would create a composite modality of Type 2, not Type 1. It would then be of no use in tracking the consequences of personal obligations. We want instead an action expression which only claims to give one of the *consequences* of the performance of the obligatory action. The key operator in such an action expression will need to be of Type 1, and as a result it might be misleading to call it an *action* operator. For lack of a better term, we might perhaps call such a Type 1 operator an *outcome* operator. But whatever we may wish to call it, a Type 1 operator closely related to our Type 2 action operator will be needed for a rich logic of action for the same reasons that both Type 1 and Type 2 obligation operators are needed for a rich deontic logic. With the Type 2 action operator, but without the corresponding Type 1 outcome operator, we could describe the actions performed but could not describe the further consequences of those actions in any way that traced them to the actions of an agent. To cite just one simple kind of example: when the agent made a left turn we would be unable to express the less specific fact that she turned unless we were willing to take the making of the left turn and the making of the turn to be distinct (but somehow mysteriously related) actions, or were willing to consider the mere turning as something which just happened to her and was not any of her doing. An outcome operator enables us to identify the mere turning as in some sense the agent's doing, without ascribing to it the status of an action separate from her turning left.

One leading candidate for the role of the personal action operator Δ_α is the dstit_α operator devised independently by von Kutschera (1986) and Horty (1989) (see also Horty and Belnap 1995). This operator does not obey the rule RM. As a consequence, using the dstit_α operator as Δ_α, we would find that the combination $\mathcal{O}^1\Delta_\alpha p$ would not obey the rule RM (or its converse RM_c) either, and hence the composite deontic modality $\mathcal{O}^1\Delta_\alpha$

would not be of Type 1. But the closely related cstit$_\alpha$ operator devised by Chellas (1964), independently reinvented in Brown 1990, and also discussed in Horty and Belnap 1995, is of Type 1 and can serve nicely as an outcome operator in tandem with the action operator dstit$_\alpha$.

4. Explicit Prohibition and Tacit Permission

Another aspect of the logic of action is relevant for deontic logic. A sensitive treatment of the logic of action, as in Belnap and Perloff (1988) and Belnap and Perloff 1992, or Horty and Belnap 1995, indicates that the notion of refraining from doing something is a moderately complex one. Specifically, refraining from speaking (for example) is more complex than merely not speaking. Not speaking might not be deliberate: for example it might result from inattention or, worse, from inability to speak. We don't want to say that the day-dreaming orator, the sleeping gossip, or the deaf-mute refrains from speaking, much less that the dog does, and still less that a rock does. Refraining from speaking, to continue the example, is something more like deliberately seeing to it that one doesn't speak (though one could).

This is of consequence for deontic logic, because it seems reasonable to hold with Belnap and Bartha (1995) that, in at least one sense of "permitted", being permitted to perform a certain action amounts to not being explicitly forbidden to perform that action, and that being explicitly forbidden to perform, or explicitly prohibited from performing, amounts to being explicitly obligated to refrain. I shall call this the notion of *tacit* permission (it might also be called *default* permission), to be contrasted with *explicit* permission, which will be taken up in Section 6.

We then have

$$\mathscr{P}\Delta_\alpha p \equiv \neg \mathscr{F}\Delta_\alpha p \equiv \neg \mathscr{O}\mathscr{R}_\alpha p$$

The view that refraining is not merely the negation of doing

$$\mathscr{R}_\alpha p \not\equiv \neg \Delta_\alpha p$$

will then have the effect that tacit permission—i.e. not being explicitly forbidden—is not simply the dual of obligation:

$$\mathscr{P}\Delta_\alpha p \equiv \neg \mathscr{O}\mathscr{R}_\alpha p \not\equiv \neg \mathscr{O} \neg \Delta_\alpha p$$

Moreover, tacit permission to do something—not being explicitly forbidden to do it—is not simply the dual of obligation to do it:

$$\mathscr{P}\Delta_\alpha p \equiv \neg \mathscr{O}\mathscr{R}_\alpha p \not\equiv \neg \mathscr{O}\Delta_\alpha \neg p$$

The principle that *ought* implies *may* will not take the form of obligation implying its dual, then, but will be considerably more subtle.

Belnap follows the view adopted earlier in Pörn 1970, that refraining from doing p is itself a kind of action and, more specifically, it is seeing to it that one does *not* see to it that p:

$$\mathscr{R}_{\alpha}p \equiv \Delta_{\alpha} \neg \Delta_{\alpha}p.$$

On that particular account of refraining (which I adopt in this paper), we then get as our account of tacit permission:

$$\mathscr{P}\Delta_{\alpha}p \equiv \neg \mathscr{F}\Delta_{\alpha}p \equiv \neg \mathscr{O}\Delta_{\alpha} \neg \Delta_{\alpha}p.$$

But the specifics of the account of refraining are not critical here. The key point is just that refraining is some kind of acting. The details of that kind of acting are not needed for the conclusion that tacit permission is not the dual of obligation. We could have any non-trivial operator $\circledast$ other than negation in place of $\neg \Delta_{\alpha}$, and the conclusion would be essentially the same: $\mathscr{P}\Delta_{\alpha}$, i.e. $\neg \mathscr{O}\Delta_{\alpha} \circledast$, is not the dual of $\mathscr{O}\Delta_{\alpha}$. Moreover, $\mathscr{O}$ is not definable in terms of $\mathscr{P}$, as it would be if the two were duals.

5. Conditional and Unconditional Obligation

Having come to recognize that unconditional obligation is far more complex and subtle than most treatments would acknowledge, we are in a position now to anticipate similar complexity in the notion of conditional obligation. First, perhaps it would be useful to set aside one class of conditional expressions involving obligation, a class of sentences we might call *simple obligation conditionals*. Here I have in mind sentences like

- If Dawn is my mother, then I ought to take a book to Dawn; and
- If today is Tuesday, then I ought to pay my gas bill today.

Simple obligation conditionals of this sort will be adequately formalized as ordinary conditionals whose consequents happen to be expressions of obligation. In some cases, they may be based on the existence of a differently described obligation (e.g. my obligation to take a book to my mother) and may be intended merely to indicate a condition under which that obligation can be redescribed (e.g. as an obligation to take Dawn a book). In other cases, there may be a mere material conditional involved, with no hint of any deontically significant connection between antecedent and consequent. Perhaps there are other sorts of sentences for which this simple rendering is appropriate. Sentences capable of this simple rendering are, I believe, not normally under consideration when we undertake to discuss conditional obligation. So from here forward, let us set aside such simple obligation conditionals.

Let us introduce the temporary notation $\mathcal{O}(p/q)$ to translate the claim that p is obligatory on condition that q. From the very first discussions of conditional obligation, in von Wright 1956, the suggestion has been made, and widely accepted, that unconditional obligation can be treated as a limiting case of conditional obligation, by letting $\mathcal{O}p$ be an abbreviation for $\mathcal{O}(p/\top)$, where $\top$ is any tautology. This is often taken to be obviously correct, or at worst to be a convenient and harmless technical device.

But a converse suggestion, made by Anderson (1959 and 1967), should alert us to a problem with this proposed reduction of unconditional to conditional obligation. Anderson suggested that $\mathcal{O}(p/q)$ should be understood as $\Box(q \rightarrow \mathcal{O}p)$, where $\Box$ is some species of necessitation operator. This would, of course, define conditional obligation in terms of unconditional obligation. The notable feature of this suggestion, apart from its reversing the order of definition proposed by von Wright, is that it introduces an element of generality into the account of conditional obligation. An unconditional obligation may be highly localized in time and circumstance: right now, under just my current conditions, I may have an obligation to take a book to my mother, having just promised to do so. In a few minutes this obligation will have been fulfilled and discharged. At that point, that obligation will cease to be in force. Anderson's treatment brings to mind the thought that conditional obligations, or at least some conditional obligations, will not be like this—that, in particular, they represent a kind of obligation which can in principle arise again and again, whenever the appropriate condition occurs. Roughly, the idea is that "I ought to take a book to my mother, on condition that I have promised to do so" is related to "I ought to take a book to my mother" (in a context in which I have promised to do so) in something of the same fashion as "all men are mortal" is related to "Socrates is mortal" (in a context in which "Socrates" names a man). The latter is particular; the former is general. As a consequence, a rule of material detachment (or something very close to one) should apply to statements of conditional obligation.

More precisely, I am suggesting that statements of conditional obligation are ordinarily to be understood as involving tacit generalization over a span of time. The point of a statement of conditional obligation such as "You ought to return a book on condition that you have borrowed it" is to make it clear that *whenever* you borrow a book, you thereby incur a specific obligation to return it; that when you return that book, that specific obligation ceases; but *also* that should you subsequently borrow another book, you will thereby incur another specific obligation. This conditional obligation will be triggered each time its condition is satisfied, generating a new unconditional obligation each time. This stands in contrast to a simple obligation conditional such as "If you borrowed that book, you ought to return it'', which can be interpreted as applying only to the present moment, and can therefore be triggered only once, if at all.

Consider, however, what would be involved in accepting von Wright's reduction. Consider a situation in which I have promised to take my mother a book. Because of this, I am now under an obligation to take my mother a book. The reasoning so far can reasonably be seen as the application of a principle of material detachment to a standing conditional obligation:

> I ought to take my mother a book
> on condition that
> I have promised to do so.
> I have promised to do so.
> Therefore:
> I ought to take my mother a book.

On von Wright's account, however, the acceptance of the reduction of unconditional to conditional obligation should lead us to extend this line of reasoning as follows:

> Therefore:
> I ought to take my mother a book
> on condition that
> today is Tuesday or is not Tuesday.

Thus a general conditional obligation with a contingent triggering condition has been transformed, once that triggering condition is first fulfilled, into one whose triggering condition is *always* fulfilled, i.e. from a conditional obligation which is sometimes triggered into one which is always triggered.

If von Wright is correct, and if the rule of material detachment still applies, then I will still (or again) be obliged to take my mother a book. Indeed I will be condemned to a life of perpetual obligation to take books to my mother in spite of having already fulfilled my original promise.

In Section 8 below we shall see reasons why a simple principle of material detachment for conditional obligation is inaccurate. Nonetheless a slightly modified principle, one which takes into account the question whether the triggered obligation would still be fulfillable, will still apply, and in normal circumstances will give much the same results. Hence this difficulty with von Wright's account will still pertain.

Evidently von Wright did not take into account the question how our obligations vary over time. In fairness, Anderson probably didn't either, but the necessity operator in Anderson's account can at least prompt the thought that some generality is involved in conditional obligation, and a little further thought suggests that it is generality with respect to circumstance, including generality with respect to time.

We will be better able to consider just how to express such generality after we introduce a comprehensive formal semantics for unconditional obligation as background for our discussions of obligations and conditional obligations, taking into account the points mentioned so far.

6. The Problem of "Positive" Permission

So-called "positive" permission, which I prefer to call *explicit* permission, has been a traditionally difficult notion to incorporate into deontic logic. We sometimes issue specific and explicit permissions, and feel that something has been accomplished thereby. But suppose our only account of permission is one more or less like the one suggested in Section 4, according to which permission consists in not being forbidden. Then one would think that permission would be automatic in situations in which no corresponding prohibition has been issued, and would be inconsistent otherwise. What, then, would be the point of explicitly permitting some action?

For example, suppose a mother has issued no prohibition against eating cookies, and now voices a permission to eat a cookie. If permission to eat cookies consisted simply in not being forbidden to do so, then one would think that the children were already permitted to eat a cookie, and hence that the mother's new pronouncement would be more in the nature of a reminder than of a genuine statement of permission, i.e. one which *grants* permission. On the other hand, if the mother had forbidden the eating of cookies or, perhaps more plausibly, had issued a general injunction against eating between meals, then in issuing the permission she would seem to be contradicting herself: the children are, but now also are not, forbidden to eat cookies. Intuitively, of course, there is no special problem here, but the point is simply that it has been difficult for simple formal treatments of deontic logic to make room for our intuition in this matter.

So one problem for a formal account of explicit permission is to avoid treating a non-redundant explicit permission as inconsistent with the prior prohibitions to which it offers an exception. This and related problems have provoked some authors (e.g. Nute (1997)) to explore logics of defeasible deontic reasoning and others (e.g. Prakken and Sergot (1996)) to invoke the notion of sub-ideal worlds introduced in Jones and Pörn 1985. Most accounts of defeasible reasoning seem to me *ad hoc*, however, and accounts in terms of ideal and sub-ideal worlds seem too coarse-grained. I shall propose another solution, which I find more straightforward.

In Section 9, I will propose an account of explicit permissions as exceptions to general prohibitions. To make that possible without contradiction and without introducing some form of defeasible logic, it will be necessary to provide a special status for explicit permissions in the semantics. So let us turn now to the semantics.

7. Semantics

7.1. Models

We need to be able to discuss actions of free agents over time. To do so, we need to incorporate a temporal logic, but more specifically an indeterminist temporal logic—a logic of branching time—of the sort discussed in Thomason 1970, Thomason 1981, and in more technical detail in Zanardo 1985, Zanardo 1991, Zanardo 1996, Brown and Goranko 1999, and used as a basis for the logic of action in Belnap and Perloff 1988, Belnap and Perloff 1992, Horty and Belnap 1995, Brown 1996, and Brown 1999. Within such a system, evaluation of formulas occurs, not with respect to moments of time, but rather with respect to the elements of a set $\mathbb{B}$ of *branches* of time (equivalently, but more cumbersomely: moment/history pairs in which the moment falls within the history). Propositions correspond to sets of evaluation-points, i.e. to sets of branches of time (sets of moment/history pairs; subsets of $\mathbb{B}$). Taking branches as the primary constituents of branching time, as I shall, we consider moments to be equivalence classes of coeval branches, i.e. we identify the moment with the set of branches which can be said to originate at that moment. Since a moment is a set of branches, it can be considered a special sort of proposition—a proposition declaring the time, as it were.

The various choices open to an agent at a moment correspond to various subsets of the branches originating at that moment, and thus each of them can also be considered a proposition. Within each choice, the differing branches represent various alternative ways the details of the execution of that choice might work out, including details corresponding to the independent choices being made simultaneously or subsequently by other agents.

Models, in addition to providing a set $\mathbb{B}$ of branches structured into a tree, must also provide a set $\mathbb{A}$ of agents and a choice function $\mathbb{C}$ which must specify, for each agent at each moment in time, what choices are open to that agent at that moment, i.e. how the set of branches originating at that moment is partitioned into a family of choices for that agent at that moment. Since the moment is identified with the set of branches said to originate at that moment, we can say that for each given agent α, $\mathbb{C}$ partitions the moment itself into actions available to α at that moment.

Finally, models must provide an impersonal deontic function $\mathbb{D}$, an explicit permission function $\mathbb{E}$, and a valuation $\mathbb{V}$. $\mathbb{D}$ will specify, for each branch of time (but without regard to any agent), which propositions correspond to ones that are (impersonally) obligatory, along that branch. $\mathbb{E}$ will specify, for each branch of time, what explicit permissions have been issued. $\mathbb{V}$ will specify which atomic formulas are true along which branches of the model. Just as we handle personal obligation as a species of impersonal obligation, we will also handle personal explicit permission, i.e. explicit permissions issued to a particular agent, as a species of impersonal explicit permission, with the intended agent specified in the statement of what is impersonally permitted.

To express the claim that branches b and c in $\mathbb{B}$ are *coeval*, i.e. begin at the same moment, I write that $\mathsf{b} \sim \mathsf{c}$. The equivalence relation $\sim$ then partitions the set $\mathbb{B}$ of branches into moments. We write that $\mathsf{b} < \mathsf{c}$ to express the claim that the branch b is a (proper) *extension* of c, with the branch b starting at an earlier moment than c. When $\mathsf{b} < \mathsf{c}$ we also say that c is a (proper) *sub-branch* of b.

Throughout this paper, I reserve early lowercase Greek letters for agents (in the semantics) and names of agents (in the syntax); the lowercase Greek "ϕ" will be used as a schematic variable over agents. Open (blackboard bold) capitals are used for the sets and functions that serve as the chief constituents of models. Sans serif characters are used for various other objects associated with models, e.g. elements of sets. With these understandings, we can define the class of models appropriate for our discussion as follows.

By a *branch frame*, I mean a structure $\langle \mathbb{B}, \sim, < \rangle$ such that:

(1) $\mathbb{B}$ is a non-empty set (of *branches* of time);
(2) $\sim$ is an equivalence relation (reflexive, symmetric and transitive) on $\mathbb{B}$;
(3) for each $\mathsf{b} \in \mathbb{B}$, $<$ is a linear ordering (antisymmetric, transitive and comparable) on $\{\mathsf{c} \in \mathbb{B} : \mathsf{c} < \mathsf{b}\} \cup \{\mathsf{b}\} \cup \{\mathsf{c} \in \mathbb{B} : \mathsf{b} < \mathsf{c}\}$;
(4) $(\forall \mathsf{b}, \mathsf{c} \in \mathbb{B} : \mathsf{b} \sim \mathsf{c})[\neg(\mathsf{b} < \mathsf{c})]$; (disjointness of $\sim$ & $<$)
(5) $(\forall \mathsf{a}, \mathsf{b}, \mathsf{c} \in \mathbb{B} : \mathsf{a} < \mathsf{b} \sim \mathsf{c})(\exists \mathsf{d} \in \mathbb{B} : \mathsf{a} \sim \mathsf{d})[\mathsf{d} < \mathsf{c}]$; (weak diagram completion)
(6) $\forall \mathsf{c}, \mathsf{d} \in \mathbb{B} : \mathsf{b} \sim \mathsf{c} \neq \mathsf{b})(\exists \mathsf{a} \in \mathbb{B} : \mathsf{b} < \mathsf{a})(\forall \mathsf{d} \in \mathbb{B} : \mathsf{c} < \mathsf{d}) \neg(\mathsf{a} \sim \mathsf{d})$. (maximality of branches)

When $\mathsf{b} \sim \mathsf{b}'$, we say that b' is *coeval with*, or *an alternative to* b; when $\mathsf{b} < \mathsf{b}'$, we say that b' is a *sub-branch of* b, and that b *extends* b'.

On any branch frame, we define some other relations between branches:

$\mathsf{b} > \mathsf{b}$	iff	$\mathsf{c} < \mathsf{b}$;	$\mathsf{b} \leq \mathsf{c}$	iff	$\mathsf{b} < \mathsf{c} \vee \mathsf{b} = \mathsf{c}$;
$\mathsf{b} \prec \mathsf{c}$	iff	$(\exists \mathsf{a} \in \mathbb{B})[\mathsf{b} \sim \mathsf{a} < \mathsf{c}]$;	$\mathsf{b} \preceq \mathsf{c}$	iff	$\mathsf{b} \prec \mathsf{c} \vee \mathsf{b} \sim \mathsf{c}$.

We take propositions to be adequately represented, for present purposes, by sets of branches, so the set $\mathbb{P}$ of *propositions* is $\mathscr{P}\mathbb{B}$. We also define, for each branch b, some related propositions (sets of branches):

$\mathsf{F}(\mathsf{b}) =_{df} \{\mathsf{c} \in \mathbb{B} : \mathsf{b} < \mathsf{c}\}$; (the *future* of b)
$\mathsf{P}(\mathsf{b}) =_{df} \{\mathsf{c} \in \mathbb{B} : \mathsf{c} < \mathsf{b}\}$; (the *past* of b)
$\mathsf{H}(\mathsf{b}) =_{df} \mathsf{P}(\mathsf{b}) \cup \{\mathsf{b}\} \cup \mathsf{F}(\mathsf{b})$; (the *history* of b)
$\mathsf{M}(\mathsf{b}) =_{df} \{\mathsf{c} \in \mathbb{B} : \mathsf{c} \sim \mathsf{b}\}$; (the (initial) *moment* of b)

We let $\mathbb{M}$ be the set of all moments, i.e. $\{\mathsf{M}(\mathsf{b}) : \mathsf{b} \in \mathbb{B}\}$.

The rationale for the constraints 1–6 is discussed in Zanardo 1996, and constraint 1 is also discussed in Brown 1999.

By a *branch choice frame*, I mean a structure $\langle \mathbb{A}, \mathbb{B}, \mathbb{C}, \sim, < \rangle$ in which the component structure $\langle \mathbb{B}, \sim, < \rangle$ is a branch frame, and:

(7) $\mathbb{A}$ is a non-empty set of *agents*;
(8) $\mathbb{C}$ (the *choice function*) provides, for each branch $\mathsf{b} \in \mathbb{B}$ and each agent $\alpha \in \mathbb{A}$, a partition $\mathbb{C}(\alpha, \mathsf{b})$ of $\mathsf{M}(\mathsf{b})$ such that:
(8.1) $(\forall \mathsf{b}, \mathsf{c} \in \mathbb{B})(\forall \alpha \in \mathbb{A})[\mathsf{b} \sim \mathsf{c} \Rightarrow \mathbb{C}(\alpha, \mathsf{b}) = \mathbb{C}(\alpha, \mathsf{c})]$;
(8.2) for each branch b, for each agent $\alpha \in \mathbb{A}$, let A_α be one choice from $\mathbb{C}(\alpha, \mathsf{b})$; then $\cap\{\mathsf{A}_\alpha : \alpha \in \mathbb{A}\} \neq \varnothing$; (independence of agents)
(8.3) $(\forall \mathsf{b}, \mathsf{c}, \mathsf{d}, \mathsf{e} \in \mathbb{B} : \mathsf{b} < \mathsf{d} \sim \mathsf{e} > \mathsf{c} \sim \mathsf{b})$ $(\forall \alpha \in \mathbb{A})$ $(\forall \mathsf{A} \in \mathbb{C}(\alpha, \mathsf{b}))$ $[\mathsf{b} \in \mathsf{A} \Rightarrow \mathsf{c} \in \mathsf{A}]$. (no choice between undivided branches)

In Horty and Belnap 1995 (among other places) one can find a discussion of the rationale for constraints 8.1–3.

By a *branch deontic frame*, I mean a structure $\langle \mathbb{A}, \mathbb{B}, \mathbb{C}, \mathbb{D}, \mathbb{E}, \sim, < \rangle$ in which the structure $\langle \mathbb{A}, \mathbb{B}, \mathbb{C}, \sim, < \rangle$ is a branch choice frame, and:

(9) $\mathbb{D}$ (the *obligation function*) is a function from branches to sets of propositions, i.e. from $\mathbb{B}$ to $\mathscr{P}\mathscr{P}\mathbb{B}$, subject to the following constraints:

(9.1) $(\forall \mathsf{b} \in \mathbb{B})(\forall \mathsf{O} \in \mathbb{D}(\mathsf{b}))(\exists \mathsf{c} : \mathsf{b} \prec \mathsf{c})[\mathsf{c} \notin \mathsf{O}]$; (non-triviality)
(9.2) $(\forall \mathsf{b} \in \mathbb{B})(\forall \mathsf{O} \in \mathbb{D}(\mathsf{b}))(\exists \mathsf{c} : \mathsf{b} \prec \mathsf{c})[\mathsf{c} \in \mathsf{O}]$; (satisfiability)
(9.3) $(\forall \mathsf{b} \in \mathbb{B})(\forall \mathsf{O} \in \mathbb{D}(\mathsf{b}))[\mathsf{b} \notin \mathsf{O}]$; (non-satisfaction)
(9.4) $(\forall \mathsf{b} \in \mathbb{B})$ $(\forall \mathsf{O} \in \mathfrak{p}\mathbb{B} : \mathsf{b} \notin \mathsf{O})$ $[((\exists \mathsf{d} < \mathsf{b})$ $(\forall \mathsf{c} : \mathsf{d} < \mathsf{c} < \mathsf{b})$ $[\mathsf{O} \in \mathbb{D}(\mathsf{c})] \wedge (\exists \mathsf{a} : \mathsf{b} \prec \mathsf{a})$ $[\mathsf{a} \in \mathsf{O}])(\exists \mathsf{a} : \mathsf{b} \prec \mathsf{a})[\mathsf{a} \notin \mathsf{O}]) \Rightarrow \mathsf{O} \in \mathbb{D}(\mathsf{b})]$. (persistence)
(10) $\mathbb{E}$ (the *explicit permission function*) is a function from branches to sets of propositions, i.e. from $\mathbb{B}$ to $\mathscr{P}\mathscr{P}\mathbb{B}$.

Finally, a *model* is any octuple $\mathfrak{m} = \langle \mathbb{A}, \mathbb{B}, \mathbb{C}, \mathbb{D}, \mathbb{E}, \sim, <, \mathbb{V} \rangle$ in which the component structure $\langle \mathbb{A}, \mathbb{B}, \mathbb{C}, \mathbb{D}, \mathbb{E}, \sim, < \rangle$ is a branch deontic frame, and:

(11) $\mathbb{V}$ is a *valuation* assigning, at each moment $\mathsf{m} \in \mathbb{M}$, to each sentential constant s, a truth value $\mathbb{V}(\mathsf{m}, s) \in \{\top, \bot\}$.

The rationale behind constraints 9.1–9.4 is discussed in Brown 1996 and Brown 1999. Briefly, *non-triviality* ensures that we don't count the tautological and the inevitable among the things that are obligatory; *satisfiability* ensures that nothing counts as an obligation if it cannot eventually be fulfilled, *non-satisfaction* ensures that obligations cease once they are fulfilled, and *persistence* that once created, obligations continue until they are fulfilled or become unfulfillable.

To indicate that a formula p is *satisfied* (*true*) along a branch b in a model $\mathfrak{m}$, we write that b, $\mathfrak{m} \models p$; otherwise b, $\mathfrak{m} \not\models p$. We say that p is *valid*, or is a *validity*, iff it is satisfied along every branch in every model. For each formula p, we let $\|p\|\mathfrak{m}$ be the set of all branches in the model $\mathfrak{m}$ at which p is satisfied in $\mathfrak{m}$; i.e. $\|p\|\mathfrak{m}$ is the proposition (in $\mathfrak{m}$) expressed by p. The satisfaction conditions (within a tacitly specified model $\mathfrak{m}$) can then be given as follows, for a wide variety of useful operators:

7.2. Truth-Functional Operators

$b \models s$	iff	$\mathbb{V}(M(b), s) = \top$ (where s is a sentential constant);
$b \models \neg p$	iff	$b \not\models p$;
$b \models (p \wedge q)$	iff	$b \models p$ & $b \models q$;
$b \models (p \vee q)$	iff	$b \models p$ or $b \models q$;
$b \models (p \rightarrow q)$	iff	$b \models p \Rightarrow b \models q$;
$b \models (p \leftrightarrow q)$	iff	$b \models p \Leftrightarrow b \models q$.

7.3. Basic Temporal Operators

$b \models \Diamond p$	iff	$(\exists c : b \sim c)[c \models p]$; it is possible for p to be true (right now);
$b \models \Box p$	iff	$(\forall c : b \sim c)[c \models p]$; p is true, no matter what anyone does; p is settled true;
$b \models \overleftarrow{\Diamond} p$	iff	$(\exists a < b)[a \models p]$; it was once true in the past that p;
$b \models \overleftarrow{\Box} p$	iff	$(\forall a < b)[a \models p]$; it has always been true in the past that p;
$b \models \overrightarrow{\Diamond} p$	iff	$(\exists b : b < c)[c \models p]$; it will sometimes in the future be true that p;
$b \models \overrightarrow{\Box} p$	iff	$(\forall b : b < c)[c \models p]$; it will always be true in the future that p.
$b \models p$ since q	iff	$(\exists a < b)[[a \models q]$ & $(\forall c < b : a < c)[c \models p]]$; q was once true and p has been true ever since;
$b \models p$ until q	iff	$(\exists a > b)[[a \models q]$ & $(\forall c > b : a > c)[c \models p]]$; q will sometimes be true, and p will be true until then;

As was indicated earlier, when $\Box p$ is true, we follow Belnap and Perloff (1988) and Belnap and Perloff (1992) in saying that p is "settled true", meaning that no matter how the future goes from here, p is true relative to that future (that branch).

7.4. Basic Obligation and Permission Operators

$\mathsf{b} \models \mathcal{O}^2 p$ iff $\|p\| \in \mathbb{D}(\mathsf{b})$; p precisely expresses one current impersonal obligation;

$\mathsf{b} \models \mathcal{O}^1 p$ iff $(\exists \mathsf{O} \in \mathbb{D}(\mathsf{b}))(\forall \mathsf{c} \in \mathsf{O} : \mathsf{b} \prec \mathsf{c})[\mathsf{c} \models p]$; there is a current impersonal obligation whose fulfillment entails the truth of p;

$\mathsf{b} \models \Pi p$ iff $\|p\| \in \mathbb{E}(\mathsf{b})$; p precisely expresses one current impersonal permission.

We now have appropriate Type 1 and Type 2 impersonal and unconditional obligation operators $\mathcal{O}^1$ and $\mathcal{O}^2$, respectively.

By changing the first quantifier in the truth conditions for $\mathcal{O}^1$, we could characterize another deontic operator which, when applied to p, would have the effect of saying that each current impersonal obligation is one whose fulfillment entails the truth of p. Such an operator would be syntactically independent of the ones given here, and would be useful in a partial axiomatization of our system, but is not needed for our immediate purposes.

7.5. Basic Action Operator

$\mathsf{b} \models \delta_\alpha p$ iff $(\exists \mathsf{A} \in \mathbb{C}(\alpha, \mathsf{b}) : \mathsf{b} \in \mathsf{A})(\forall \mathsf{c} \in \mathsf{A})[\mathsf{c} \models p]$; agent α's chosen action entails the truth of p; i.e. $\text{cstit}_\alpha p$.

This operator δ_α is just the cstit_α operator mentioned earlier and studied in Chellas 1964, Brown 1990, and Horty and Belnap 1995. I choose the notation used here, rather than the "cstit_α" notation, for compactness of expression, while preserving the sense of connection to the general action operator Δ_α to be defined shortly.

7.6. Defined Temporal and Action Operators

We can now define some additional useful temporal operators:

$\bullet\!\overrightarrow{\Box} p =_{df} p \wedge \overrightarrow{\Box} p$; it is now and always will be true that p;
$\bullet\!\overrightarrow{\Diamond} p =_{df} p \vee \overrightarrow{\Diamond} p$; it is now or sometime will be true that p;
$\overleftarrow{\Box}\!\bullet p =_{df} p \wedge \overleftarrow{\Box} p$; it is now and always has been true that p;
$\overleftarrow{\Diamond}\!\bullet p =_{df} p \vee \overleftarrow{\Diamond} p$; it is now or once was true that p.

We can also define some other modalities of the logic of action:

$\mathscr{A}_\alpha p =_{df} \Diamond \delta_\alpha p$; agent α has a choice available which would assure the truth of p;

$\Delta_\alpha p \quad =_{df} \quad \delta_\alpha p \wedge \neg \Box p$; agent α's action assures the truth of p, (Horty and Belnap's positive condition) but p is not otherwise settled true (Horty and Belnap's negative condition); agent α deliberately sees to it that p;

$\mathscr{R}_\alpha p \quad =_{df} \quad \Delta_\alpha \neg \Delta_\alpha p$; agent α deliberately sees to it that s/he does not deliberately see to it that p;

$\mathscr{L}_\alpha p \quad =_{df} \quad \neg \Delta_\alpha \neg p$; agent α doesn't deliberately see to it that $\neg p$; agent α deliberately allows p to be true.

As defined here, the ability operator $\mathscr{A}_\alpha$ is a Type 1 operator, while the remaining three defined operators are of Type 2 because of the negative condition involved in the definition of the Δ_α operator. Indeed that negative condition was introduced by Horty and Belnap precisely to convert the Type 1 operator δ_α into a Type 2 operator, thus preventing it from being possible to assert that an agent deliberatively saw to the truth of tautologies, accidental truths and other propositions whose truth was similarly none of the agent's doing.

There is no need to introduce parallel definitions (by analogy with the relation between the Δ_α and δ_α operators) for new Type 1 operators of refraining and allowing, since it is demonstrable that $\delta_\alpha \neg \delta_\alpha p$, which would be used to define the Type 1 refraining operator, is logically equivalent to the simpler $\neg \delta_\alpha p$; similarly the Type 1 allowing operator would just be $\neg \delta_\alpha \neg p$.

7.7. Defined Personal Deontic Operators

We can now approach the problem of defining appropriate personal (but still unconditional) deontic modalities. We will then proceed to discuss conditional obligation in Section 8.

$\mathscr{O}^1_\alpha p \quad =_{df} \quad \mathscr{O}^1 \delta_\alpha p$; agent α has an obligation, every means of fulfilling which would require acting in a way that would assure the truth of p; more precisely: there is an obligation that agent α act in a way that will entail the truth of p;

$\mathscr{O}^2_\alpha p \quad =_{df} \quad \mathscr{O}^2 \Delta_\alpha p$; agent α has an obligation to see to it that p; more precisely: there is an obligation that agent α see to it that p;

$\mathscr{F}^1_\alpha p \quad =_{df} \quad \mathscr{O}^1 \neg \delta_\alpha p$; agent α has an obligation, every means of fulfilling which would require not choosing any action which would entail that p; more precisely: there is an obligation that agent α not act in any way that will entail the truth of p;

$\mathscr{F}^2_\alpha p \quad =_{df} \quad \mathscr{O}^2 \mathscr{R}_\alpha p$; agent α has an obligation to refrain from seeing to it that p; more precisely: there is an obligation that agent α refrain from seeing to it that p;

$\mathscr{P}^1_\alpha p \quad =_{df} \quad \neg \mathscr{F}^1_\alpha p$; i.e. $\neg \mathscr{O}^1 \neg \delta_\alpha p$; every obligation agent α has is one whose fulfillment is compossible with choosing an action which would entail that p; more precisely: there is no obligation that agent α not act in any way that will entail the truth of p;

$\mathscr{P}^2_\alpha p \quad =_{df} \quad \neg \mathscr{F}^2_\alpha p$; agent α is not obligated to refrain from seeing to it that p; more precisely: there is no obligation that agent α refrain from seeing to it that p.

It is worth reminding ourselves here that the basic notions of unconditional obligation provided thus far are appropriate only to the discussion of simply dischargeable obligations. The various senses of *forbidden* in this paper therefore inherit this specific character, since they are defined in terms of obligation. But the concepts of permission just introduced as negations of specific obligations, are general in character: to be free of any specific prohibition against bringing it about that p is to enjoy a general permission to bring it about.

For our treatment of explicit permission, however, we will need to set explicit permissions against a background of general prohibitions. I will suggest that such general prohibitions can best be understood using the concept of a general conditional obligation. So let us now turn to the subject of conditional obligation.

8. Conditional Obligation

With these concepts in place, let us attempt to define conditional obligation, keeping two considerations in mind:

(1) Expressions of conditional obligation are normally general, intended to cover any possible situation (within some tacitly given temporal range) in which the triggering condition may be met.
(2) We cannot have unsatisfiable obligations.

We most often have occasion to note this second point in connection with obligations which have for a time been in force but which for one reason or another become unsatisfiable: when I return the book I borrowed, my obligation becomes fulfilled, is then no longer fulfillable, and ceases; when someone else helps the accident victim before I can, I am relieved of my now unfulfillable obligation; when my friend dies before I can apologize, my obligation ceases; when I become bankrupt, my debts are canceled; etc. But there is another aspect to the impossibility of unsatisfiable obligations: an obligation cannot come into force if there is no chance of satisfying it. This is urged in the slogan "*ought* implies *can*" (to be distinguished from the slogan "*ought* implies *may*"). In connection with conditional obligation, this has the effect that it would be inaccurate to assert that there was a conditional obligation $\mathcal{O}(p/q)$ if we considered that the truth of q was the sole condition for triggering the obligation in question, and if it could happen that the trigger q should come to be true in circumstances in which the resulting "obligation" $\mathcal{O}p$ was unsatisfiable. More specifically, this means that in any true claim of the form $\mathcal{O}(p/q)$, the actual condition under which an obligation that p is created must be the double condition

1. that q be true *and*
2. that it be possible for any obligation that p to become fulfilled.

I call the first conjunct the *trigger*, or the *triggering condition*, of the conditional obligation, and call the second conjunct the *satisfiability condition*.

With these considerations in mind, let us look first at personal conditional obligation. I would argue that there is a use for each of several notions of conditional obligation that we can define consistently with these conditions just stated, and that on different occasions similar-looking sentences expressing conditional obligation in natural language may well require different renderings into our formal language, calling upon its full expressive power.

When I say that

If I'm going, I ought to let Sue know

this may be a very circumstantial claim, not intended to generalize in any obvious way to future or past occasions, but nonetheless intended to apply no matter what I now decide about going. At the same time, it may well be intended that letting Sue know should be the whole content of the obligation that would arise, so that Type 2 obligation is apt. In such a case the claim may be properly translated simply as of the form:

$$\mathcal{O}^2_\alpha(p/q) =_{df} \Box((q \wedge \Diamond \vec{\Diamond} \Delta_\alpha p) \rightarrow \mathcal{O}^2_\alpha p).$$

Here the generality involved is, crudely speaking, generalization over all currently available choices. The extra conjunct $\Diamond \vec{\Diamond} \Delta_\alpha p$ in the antecedent of the conditional is the satisfiability condition. It expresses the constraint—normally presupposed and left tacit—that the obligation which might eventuate would be satisfied at some future point in some currently available branch, and therefore is satisfiable. Suppose, for example, that the choices open to me are to go by train, or by plane, or by bus, or to refrain from going at all, and suppose that if I go by bus, there will be neither time nor opportunity to let Sue know. In that case, I will be obligated to let Sue know if I go by plane or by train, but not otherwise. We might, of course, say that even if I go by bus I *ought* to let Sue know, though that's impossible. I suggest that such uses of "ought" represent a comparison to some more utopian situation—"it would have been better if it had been possible to …". Utopian *ought*s, if we may call them that, no doubt play an important role in our normative reasoning, perhaps by providing ideals which can serve as a moral compass to help us navigate through the all too often disorienting moral terrain in which we find ourselves. But I suggest they do not express implacable current obligations for whose non-fulfillment we can be blamed even when they are unfulfillable.

When I say instead that

If I'm going I should let somebody know

I may be expressing a similarly circumstantial claim, but it is unlikely that letting somebody know is intended to be a sufficient condition for the fulfillment of the associated obligation. Letting a total stranger know will probably not count. In that case a Type 1 formulation is needed:

$$\mathcal{O}^1_\alpha(p/q) =_{df} \Box((q \wedge \Diamond \overrightarrow{\Diamond} \delta_\alpha p) \to \mathcal{O}^1_\alpha p).$$

Commonly, however, we might say something like

If you ever get to Syracuse, you should look me up.

Such a statement is intended to apply throughout the future, but probably is not intended to apply to the past. Since we just met, I can hardly suggest that when you were in Syracuse last year you should have looked me up. At the same time, it may well be that looking me up exactly captures the nature of what I wish to say you ought to do if you come to Syracuse. In that case, we need a Type 2 formulation like this:

$$\overrightarrow{\mathcal{O}}^2_\alpha(p/q) =_{df} \Box^\bullet \overrightarrow{\Box}((q \wedge \Diamond \overrightarrow{\Diamond} \delta_\alpha p) \to \mathcal{O}^2_\alpha p).$$

A similarly future-oriented, but Type 1, claim might be intended by a sentence such as:

If you ever get to Syracuse, you should come to the University,

where the real intent may be to say that you should come to a particular department, for a non-trivial stretch of time, and make yourself known to one of a limited number of key people there, etc. Coming to the University will be one salient aspect of fulfilling the obligation in question, but will not itself constitute fulfilling it. Accordingly, we may need a Type 1 rendering, as follows:

$$\overrightarrow{\mathcal{O}}^1_\alpha(p/q) =_{df} \Box^\bullet \overrightarrow{\Box}((q \wedge \Diamond \overrightarrow{\Diamond} \delta_\alpha p) \to \mathcal{O}^1_\alpha p).$$

We can also have occasion to say something like this:

When your partner dies, you should go to the funeral

which is evidently intended to cover past, present, and future contingencies, and thus might properly be translated as of the form

$$\overleftrightarrow{\mathcal{O}}^2_\alpha(p/q) =_{df} \overleftarrow{\Box}^\bullet \Box^\bullet \overrightarrow{\Box}((q \wedge \Diamond \overrightarrow{\Diamond} \Delta_\alpha p) \rightarrow \mathcal{O}^2_\alpha p).$$

A less specific statement, such as

When your partner dies, you should do something about it

might be rendered with a corresponding formula of Type 1:

$$\overleftrightarrow{\mathcal{O}}^1_\alpha(p/q) =_{df} \overleftarrow{\Box}^\bullet \Box^\bullet \overrightarrow{\Box}((q \wedge \Diamond \overrightarrow{\Diamond} \delta_\alpha p) \rightarrow \mathcal{O}^1_\alpha p).$$

Since, on the account given here, impersonal conditional obligation is involved in any personal conditional obligation, these same examples illustrate the uses for various sorts of expressions of impersonal conditional obligation:

$$\mathcal{O}^2(p/q) =_{df} \Box((q \wedge \Diamond \overrightarrow{\Diamond} p) \rightarrow \mathcal{O}^2 p);$$
$$\mathcal{O}^1(p/q) =_{df} \Box((q \wedge \Diamond \overrightarrow{\Diamond} p) \rightarrow \mathcal{O}^1 p);$$
$$\overrightarrow{\mathcal{O}}^2(p/q) =_{df} \Box^\bullet \overrightarrow{\Box}((q \wedge \Diamond \overrightarrow{\Diamond} p) \rightarrow \mathcal{O}^2 p);$$
$$\overrightarrow{\mathcal{O}}^1(p/q) =_{df} \Box^\bullet \overrightarrow{\Box}((q \wedge \Diamond \overrightarrow{\Diamond} p) \rightarrow \mathcal{O}^1 p);$$
$$\overleftrightarrow{\mathcal{O}}^2(p/q) =_{df} \overleftarrow{\Box}^\bullet \Box^\bullet \overrightarrow{\Box}((q \wedge \Diamond \overrightarrow{\Diamond} p) \rightarrow \mathcal{O}^2 p);$$
$$\overleftrightarrow{\mathcal{O}}^1(p/q) =_{df} \overleftarrow{\Box}^\bullet \Box^\bullet \overrightarrow{\Box}((q \wedge \Diamond \overrightarrow{\Diamond} p) \rightarrow \mathcal{O}^1 p);$$

Anderson (1959) held that a number of rules and schemata should be valid in any adequate account of conditional obligation. Among them, and included by Åqvist (1984) as his C8, is the rule of Material Detachment. But this rule is invalidated by all the accounts given here. This might seem to be a misfortune, since it would at first seem that the whole point of a conditional obligation $\mathcal{O}(p/q)$ would be lost if we cannot depend on it to generate a (no longer merely conditional) obligation $\mathcal{O}(p)$ when its triggering condition q is met. And that is nearly right; it leaves out only the detail that the unconditional obligation in question must be *possible* in a given set of circumstances in order for those circumstances to trigger it. More specifically, the proposed unconditional obligation must be *satisfiable* in the circumstances.

So a slightly revised version of the Principle of Material Detachment is valid, which says that from $\mathcal{O}(p/q)$ and q we may infer that $\mathcal{O}p$, *provided* that it is possible to satisfy $\mathcal{O}p$. We might call this the rule of Qualified Material Detachment.

The account of conditional obligation given here has been shown in Brown (forthcoming) to handle well a number of benchmark examples which have posed problems for other accounts of conditional obligation.

In particular, because none of the versions of conditional obligation given here support a rule of deontic detachment, none of them fall prey to Chisholm's "paradox", first discussed in Chisholm 1964. Chisholm's example concerns the consequences to be derived in a situation in which:

I ought to visit my mother.
If I'm going to visit my mother,
    I ought to tell my mother I'm coming.
If I'm not going to visit my mother,
    I ought to tell my mother I'm not coming.
I don't visit her.

The last premiss is contrary-to-duty, relative to the first. If we accept deontic detachment, the first two premises yield an obligation to tell my mother I'm coming. If we also accept material detachment, then the last two premises yield an obligation to tell my mother I'm not coming. Together, then, the four premises would yield conflicting obligations of a sort we feel intuitively should not arise in this situation.

In fact, our account of conditional obligation doesn't support either deontic or material detachment, and therefore doesn't give rise to the unintuitive conflict. The fact that it doesn't support material detachment would not by itself solve the problem, since we could easily strengthen the problem by adding a fifth, entirely plausible, premise:

It is possible to tell my mother I'm coming.

With this additional premise in place, the material detachment would go through. The deontic detachment would still be blocked, however, and no plausible premise could be added that would make that inference go through.

It has been urged by Prakken and Sergot, in Prakken and Sergot 1996, that we should also be concerned with another aspect of Chisholm's "paradox", which they call the "pragmatic oddity". This is the oddity of being left in a position of having to acknowledge both that I ought to visit my mother and that I ought to tell her I'm not coming. With the fifth premise in place, that can happen. However, I think it is entirely correct to label this a *pragmatic* oddity. This "oddity" disappears entirely if we add the appropriate pragmatic cues to avoid inappropriate conversational implicatures:

I really ought to visit my mother but,
given that I'm *not* going to visit her,
I ought to tell her I'm not coming.

9. Prohibition and Explicit Permission

As we prepare to consider general prohibitions and explicit permissions, it is worth emphasizing two aspects of the account just given of conditional obligations. First, they involve a satisfiability condition that is normally tacit in their natural-language expressions. Second, the conditional obligations are general, though when triggered and satisfiable they give rise to specific obligations.

I suggest that general prohibitions also involve an element normally left tacit in their natural-language expressions, and more specifically a condition: the condition that no countervailing explicit permission be in force. Thus I suggest that the prohibition that might normally be expressed by saying

> You ought not to drive through an intersection
> while the traffic light is red

can be understood as involving a tacit *unless*, as follows:

> You ought not to drive through an intersection
> while the traffic light is red
> unless explicitly permitted to do so
> (by a duly constituted authority).

Formally, a typical expression of general prohibition can be expected to look something like the following, using our Type 2 impersonal permission operator Π introduced in section 7.4:

$$\mathcal{O}^2(\mathcal{R}_\alpha p / \neg \Pi \Delta_\alpha p)$$
$$\text{i.e. } \Box((\neg \Pi \Delta_\alpha p \wedge \Diamond \vec{\Diamond} \mathcal{R}_\alpha p) \rightarrow \mathcal{O}^2 \mathcal{R}_\alpha p).$$

Because the introduction of this condition invokes our treatment of conditional obligation, it automatically brings with it the needed element of generality.

Now consider the relationship between a prohibition and a corresponding explicit permission, as thus construed. The permission is not redundant, since it stands in relation to a relevant prohibition whose applicability it influences, nor does it contradict the prohibition, since the prohibition has in effect anticipated the possibility of such exceptions, and has provided a place for them.

Variants on this formulation will be appropriate, depending on the intention involved, as in the case of conditional obligation. In particular, the generality involved might extend into the future or into the future and the past, and might quite easily also include generality with respect to agent. This last introduces a new complexity: it represents the first time we have seriously entertained including overt quantification, or its equivalent, within our formal language. For convenience, we shall use a schematic variable ϕ over agents to effect this generality here, giving rise to formulations like this:

$$\mathcal{O}^2(\mathcal{R}_\phi p / \neg \Pi \Delta_\phi p).$$

Note that the formulations given have all been of Type 2. So far, there appears to be no need to provide corresponding Type 1 formulations. Explicit permission, precisely because it is explicitly given, would appear to be by its nature of Type 2, and has been so treated in the satisfaction conditions. It then appears that if a general prohibition is to be understood in a spirit which permits explicitly granted exceptions, it too must be of Type 2 in order for the exceptions to be exceptions to it rather than to some other prohibition. Perhaps this reasoning will prove over-hasty after we have had time to digest and treat various examples. If so, Type 1 correlates could easily be provided.

10. Material for Future Investigation

We generally expect some symmetry between permission and prohibition. This naturally prompts the question as to whether our account of explicit permission against a background of general prohibition should be paralleled by an account of explicit prohibition against a background of general permission.

At first glance this seems reasonable and one would expect it to be easy to do. We should just define conditional permission, provide various sorts of generality for different conditional permission operators, and build in an explicit escape clause which invalidates the general permission in the face of specific prohibitions.

The problem is, however, more complex than at first appears. The parallel between permission and prohibition is marred by the fact that, at least in the treatment offered thus far, explicit permission and explicit obligation are provided semantically by primitive functions $\mathbb{E}$ and $\mathbb{D}$, but prohibition is taken as a more complex defined notion. Moreover, as treated here, obligation and permission are, at bottom, impersonal, while prohibition has been treated in a way that involves the agentive notion of refraining. We can make some sense of the notion of an impersonal *ought*, and indeed our analysis in Section 3 seems to force it upon us. But it is difficult to see what could be meant by an impersonal *prohibited* other than just *ought not*. While the Type 1 account of personal prohibition does apply a Type 1 impersonal *ought not* to a Type 1 outcome expression, and is thus parallel to the treatment of Type 1 personal explicit permission, no such parallel works for Type 2 expressions. The Type 2 impersonal *ought not* is not involved in the definition of Type 2 personal prohibition.

It now appears that the symmetry between permission and prohibition is not as great as a casual glance would have suggested. Further investigation of these matters will be needed.

References

Anderson, A. R. 1959. On the logic of commitment. *Philosophical Studies*, vol. 10, pp. 23–27.

Anderson, A. R. 1967. The formal analysis of normative systems. In Nicholas Rescher (ed.), *The Logic of Decision and Action*, pp. 147–213. University of Pittsburgh Press, Pittsburgh.

Åqvist, L. 1984. Deontic logic. In D. Gabbay and F. Guenthner (eds.), *Extensions of Classical Logic*, vol. II of *Handbook of Philosophical Logic*, pp. 605–714. D. Reidel, Dordrecht.

Belnap, Jr, N. D. and Bartha, P. 1995. Marcus and the problem of nested deontic modalities. In W. Sinnott-Armstrong, D. Raffman, and N. Asher (eds.), *Morality and Belief: A Festschrift in Honour of Ruth Barcan Marcus*, pp. 174–197. Cambridge University Press, New York.

Belnap, Jr, N. D. and Perloff, M. 1988. Seeing to it that: A canonical form for agentives. *Theoria*, vol. 54, pp. 175–199.

Belnap, Jr, N. D. and Perloff, M. 1992. The way of the agent. *Studia Logica*, vol. 51, pp. 463–484.

Brown, M. A. 1990. Action and ability. *Journal of Philosophical Logic*, vol. 19, 95–114.

Brown, M. A. 1996. Doing as we ought: Towards a logic of simply dischargeable obligations. In M. A. Brown and J. Carmo (eds.), *Deontic Logic, Agency and Normative Systems: ΔEON '96*, pp. 47–65. Springer-Verlag, Berlin.

Brown, M. A. 1999. Agents with changing and conflicting commitments: A preliminary study. In P. McNamara and H. Prakken (eds.), *Norms, Logics and Information Systems: New Studies in Deontic Logic and Computer Science*, vol. 49 of *Frontiers in Artificial Intelligence and Applications*, pp. 109–125. IOS Press, Amsterdam.

Brown, M. A. forthcoming. Conditional and unconditional obligation for agents in time. In M. de Rijke, K. Segerberg, H. Wansing, and M. Zacharyaschev (eds.), *Advances in Modal Logic II*. CSLI, Stanford.

Brown, M. A. and Goranko, V. 1999. An extended branching-time ockhamist temporal logic. *Journal of Logic, Language and Information*, vol. 8, pp. 143–166.

Chellas, B. F. 1964. *The Logical Form of Imperatives*. Perry Lane Press, Palo Alto.

Chellas, B. F. 1980. *Modal Logic: An Introduction*. Cambridge University Press, Cambridge.

Chisholm, R. 1964. The ethics of requirement. *American Philosophical Quarterly*, vol. 1, pp. 147–153.

García, J. 1986. The *tunsollen*, the *seinsollen*, and the *soseinsollen*. *American Philosophical Quarterly*, vol. 23, pp. 267–276.

Horty, J. F. 1989. An alternative stit-operator. Tech. rep., University of Maryland.

Horty, J. F. 1996. Combining agency and obligation (preliminary version). In M. A. Brown and J. Carmo (eds.), *Deontic Logic, Agency and Normative Systems: ΔEON '96*, pp. 98–122. Springer-Verlag, Berlin.

Horty, J. F. and Belnap, N. D. 1995. The deliberative stit: A study of action, omission, ability, and obligation. *Journal of Philosophical Logic*, vol. 24, 583–644.

Jones, A. J. I. and Pörn, I. 1985. Ideality, sub-ideality and deontic logic. *Synthese*, vol. 65, pp. 275–290.

Kanger, S. 1971. New foundations for ethical theory. In R. Hilpinen (ed.), *Deontic Logic: Introductory and Systematic Readings*, vol. 33 of *Synthese Library*, pp. 36–58. D. Reidel, Dordrecht. Reprinted from a 1957 technical report, University of Stockholm.

Krogh, C. and Herrestad, H. 1996. Getting personal: Some notes on the relationship between personal and impersonal obligation. In M. A. Brown and J. Carmo (eds.), *Deontic Logic, Agency and Normative Systems; ΔEON '96*, pp. 134–153. Springer-Verlag, Berlin.

von Kutschera, F. 1986. Bewirken. *Erkenntnis*, vol. 24, pp. 253–281.

Makinson, D. 1986. On the formal representation of rights relations. *Journal of Philosophical Logic*, vol. 15, pp. 403–425.

Makinson, D. 1999. On a fundamental problem of deontic logic. In P. McNamara and H. Prakken (eds.), *Norms, Logics and Information Systems: New Studies in Deontic Logic and Computer Science*, vol. 49 of *Frontiers in Artificial Intelligence and Applications*, pp. 29–53. IOS Press, Amsterdam.

Nute, D. (ed.) 1997. *Defeasible Deontic Logic*. Kluwer Synthese Library, Dordrecht.

Pörn, I. 1970. *The Logic of Power*. Basil Blackwell, Oxford.

Prakken, H. and Sergot, M. J. 1996. Contrary-to-duty obligations. *Studia Logica*, vol. 57, no. 1, pp. 91–115.

Ross, A. 1941. Imperatives and logic. *Theoria*, vol. 7, pp. 53–71.

Segerberg, K. 1971. *An Essay in Classical Modal Logic*. University of Uppsala, Uppsala.

Sergot, M. J. 1999. Normative positions. In P. F. McNamara and H. Prakken (eds.), *Norms, Logics and Information Systems: New Studies in Deontic Logic and Computer Science*, vol. 49 of *Frontiers in Artificial Intelligence and Applications*, pp. 289–308. IOS Press, Amsterdam.

Thomason, R. H. 1970. Indeterminist time and truth value gaps. *Theoria*, vol. 36, pp. 264–281.

Thomason, R. H. 1981. Deontic logic as founded on tense logic. In R. Hilpinen (ed.), *New Studies in Deontic Logic*, vol. 152 of *Synthese Library*, pp. 165–176. D. Reidel, Dordrecht.

von Wright, G. H. 1956. A note on deontic logic and derived obligation. *Mind*, vol. 65, pp. 507–509.

Zanardo, A. 1985. A finite axiomatization of the set of strongly valid ockhamist formulas. *Journal of Philosophical Logic*, vol. 14, pp. 447–468.

Zanardo, A. 1991. A complete deductive system for since-until branching-time logic. *Journal of Philosophical Logic*, vol. 20, pp. 131–148.

Zanardo, A. 1996. Branching-time logic with quantification over branches: The point of view of modal logic. *Journal of Symbolic Logic*, vol. 61, pp. 1–39.

Department of Philosophy,
Syracuse University,
Syracuse NY, USA
E-mail: mabrown@syr.edu