Compliance and Command II, Imperatives and Deontics

In this part of the paper, I am interested in providing a semantics and logic for deontic sentences and working out their connection with the previous semantics and logic for imperatives.

The standard approach to deontic logic is in terms of possible worlds. It is supposed that with each world is associated a set of ideal worlds. A statement O(A) to the effect tat A is obligatory is then taken to be true if A is true in all ideal worlds and a statement P(A) to the effect that A is permissible is taken to be true if A is true in some ideal world.

I believe the possible worlds approach to be fundamentally misguided. The main problem is that obligation and permission relate most directly to action. At the end of the day what we want to know is what it is obligatory or permissible *to do*. But the intensional treatment of the embedded clause A prevents the deontic statements from providing, in this way, a guide to action. For A merely represents the possible outcomes of some action or actions. But as I have argued in Fine [2014], given a set of outcomes, there is, in general, no satisfactory way to determine the actions from which they arose.

The present approach, by contrast, is action- rather than outcome-oriented. The deontic operators apply directly to expressions that indicate a range of possible actions rather than a range of possible outcomes or worlds. Indeed, under the semantics for imperatives in part I, an imperative will indicate a range of actions, those in compliance with the imperative; and so we may take the deontic operators to have direct application to imperatives. To say when a deontic sentence is true we must therefore distinguish, not a preferred set of worlds but a preferred set of actions, something I call a 'code of conduct'; and the conditions under which a deontic sentence is true relative to a code of conduct will then be somewhat different from the conditions under which such a sentence is true relative to a sphere of worlds.

The plan of the paper is as follows. I begin by making some distinctions and stipulations which will be useful in the rest of the paper (§1); I introduce and explain the key notion of a code of conduct, relative to which deontic formulas are to be be interpreted (§2); I give the clauses for when a deontic formula is true or false relative to a code of conduct (§3) and spell out some of the consequences of these clauses, especially in regard to the contrast with the standard possible worlds semantics for deontic logic (§4); I consider various ways of reformulating the criterion of validity for deontic formulas and point, in particular, to a very close connection between this criterion and the criterion of validity for imperative inference proposed in part I (§5); I consider some of the characteristic inferences that are or fail to be valid (§6) and outline a system of deontic logic within the truthmaker approach (§7); I show how one might deal with the problem of deontic updating within the truthmaker framework (§8); and I conclude with a brief formal appendix.

I assume the reader is familiar with the basic material from part I, including the truthmaker semantics for imperatives and the definition of validity for imperative inference; and it would also be helpful for her to have some knowledge of the standard possible worlds semantics for deontic logic.

§1 Preliminaries

Let me begin with some remarks on syntax. The reader will recall that imperative formulas are constructed from imperative atoms α_1, α_2, ... using the usual array of connectives - $\vee$, $\wedge$ and $\neg$ - and also the verum constant $\top$. We will now take the deontic operators O and P, for obligation and permission, to have application to imperatives. Thus when X is an imperative, O(X) and P(X) will be deontic sentences. Thus if X is the imperative 'stop', then O(X) might be taken to be the sentence 'It is obligatory to stop' and P(X) to be the sentence 'it is permitted to stop'.

This account of the logical grammar of the deontic formulas O(X) and P(X) is not meant to have any implications for the correct grammatical analysis of the corresponding sentences of ordinary language. One might, of course, see the sentences 'it is obligatory to go' (or 'you ought to go') and 'it is permitted to go' (or 'you may go') as containing the imperative 'go'. But there is no need to regard them in this way and many other grammatical analyses of these sentences are possible. The present account of the logical grammar merely allows us to relate imperative sentences to deontic sentences in a natural and straightforward way.

However, one aspect of our account is of more significance. For the deontic operators cannot sensibly be taken to apply to arbitrary indicative sentences, as in standard deontic logic. Thus we have no use for O(S) or for P(S) when S is an arbitrary indicative sentence. Indeed, the sentences 'it is obligatory to it rains' or 'it is permitted to it rains' are not even grammatical. Only actions, on the present view, can properly be said to obligatory or permitted; and if one wants to apply the deontic operators to arbitrary indicative sentences, then it must be by way of some transformation of these sentences into terms or sentences for actions, as in 'it is obligatory to see to it that it rains'.

This has the consequence that, in contrast to standard systems of deontic logic, iterative statements of obligation and permission, such as O(O(X)) or P(O(X)) or P(P(X)), will not be well-formed, since the operators O and P can only properly be taken to have application to imperative sentences. However, we may still form truth-function compounds of deontic sentences, as in $O(X) \vee O(Y)$ or $O(X) \supset P(X)$, just as we may form truth-functional compounds of imperative sentences.

There are a number of (more or less familiar) ways to interpret deontic sentences, the distinction between which will be important in what follows.

There is, in the first place, the distinction between the performative and descriptive interpretation of these locutions. Under the performative interpretation, 'you ought to go' is used to *place* one under an obligation and 'you may go' is used to *grant* a permission whereas, under the descriptive interpretation, 'you ought to go' is merely used to state that one is under an obligation and 'you may go' is used merely to state that one is so permitted (the descriptive use goes more naturally with the past tense, as in 'you were permitted (obliged) to go'). For the most part, I shall be concerned with the descriptive use of these sentences but, in the section on updates, I shall also be interested in their performative use, in which they are more akin to imperatives. I shall also assume, pace the emotivists and some expressivists, that there is a straightforward sense in which the deontic sentences, under their descriptive use, are capable of being true or false.

There is, in the second place, a familiar distinction between a strict and weak sense of permission. Roughly speaking, an action is permitted in the weak sense if it is not forbidden, whereas something more is required for an action to be permitted in the strict sense. It must have been *singled out* as being permitted - by being expressly permitted or perhaps in some other, less direct, way. The semantics I provide will initially be for the strict sense of permission, although I shall later extend it to the weak sense.

I shall assume that if a compound action, such as turning on the gas and lighting the stove, is permitted then so are the component actions, of turning on the gas and of lighting the stove. There is perhaps a sense in which one is not permitted to turn on the gas, since one is not permitted to turn on the gas without doing anything else, which is to say that one is obliged to do something else, viz. light the stove, if one turns on the gas. But we shall find it convenient to use permission in the less restrictive sense, so that an action can be permitted even though its performance requires one to do something else.

Finally, there is also a distinction, in a somewhat different sense, between a strict and weak sense of obligation. Suppose one is obliged to shut the door. Then is one also obliged to shut the door or burn the house down? There is a sense in which one is. This is a sense in which it is consistent to say that one is obliged to shut the door or burn the house down but not permitted to burn the house down. There would also appear to be a sense in which one is not. This is a sense in which one's being obliged to shut the door or burn the house down implies that one is permitted to burn the house down. I call the first sense of obligation *bounded* and the second *unbounded* or *free*. In what follows I shall be interested in giving a semantics for both sense of obligation.

§2 Codes of Conduct

We shall take a deontic statement to be true or false relative to a code of conduct. For our purposes, we may take a code of conduct to be a prescription, i.e. a set of actions. But which actions?

There are two relevant conditions. The first is that each action should be *permissible* according to the code. The second is that each action should be *adequate*, i.e. the performance of the action should be sufficient to discharge all of one's obligations in regard to the code. The first condition (Permissibility) means that the action should not contain too much, it should not encroach on what is impermissible, and the second (Adequacy) means that the action should not contain too little, it should not fall short of what is obligatory. Thus an action meeting both conditions will conform to what one might call the 'Goldilock's principle', it will strike a middle course between containing enough and not containing too much (although perhaps Aristotle should be given credit for the principle).

An action meeting both conditions will be said to be *ideal*. An ideal action is the analogue of a deontic alternative in the possible worlds semantics for deontic logic; for a deontic alternative is an ideal world, one in which all obligations are discharged. However, for us, there is no requirement that an ideal action should be a complete action, let along a complete state of the world, and there is not even any requirement that it should be consistent. Moreover, for us, the two conditions (of Permissibility and Adequacy) are independent whereas, in the standard modal setting, a world will be permissible just in case it is a world in which all one's obligations are discharged.

We might say that a code *sanctions* those actions which are ideal with respect to the code;

and a code of conduct may be identified with the set of the actions that it sanctions. The actions sanctioned by a code will often be an action-stream, composed of many individual actions, and, for this reason, we might talk instead of a ‘ course of action’. However, the space of actions is closed under fusion; and so, strictly speaking, a course of action is just another action.

A code of conduct is the analogue of the ‘sphere’ of deontic alternatives within the standard possible worlds semantics for deontic logic. Indeed, with each code of conduct may be associated a deontic sphere, consisting of all those worlds which are compatible with some course of action sanctioned by the code. If we think of each course of action sanctioned by the code in terms of its intensional content, i.e. the set of worlds containing the action, then the associated deontic sphere will be the union or ‘disjunction’ of all these contents.

However, the correspondence is far from being one-one. Thus one code might consist of my eating an apple and another of my eating an apple and of my eating an apple and a pear. These are distinct codes, since the second sanctions my eating an apple and a pear while the first does not, but the corresponding deontic spheres are the same. Thus codes of conduct are much more fine-grained than the corresponding deontic spheres and can be expected - at least, in principle - to deliver different results about what is obligatory or permissible.

Codes of conduct are also loosely related the “ to-do lists”, championed by Portner ([2005]) and others. But, as will become clear, we have a rather different view of their nature and semantic role.

There are a number of different ‘consistency’ conditions one might wish to impose on codes of conduct:

<u>Non-Emptiness</u> Each code should sanction at least one action.

If this were not so then there would be no action (even impossible or necessary) sanctioned by the code. This condition does not mean that the code must have any real substance since it is met by the ‘minimal’ code {□}, under which it is guaranteed that one will do what is permissible (via the performance of the null action □) and thereby discharge all of one’s obligations.

Given that a code of conduct is non-empty, we might also want to restrict its content:

<u>Non-Anarchy</u> The full action ■ is not sanctioned by any conduct.

If this condition were violated it would then mean that everything whatever (even the impossible) was permitted.

<u>Consistency</u> Each code should sanction at least one possible conduct.

If this condition were violated then the code of conduct would either be empty or consist entirely of impossible courses of action.

A stronger condition still (in the presence of Non-Emptiness) is:

<u>Complete Consistency</u> Each action sanctioned by a code is possible.

It has sometimes been thought that consistency constraints of this sort are normative rather than logical in character. This may be so. But their failure will restrict the capability of a code to serve as a guide to conduct. If a code sanctions no action, then it can provide no guidance at all; if a code sanctions only impossible actions, then it can only serve as a guide in so far as one is capable of performing some consistent part of a action that it sanctions; and if a code sanctions some impossible action, then one is no longer capable of performing all of the action that it sanctions. It is only when the conditions of Non-emptiness and Complete Consistency are both satisfied, that the code can straightforwardly serve as a guide to conduct.

Say that an action *b lies between* two others, *a* and *c*, if *a* is a part of *b* and *b* a part of *c*. Then the fact that codes of conduct conform to the Goldilock's principle means that:

Convexity Any action that lies between two courses of action sanctioned by a code is also sanctioned by the code.

For suppose that both *a* and *c* are ideal, i.e. permissible and adequate. Then *b* is permissible since it is a part of a permissible action *c*; and *b* is adequate since it contains an adequate action *a*.

It might be thought to be implausible that all codes of conduct should conform to Convexity. Suppose, for example, that the code sanctions pressing button A and also pressing buttons A, B and C. Then why should it also sanction pressing button A and B? But we may appeal here to the argument used in §5 of Part I to justify the appeal to reverse entailment. For if pressing buttons A and B is not ideal, then that can only be because it is not pressing button A that is ideal, but pressing A without B, in which case the convexity condition will have no application.

Say that a (course of) action is *complete* if it is possible and every action is either incompatible with it or part of it. Thus a complete course of action is the action-theoretic counterpart of a possible world. A further closure condition one might then wish to impose is:

Completeness Every possible action sanctioned by a code is included in complete course of action sanctioned by the code.

This condition is plausible for weak permission; if the action *c* is implicitly permitted then some completion of the action must also be implicitly permitted. However, it has no plausibility for strict (explicit) permission. Suppose that the null action *a* is sanctioned. Then this tells us nothing about whether shutting the door or leaving it open is strictly permitted.

Let us note, finally, that we might define a natural relation of part-whole among codes. For we may suppose that one code is a part of another if it analytically entails the other, i.e. if every action sanctioned by the first contains a action sanctioned by the second and every action sanctioned by the second is contained in a action sanctioned by the first. It is readily shown, given that codes are convex sets, that the relation of part-whole is antisymmetric; and it is also readily shown that the fusion of any codes will also be a code. Thus the space of codes will form a state space and, as we shall see, the codes may themselves be regarded as truthmakers or falsemakers for deontic statements.

§3 Truthmaker Semantics for Deontic Logic

We shall provide a truthmaker semantics for deontic statements. The atomic deontic statements will be of the form O(X) or P(X), for X an imperative. Recall from part I that the prescriptive content *X* of X is the set of actions in compliance with X. For the purpose of providing a semantics, we suppose given a code of conduct *C* and then go on to specify when O(X) or P(X) is true in terms of an appropriate relation between the content *X* and the code *C*.

Recall the notions of subsumption and subservience from part I. Given prescriptive contents *X* and *Y*, *X* will subsume *Y* if every action in compliance with *X* contains an action in compliance with *Y* and *Y* will subserve *X* if every action in compliance with *Y* is contained in an action in compliance with *X*. I would now like to suggest the following account for when statements of obligation or permission are true:

(i) O(X) is true iff *C* subsumes *X*;
(ii) P(X) is true iff *X* subserves *C*.

Or to state the clauses explicitly:

(i)′ O(X) is true iff every ideal course of action contains an action in compliance with X;
(ii)′ P(X) is true if every course of action in compliance with X is contained in an ideal action.

These clauses may be extended to truth-functional compounds of atomic deontic statements in the standard (classical) way; and we might then say that the deontic statements S_1, S_2, ... (*classically*) *entail* the deontic statement T if T is true whenever (i.e. in any model in which) S_1, S_2, ... are true.

We may informally justify each of these clauses above; and, to this end, let us assume, since it is not currently in contention, that an imperative X will be permitted just in case each of the actions in compliance with X is permitted.

First, in regard to the right-to-left direction of (i), suppose that *C* subsumes *X*, i.e. that each ideal action c_1, c_2, ... contains an action a_1, a_2, ... in compliance with X. Now surely one is obliged to discharge one's obligations in a permissible manner, i.e. one is obliged to perform some one of c_1, c_2, ... But then surely one is obliged to perform some one of their parts a_1, a_2, ...

Second, for the other direction, suppose that X is obligatory. Consider any ideal action *c*. Thus one may discharge one's obligations by performing c. Suppose now that no action in compliance with X is a part of *c*. How then can X be obligatory, given that the performance of an action in compliance with X will go no way towards discharging this particular way of discharging one's obligations?

Third, in regard to the right-to-left direction of (ii), suppose *X* subserves *C*, i.e. that every action *a* in compliance with X is part of an ideal action *c*. Then since *c* is permissible, so is *a* (under the loose understanding of permission which we have adopted).

Finally, for the other direction, suppose X is permissible. Then every action in compliance with X is permissible. Assuming:

(*) Every permissible action is part of an ideal action

we can infer that every action in compliance with X is part of an ideal action.

The above justifications are not entirely unproblematic. But there is one major line of questioning which I believe may be resisted. For might question assumption (*) on the grounds that, when the obligations within a code of conduct conflict, there will be no ideal action and so, a fortiori, an action (such as the null action) may be permissible without being contained in an ideal action. However, the only reason to deny that there is any ideal action in this case is that one accepts the Complete Consistency condition above, that any action sanctioned by the code should be consistent. But it seems to me that in the case of conflicting obligations, we might simply allow a code of conduct to contain inconsistent courses of action. It will not then follow, within our framework as it does within the standard possible worlds framework, that everything whatever is obligatory, since an inconsistent action will not, in general, contain every other action.

One great advantage of the present approach is that we can provide a unified account of what is obligatory and what is permitted by reference to a single code of conduct. Once we give up assumptions, such as (*) above, it may be necessary to 'fracture' a code of conduct into two components which separately specify what is relevant to the permissibility or obligatoriness of a given content. One is then able to deal with certain anomalous cases but at great expense in elegance and simplicity and with little gain in generality.

§4 Some Remarks on the Semantics

We make some remarks on the form of the clauses, how they might be extended to other deontic operators, and how they compare with the standard possible worlds clauses.

(1) From a formal point of view, the clauses for the operators O and for P are both ∀∃ (for-all/for-some) in form.[1] This is in marked contrast to the clauses for the operators in the standard possible worlds semantics, according to which:

O(S) is true iff S true in all deontic alternative worlds

P(S) is true iff S is true in some deontic alternative world

which are respectively universal and existential in form. Indeed, the quantifier relevant to the embedded clause X in (i)′& (ii)′ is existential in the case of obligation and universal in the case of obligation; and so one might claim, with some justice, that the standard treatment gets the correspondence with the quantifiers completely backwards!

We also lose the duality in the clauses (i)′& (ii)′ that parallels the duality in the two quantifiers. But oddly enough, the clauses are dual in another respect, since there is a reversal of mereological role and order in going from the one clause to the other.

(2) Clauses (i) and (ii) above are appropriate for *bounded* obligation and for *strict* permission. The clause for unbounded obligation may be obtained by combining the two clauses:

(iii)$^{+}$ $O^{P}(X)$ is true if *C* subsumes *X* and if *X* subserves *C*;

or, more explicitly:

(iii)$^{+}$ $O^{P}(X)$ is true if every ideal course of action contains an action in compliance with X and every action in compliance with X is contained in an ideal course of action.

Thus $O^{P}(X)$ will have the same truth-conditions as the conjunction $O(X) \wedge P(X)$ and an unbounded statement of obligation will serve a dual purpose, both stating what one is obliged to do and also specifying the permissible actions by which the obligation might be discharged.[2] There is an interesting question as to whether ordinary language 'ought' and its cognates are to be understood in a bounded or unbounded sense. Certainly, 'you should post the letter or destroy it' in some sense implies 'it is permissible to destroy it'. But I am inclined to think of this as some kind of pragmatic implication. There seems to be no contradiction involved in saying 'you should post the letter or destroy it and, since you should not destroy it, you should post it.' If this

---

[1]Aloni ([2007], 76) provides a similar (albeit modal form) of the ∀∃ clause for permission but adopts the quantificational dual ∃∀-form for obligation.

[2]Aloni ([2007], 86) introduces an imperative analogue of O^{P}, for which she adopts a somewhat similar clause.

is right, then O^P captures a pragmatically strengthened meaning of ought-statements rather than their strict literal meaning.

We should also note that the two conditions, that *C* subsumes *X* and that *X* subserves *C*, are the two conditions required for the content *X* to be part of the content *C*. Thus, under the proposed semantics, an unbounded statement of obligation will play the role of specifying part of the content of the code of conduct. Given our previous semantics for imperatives, we may also see the truth-conditions for unbounded obligation as relating directly to imperative inference since, for an obligation statement $O^P(X)$ to be true is for the embedded imperative X to follow from the implicit code of conduct.

(3) The clause for weak permission P may be obtained by replacing talk of part-whole with talk of compatibility:

(ii)′ P(X) is true if every action in compliance with X is compatible with a course of action sanctioned by the code.

Thus strict permission requires that the condoned actions should be ruled *in* by the code of conduct whereas weak permission only requires that they not be ruled *out*.

Clause (ii)′ would appear to be a variant rather than a special case of clause (ii). But there is a way of seeing it as a special case. For suppose we replace a code of conduct *C* by the set *C**of complete courses of action compatible with some member of *C*. Then for an action to be compatible with a conduct sanctioned by *C* is simply for it to be a part of a course of action sanctioned by *C**. Thus we can also see the semantics for weak permission as arising from a conception of codes in which only complete courses of action are sanctioned.

Under the present approach, strict permission would appear to be the more straightforward notion; weak permission must be obtained by replacing mereological with modal relationships or by imposing special conditions on codes of conduct. This is in contrast to the standard possible worlds approach, in which permission can only be understood as weak permission and under which it is difficult even to see how strict permission might be defined.

(4) Just as we might give a modal clause for weak permission, so we might also for 'weak' obligation:

(i)′ O(X) is true if *C* necessitates *X*, i.e. if it is impossible for a course of action in *C* to be performed without an action in *X* being performed.

Thus under the weak modal criterion, if one is obliged to square the circle then one is obliged to do anything whatever, whereas this will not follow under the strong mereological criterion.

As before, we may see the modal criterion as a special case of the mereological criterion, obtained by replacing the code of conduct *C* with the corresponding set *C** of completions. For *C** will necessitate *X* just in case every member of *C** contains a member of *X*.

Clause (i)′ corresponds, of course, to the standard possible worlds clause, with *C* the deontic sphere of alternatives. Thus the operator O, so understood, will be intensional; O(X) and O(X′) will have the same truth value whenever X and X′ are necessarily co-enacted.

The same is not true of clause (ii)′ for weak permission, since it embodies a free choice effect; the permissibility of X requires, not simply that *X* be compatible with *C*, but that every member of *X* be compatible with *C*. We could get a counterpart to the standard clause by replacing (ii)′ with:

(ii)′′ P(X) is true if *some* action in compliance with X is compatible with a course of action sanctioned by the code

and the resulting notion of permission would then also be intensional.

In either case, we can state the clauses for permission and obligation by reference simply to the *permissible* courses of action since, as already noted, any complete permissible action will automatically be ideal.[3]

(5) There is a familiar reduction of deontic to alethic modal logic, deriving from Anderson [1966] and Kanger [1957]. Although they have standard versions of modal logic in mind, their general method of reduction may be applied, perhaps somewhat surprisingly, to our own semantical approach.

Let us begin with permission. Each code of conduct *C* may be closed under part to give what one might call the *code of permission* $C{\downarrow}$. Thus an action is sanctioned by $C{\downarrow}$ just in case it is part of a course of action sanctioned by *C*. Intuitively, $C{\downarrow}$ consists of the actions permitted by the code; and the only aspect of *C* required to state the clause for permission is given by $C{\downarrow}$.

Let us now introduce a constant 'OK' for the imperative 'do something alright (permitted)'. The actions in compliance with OK will be the members of $C{\downarrow}$. Let us also introduce an implicational connective ⇒ for exact entailment between imperatives; X ⇒ Y is to be true if every action in compliance with X is in compliance with Y. It is then evident that P(X) will be true just in case X ⇒ OK is true. To say that X is permitted is to say that each action in compliance with X is permitted, which is to say that X exactly entails the imperative 'do something permitted'.

Alternatively, we might introduce a constant 'A-OK' for the imperative 'do something completely alright (ideal)'. The actions in compliance with A-OK will then be the members of *C*. Say that the action *a* is in *sub-compliance with* the imperative X if it is part of an action in compliance with X and say that X *sub-entails* the imperative Y if any action in sub-compliance with X is in sub-compliance with Y, which is to say that any action in compliance with X is in sub-compliance with Y. The statement P(X) of permission will then be true just in case X sub-entails A-OK or, using $\Rightarrow_*$ for sub-entailment, P(X) will be true just in case $X \Rightarrow_*$ A-OK is true.

It should be noted that both $\Rightarrow_*$ and ⇒ have the property that $X \vee Y \Rightarrow_{(*)} Z$ implies $X \Rightarrow_{(*)} Y$ and $Y \Rightarrow_{(*)} Z$ even though neither has the property that $X \Rightarrow_{(*)} Z$ implies $X \wedge Y \Rightarrow_{(*)} Z$. Thus the reduction will give us the desirable result that P(X ∨ Y) implies P(X) and P(Y) without giving us the undesirable result that P(X) implies P(X ∧ Y).

Let us turn to obligation. Each code of conduct *C* may be closed under wholes to give what one might call the *code of obligation* $C{\uparrow}$. Thus an action is sanctioned by $C{\uparrow}$ just in case it contains a course of action sanctioned by *C*. Intuitively, $C{\uparrow}$ consists of the actions that discharge all of one's obligation; and the only aspect of *C* required to state the clause for obligation is given by $C{\uparrow}$.

Let us now introduce a constant 'AR' for the imperative 'do something all right (i.e. discharge all of one's obligations)'. The actions in compliance with AR will be the members of

[3] It is perhaps worth noting that although Aloni and Ciardelli [2013] have option sets, which correspond to our actions and are identified with sets of worlds, their counterpart to our code of conduct is a set of worlds and so they do not avail themselves of the more general notion of a code of conduct.

$C\uparrow$. O(X) will then be true just in case AR ⇒ X is true. To say that X is obligatory is to say that the imperative ‘do something all right’ exactly entails X.

Alternatively, we might say that the action *a* is in *super-compliance with* the imperative X if it contains an action in compliance with X and that X *super-entails* the imperative Y if any action in super-compliance with X is in super-compliance with Y, which is to say that any action in compliance with X is in super-compliance with Y.[4] The statement O(X) of obligation will then be true just in case AR super-entails X.

We see that by suitably interpreting the ‘sanction’ constant and the conditional we may obtain Anderson/Kanger type reductions for both obligation and permission, although the interpretation of either the constant or of the conditional must be different in each case.[5] We can, of course, provide a reduction for obligation under the standard possible worlds account by interpreting O(X) as C → X, where C stands for the set of deontic alternatives and → for strict implication. But it is hard to see how a corresponding reduction for statements of permission might go. Hence Nute’s claim, "We cannot reduce any interesting notion of permission to conditionals no matter how we interpret conditionals" (Nute [1985], 179). However, if I am not mistaken, this claim may be something of an exaggeration.

(6) We have so far given conditions for when a deontic statement of the form O(X) or P(X) is true but we have not specified when a state verifies or falsifies such a statement. This would be required if we wished, for example, to say when one deontic statement analytically entailed another.

There are a number of different ways in which this might be done. The most straightforward is as follows. Each code of conduct *C* is understood to be the state that consists in its members c_1, c_2, ... being all and only the ideal courses of action. We might say, in this case, that the code *C prevails*; and so the code is, in effect, being identified with the state that it prevails.

No code, as so conceived, will be a part of any other code, and so the space of codes will have the following ‘flat’ structure:

■
/ / \
C_1 C_2 ... C_k C_{k+1} ...
\ \ / /
□

with the full state on top, the null state at the bottom, and all of the codes of conduct in between. Given that the codes C_1, C_2, ... are all possible, the only impossible state will be the full state ■ and any two codes will be incompatible. In effect, each code is a mini-world, completely settling which deontic statements do and do not hold.

---

[4]I have also talked in this connection of *inexact* compliance and entailment.

[5]Some rather different attempts at reduction using a non-standard conditional are to be found in Asher & Bonevac [2005] and Barker [2010].

We may then adopt the following clauses for when an atomic deontic statement is verified or falsified by a code *C*:

(i)$^+$ *C* verifies O(X) iff *C* subsumes *X*;

(i)$^-$ *C* falsifies O(X) iff *C* does not subsume *X*;

(ii)$^+$ *C* verifies P(X) iff *X* subserves *C*;

(ii)$^-$ *C* falsifies P(X) iff *X* does not subserve *C*.

These clauses may then be extended to truthfunctional compounds of atomic deontic statements in the usual way.

However, we might wish to provide a less demanding account of what verifies or falsifies a deontic statement, making clearer what it is about the code that is responsible for the statement's being true or false. To this end, we might associate two states with each code *C*: the state C^* of *C upwardly prevailing* in the sense that the prevailing code contains *C* as a part; and the state C_* of *C*'s *downwardly prevailing* in the sense that the prevailing code is a part of *C*. The state C^* will be part of the state D^* just in case *C* is a part of *D*; and the state C_* will be a part of D_* just in case *D* is a part of *C*.

We might now adopt the following clauses for the atomic deontic statements:

(i)$^+$ *s* verifies O(X) if *s* is of the form C^* and *C* subsumes *X*;

(i)$^-$ s falsifies O(X) if *s* is of the form C_* and *C* does not subsume *X*;

(ii)$^+$ *s* verifies P(X) if *s* is of the form C^* and *X* subserves *C*;

(ii)$^-$ *s* falsifies P(X) if *s* is of the form C_* and *X* does not subserve *C*.

Once we combine obligation and permission statements, we will need to combine states of the form C^* and C_*. We may do this by taking each state to be of the form (C^*, D_*), where what was previously C^* is now $(C^*, \Box)$ and what was previously C_* is now $(\Box, C_*)$. Part-whole is defined "point-wise". Thus (C^*, D_*) is a part of (E^*, F_*) if C^* is a part of E^* and D_* is a part of F_*. And the state (C^*, D_*) will be possible if some code *E* lies between *C* and *D*. Thus when previously we took *C* to verify a deontic statement, we now take this to be the special, and highly demanding, case in which (C^*, C_*) verifies the statement.

§5 <u>Imperatival Validity Revisited</u>

I would like to consider various different ways of formulating the criterion of validity for imperative inference in the light of the previous semantics for deontic logic. The reformulations are technically trivial and more of philosophical than technical interest.

Recall that the imperative inference X/Y was taken to be valid if two conditions were met: *X* subsumes *Y*; and *Y* subserves *X*. This definition does not take the usual form, in which the validity of an inference is a matter of some value or values being preserved in the transition from premiss to conclusion; and it is natural to wonder whether a definition of this sort can be given.

There is a way in which the answer to this question is trivial. For let us suppose that entailment (the relation that holds between X and Y when the inference X/Y is valid) is both reflexive and transitive. Then X will entail Y just in case for every Z, Z entails Y whenever Z entails X. For the left-to-right direction will follow from transitivity; and the right-to-left direction will follow from reflexivity upon setting Z = X. Taking the values in question to be properties of the form 'entailed by Z' will then give us an account of entailment in terms of the preservation of these values.

Let us now apply this 'cheap trick' to the validity of imperative inference. Say that the imperative X *conforms to* the code of conduct *C* if *X* subserves and is subsumed by *C*. Then the cheap trick will give us that X entails Y just in case Y conforms to any code of conduct to which X conforms.[6]

But this gives us a new way to understand the force of imperative inference. For we may suppose that prior to the stipulation of the premisses, there is a prevailing code of conduct (as given by the moral code, a body of law, some previous imperatives, or the like). We can then understand the imperative inference as telling us that the conclusion will conform to the prevailing code as long as the premisses do.

This is in line with thinking of there being an underlying warrant for the imperative in terms of a reason or norm and of the validity of an imperative inference as then consisting in the preservation of such a warrant (as in Vranas [2011], for example). For we may take the reason or norm to be a code of conduct and we can then take the warrant of the imperative to consist in its conformity to the reason or norm. But of course, from our own point of view, this way of thinking does not give us an actual criterion of validity since the relevant notion of validity is already presupposed in what it is to have a warrant.

The present criterion of imperative validity also makes evident the connection between imperative and deontic reasoning:

The Imperative-Deontic Link The imperative inference X/Y is analytically valid iff the deontic inference $O^P(X)/O^P(Y)$ is classically valid.

For the classical validity of the deontic inference simply amounts to Y conforming to any code of conduct to which X conforms. In effect, one reasons, in imperative inference, from the obligation implicit in the premisses to the obligation implicit in the conclusion. Thus there is a sense in which imperative inference is classical after all but under a suitable modal understanding of the premisses and conclusion.

The criterion also provides us with yet another understanding of imperative validity in terms of 'updating'. For we might think of an imperative, not as something that conforms to a prevailing code of conduct, but as a means of updating the code. Thus if the prevailing code is *C* and the imperative premiss is X, then the result *C*[X] of updating *C* with X will be $C \wedge X$. We simply combine the actions sanctioned by *C* with the actions in compliance with X. The inference X/Y will then be valid if and only if Y conforms to the updated code *C*[X]. This is something like the definition of validity in the update semantics of Veltman [1996] though, as previously noted (in §5 of Part I), we cannot assume, in the usual way, that Y's conforming to *C*[X] is a matter of *C*[X] being identical with *C*[X][Y].

This particular connection is relevant, I believe, to the inferential interface between deontic and imperative statements. For it is natural to suppose that from an imperative X (such as 'Shut the door'), we can infer the corresponding unbounded statement of obligation $O^P(X)$ ('so

---

[6]Strictly speaking, a code of conduct is a convex set. But it is readily shown that this is a harmless addition, since X will subserve (or be subsumed by) a set of actions *C* if and only if it is subserves (or is subsumed by) the convex closure of *C*. For the purposes of the application, I have also switched from imperatives to their contents.

you ought to shut the door and you may shut the door'). We do not yet have a notion of validity for which this is so since the inference involves a mix of imperative and indicative statements. Such mixed inferences will be discussed at greater length in Part III, but let us note that the present conception of updating provides us with a handle for dealing with this particular case. For we may take the inference from an imperative X to a deontic statement S to be valid if, for any code of conduct *C*, S is true in *C*[*X*]; and the inference from X to $O^P(X)$ will then be valid.

§6 Some Special Inferences

I should like to discuss some cases of valid or invalid deontic inference of special interest, dealing first with inferences involving obligation and permission separately and then with inferences involving both together. In each case, I have paid particular attention to the comparison with standard deontic logic and have given somewhat informal proofs of the various claims of validity and invalidity.

P-Inference

1. We should note right away that the rule of Simplification will hold, i.e. $P(X \vee Y)$ will entail P(X) and P(Y), in contrast to the standard semantics for deontic logic. For if every action in compliance with $X \vee Y$ is contained in an ideal action (one sanctioned by the code of conduct) then a fortiori every action in compliance with X or in compliance with Y will be contained in an ideal action.

However, $P(\alpha)$ will not entail $P(\alpha \wedge \beta)$, for even if every action in compliance with α is contained in an ideal action that is no reason to assume that any action (let alone every action) in compliance with $\alpha \wedge \beta$ is contained in an ideal action. Thus the most immediate logical problem posed by the existence of free choice permission is solved.

2. $P(X \wedge Y)$ will entail P(X), in line with the standard semantics. For suppose that every action in compliance with $X \wedge Y$ is contained in an ideal action. Take now any action *a* in compliance with X. Under the semantics, there will be an action in compliance with any imperative and hence an action *b* in compliance with Y. But then $a \sqcup b$, and hence *a*, will be contained in an ideal action.

However, P(X) will not entail $P(X \wedge X)$ in marked contrast to the standard semantics. This is because actions *a* and *b* in compliance with X may be part of an ideal action even though the action $a \sqcup b$ in compliance with $X \wedge X$ is not part of an ideal action.

Let us say that the imperative X is *definite* (with respect to a model) if there is exactly one action in compliance with X and is otherwise *indefinite*. Then for definite X, P(X) will entail $P(X \wedge X)$. This is but one of many cases in which a principle not valid for all imperatives is valid for definite imperatives.

3. The deontic formula $P(\top)$, according to which the null action is permissible, will be valid if we assume Non-Emptiness, i.e. that the code of conduct contains at least one action, since the single action, $\square$, in compliance with $\top$ will be contained in that action. But dropping that assumption, i.e. allowing for the empty code of conduct, will render $P(\top)$ invalid, since then no action sanctioned by the code will contain $\square$. Thus $P(\top)$ tells us, in effect, that some action is sanctioned by the code of conduct.

However, the deontic formula P(X) is not valid for an *arbitrary* tautology X, in contrast to standard deontic logic. For example, $P(\alpha \vee \neg\alpha)$ is not valid since otherwise, by Simplification, both $P(\alpha)$ and $P(\neg\alpha)$ would be valid. Indeed, it can be shown that P(X) will not be valid for any formula X that contains an imperative atom (even under tight restrictions on the code of conduct).

The deontic formula $\neg P(\perp)$ is also not valid, since it will be false when the sole action sanctioned by the code of conduct is ■. However, assuming Non-Emptiness and Complete Consistency (that every action sanctioned by the code of conduct is consistent), then $\neg P(X)$ will be valid for any classically inconsistent formula X, since only inconsistent actions will then be in compliance with X.

If $\neg P(X)$ is taken to be valid for classically inconsistent formula X, then, given Simplification, it will render $\neg P(X)$ valid for many formulas which are *not* classically inconsistent, in marked contrast to the standard case. For example, $\neg P(\neg X \wedge (X \vee Y))$ will be invalid, since $P(\neg X \wedge (X \vee Y))$ will entail $P((\neg X \wedge X) \vee (\neg X \wedge Y))$, which, by Simplification, will entail $P(\neg X \wedge X)$.

Even if $\neg P(X)$ is *not* taken to be valid for classically inconsistent X, we should note that $P(\perp)$ will still entail P(Y) - if the completely impossible is permitted then everything whatever is permitted. For if $P(\perp)$ is true then ■ will be an ideal action, and P(X) will then be true since every action is contained in ■. However, we will not in general have that P(X) entails P(Y) for X a classical contradiction, since a code of conduct may sanction an inconsistent action without sanctioning ■ and so, in particular, we will not have P(X) entails $P(\perp)$.

<u>O-Inference</u>

1. O(X) will entail $O(X \vee Y)$, in line with standard deontic logic. For if every ideal action contains an action in compliance with X then, a fortiori, it will contain an action in compliance with $X \vee Y$. Similarly, $O(X \wedge Y)$ will entail O(X). For if every ideal action contains an action in compliance with $X \wedge Y$ (of the form $a \sqcup b$ for *a* an action in compliance with X and *b* an action in compliance with Y) then it will thereby contain an action (viz. *a*) in compliance with X.

On the other hand, $O^P(X)$ will *not* entail $O^P(X \vee Y)$. For suppose *a* is the sole action in compliance with X and *b* the sole action in compliance with Y; and suppose *a* is the sole action sanctioned by the code of conduct. Then $O^P(X)$ will be true. But there is no reason why $O^P(X \vee Y)$ should be true. For its truth would require the truth of $P(X \vee Y)$, which, in its turn, would require the truth of P(Y). It is no doubt the 'unbounded' interpretation of obligation that stands in the way of accepting the inference from 'you ought to post the letter' to 'you ought to post the letter or burn the house down', since the latter seems to grant permission to burn the house down.

2. We have the standard principles that $O(X) \wedge O(Y)$ entails $O(X \wedge Y)$ and that $O(X \wedge Y)$ entails $O(X) \wedge O(Y)$? For if every ideal action contains an action *a* in compliance with X and an action *b* in compliance with Y then it will contain an action $a \sqcup b$ in compliance with $X \wedge Y$, and conversely.

We also have the principle that $O^P(X \wedge Y)$ entails $O^P(X) \wedge O^P(Y)$ for unbounded obligation, since $O(X \wedge Y)$ entails $O(X) \wedge O(Y)$ and $P(X \wedge Y)$ entails $P(X) \wedge P(Y)$. However, we do not have the principle that $O^P(X) \wedge O^P(Y)$ entails $O^P(X \wedge Y)$ for unbounded obligation. Indeed, we do not even have that $O^P(X)$ entails $O^P(X \wedge X)$ for, when X is an indefinite imperative, $O^P(X)$ may be true while $P(X \wedge X)$ is false.

3. The deontic formulas $O(\top)$ and $O^P(\top)$, according to which the null action is obligatory, will be valid if we assume Non-Emptiness, since $P(\top)$ is then valid and since any ideal action will contain the single action □ in compliance with $\top$. But dropping the assumption will render $O(\top)$, and hence $O^P(\top)$, invalid, since then no action sanctioned by the code will contain □.

The deontic formula O(X) is not valid for an *arbitrary* tautology X, in contrast to standard deontic logic. For example, $O(\alpha \vee \neg\alpha)$ is not valid, since there is no reason, in general, to think that a code of conduct will contain an action in compliance with or in contravention to α. Indeed, it can be shown that O(X) will not be valid for any formula X that contains an imperative atom. Of course, matters will be different if we suppose that the courses of action sanctioned by the code of conduct are all complete.

The formula $\neg O(\perp)$ is also not valid, since it will be false when the sole action sanctioned by the code of conduct is ■. However, assuming Consistency, that some ideal action is consistent, $\neg O(X)$ will be valid for any classically inconsistent formula X, since only inconsistent actions will then be in compliance with X.

$O(\perp)$ entails O(Y) - if the completely impossible is obligatory then everything whatever is obligatory. For if $O(\perp)$ is true then every ideal action will contain ■ and hence be identical to ■, in which case every ideal action will contain an action in compliance with Y.

However, as before, we will not in general have that O(X) entails O(Y) for X a classical contradiction, since a code of conduct may sanction an inconsistent action without sanctioning ■.

4. In the case of a moral - or, more generally, a normative - dilemma, we will have both O(X) and $O(\neg X)$ true or both O(X) and O(Y) true, with X is incompatible with Y. The usual attitude to cases of this sort (granted that they arise) is to question the inference from O(X) and $O(\neg X)$ to $O(X \wedge \neg X)$ or, more generally, the inference from O(X) and O(Y) to $O(X \wedge Y)$ when X is incompatible with Y. Our attitude, on the other hand, has been to allow the inference but to deny that 'explosion' thereby results, with everything whatever being obligatory, as would be the case in more standard versions of deontic logic. However, it is worth noting that we could also follow the standard route by allowing there to be several codes of conduct, perhaps not all compatible with one another, and then taking O(X) to be true under this class of codes if it true under one of the codes (much as in van Fraassen [1973]).

O/P Inference

1. Under the present semantics, O(α) does not entail P(α). For suppose *a* and *b* are in compliance with α but that *a* is the only action sanctioned by the code of conduct. O(α) will then be true but there is no reason, in general, why *b* should be part of *a*. To take an ordinary example, it may be obligatory to post the letter and hence obligatory to post the letter or burn the house down and yet not permissible to post the letter or burn the house down.

We do, of course, have the principle that $O^P(X)$ entails P(X) for unbounded obligation. And we also have it for bounded obligation and definite imperatives X, given Non-Emptiness. For if O(X) is true then every ideal action contains the single action a_0 in compliance with X and so, given that some action is ideal, every action in compliance with X, viz. a_0, will be contained in an ideal action. So this is another case in which a generally accepted principle holds for definite imperatives but not for indefinite imperatives and our intuitive judgement of validity in this case no doubt derives from considering the special case in which the imperative is definite.

2. We also do not have the standard principle that O(X) and P(Y) entails P(X ∧ Y). Indeed, if we were to set Y = ⊤, then this principle would yield, as a special case, that O(X) and and P(⊤) entails P(X ∧ ⊤), which, in its turn, given Non-Emptiness, would yield O(X) entails P(X).

However, as it to be expected, the principle that O(X) and P(Y) entails P(X ∧ Y) will hold when X and Y are definite or, indeed, when X alone is definite. For suppose that O(X) is true, so that every ideal action contains the single action a_0 in compliance with X, and that P(Y) is true, so that every action *b* in compliance with Y is contained in an ideal action. Take now any action, of the form $a_0 \sqcup b$, in compliance with X ∧ Y. Then *b* is contained in an ideal action *c*, which must contain a_0; and so *c* must contain $a_0 \sqcup b$.

3. In standard deontic logic, the deontic operators are dual in the sense that O(X) is equivalent to ¬P(¬X) and P(X) to ¬O(¬X). These duality principles fail in the present context because of two separate reasons - one arising from the strict sense of permission and the other from indefiniteness in the imperative X.

O(X) will entail ¬P(¬X) as long as the code of conduct is completely consistent. For suppose, for reductio, that O(X) and P(¬X) are both true. Then some action *c* sanctioned by the code of conduct will contain an action *a* in compliance with ¬X and *c* in its turn will contain an action in compliance with X; and so *c* will be inconsistent.

However, ¬P(α) will not entail O(¬α). For suppose that the sole action sanctioned by the code of conduct is the null action □ and that *a* is a non-null action in compliance with α and that no null action is in contravention to α. Then ¬P(α) is true since no ideal action contains *a* and O(¬α) is false since the ideal action □ contains no action in compliance with ¬α. As is clear, the entailment may fail to go through even when there is a single action in compliance with and a single action in contravention to α. The reason intuitively speaking, is that P expresses strict permission. Thus under the null code of conduct, no action (beyond □) will be permitted and no X which does not allow null compliance will be obligatory.

Suppose, however, that we take the courses of action sanctioned by the code of conduct to be complete (as under a modal interpretation of the deontic operators). Then the entailment may still fail to go through since what is not permitted may be indefinite. Thus consider the entailment from ¬P(α ∨ β) to O¬(α ∨ β), equivalent to the entailment from ¬P(α) ∨ ¬P(β) to O(¬α ∧ ¬β). If this latter entailment were to hold then it would require that ¬P(α) entail O(¬α ∧ ¬β), which clearly is not so.

However, the entailment *will* go through if we insist both that the courses of action sanctioned by the code of conduct be complete and that the imperative α be definite. For suppose P(α) is false. Then the sole action a_0 in compliance with α will not be contained in an ideal (and complete) action sanctioned by the code. But this then means that every action sanctioned by the code must contain an action in contravention to α; and so O(¬α) will be true. Indeed, in this case, there is, in effect, no quantification over the actions in compliance with α and so our clauses will reduce to the familiar clauses from the possible worlds semantics.

4. Although we cannot, in general, define P(X) as ¬O(¬X), there is a definition of P in

terms of $O^P(X)$, though not, I suspect, of O^P in terms of P. For in line with the permissive sense of imperative, we may define P(X) as $O^P(X \vee \top)$, as long as the code of conduct is non-empty. For then every ideal action will contain an action in compliance with $X \vee \top$, viz. $\square$, while $P(X \vee \top)$ will be equivalent to $P(X) \wedge P(\top)$, which will be equivalent to P(X) given Non-Emptiness.

It is not altogether clear to me that we have much use for O^P as opposed to O. Certainly, the most informative and useful way to state an obligation is by indicating the ways of realizing the obligation that are permissible. And if this is so, then we may formulate the language of deontic logic in the present setting by using a single deontic operator, just as in the standard formulations but with the difference that the operator must be the operator for obligation rather than permission.

§7 An Axiom System

I briefly outline a system of deontic logic appropriate for the deontic formulas which are valid within the truthmaker approach.

We take a *deontic formula* formula to be a truth-functional compound of formulas of the form O(X) and P(X) for X an imperative. Such a formula is then taken to be *valid* if it is true in all models whose code of conduct is non-empty. We also suppose that, in the models in question, the imperative atoms will be doubly definite - i.e., for each imperative atom α, there will be a single action in compliance with α and a single action in contravention to α.

We take two imperatives X and Y to be *inexactly equivalent* if, in each model, every action in compliance with X contains an action in compliance with Y and every action in compliance with Y contains an action in compliance with X. In axiomatizing the class of valid deontic formulas, we take for granted when two imperatives X and Y are inexactly equivalent. Indeed, X and Y will be inexactly equivalent when $X \vee Y$ analytically entails both X and Y. However, it should be noted that, since we have assigned a doubly definite content to imperative atoms, it needs to be assumed that $\lambda \wedge \lambda$ is inexactly equivalent to λ, where λ is a literal, i.e. an atom or its negation.

An imperative formula X is said to be (*syntactically*) *definite* if it is a conjunction of literals. Clearly, each syntactically definite formula X will be semantically definite, there will be a single action in compliance with X.

We then have the following axioms:

Equivalence

$O(X) \equiv O(X')$

$P(X) \equiv O(X')$

for X inexactly equivalent to X′.

Distribution

$O(X \wedge Y) \equiv O(X) \wedge O(Y)$

$P(X \vee Y) \equiv P(X) \wedge P(Y)$

Weakening

$O(X) \supset O(X \vee Y)$

$P(X \wedge Y) \supset P(X)$

Non-Triviality

$O(\top)$

$P(\top)$

Mixture

$O(X \vee Y) \wedge \neg P(X) \supset O(Y)$, for definite X

$O(X_1 \vee X_2 \vee ... \vee X_n) \wedge P(Y) \supset P(X_1 \wedge Y) \vee P(X_2 \wedge Y) \vee ... \vee P(X_n \wedge Y)$, for definite X_1, X_2,... , X_n and Y.

The theorems of the system are the truth-functional consequences of these various axioms. Soundness may then be established by a straightforward induction and completeness by using normal forms.

The resulting logic is not, of course, closed under substitution. Thus even though $P(\alpha) \supset P(\alpha \wedge \alpha)$ is a theorem, $P(\alpha \vee \beta) \supset P((\alpha \vee \beta) \wedge (\alpha \vee \beta))$ is not. The restrictions on Mixture point to the utility of distinguishing between definite and indefinite statements in reasoning that combines consideration of what is obligatory and what is permissible. However, it would be of interest to determine which system we would obtain when no such distinction is made, with a formula A being a theorem of the new system just in case each substitution instance of A is a theorem of the original system. One might also profitably consider various extensions and fragments of these systems.

§8 Updating

I should like to make a few remarks concerning the problem of deontic updating. The topic calls for much more extensive discussion, especially in regard to its connection with belief revision. But, at the very least, the present discussion will indicate how very different the problem looks from the present perspective as opposed to the usual possible worlds perspective and how, in various respects, the apparatus of truthmaking makes the problem far more tractable.

In a given context, we may suppose, certain deontic statements directed to a given agent are true and the others false. Suppose now someone in authority, perhaps myself, tell the agent what he ought to do or what he may do, where this is not something that was previously true. The problem of deontic updating is then the problem of explaining which deontic statements will then be true.

Of course, if I tell the person that he ought to do something (or may do something), it will then be true in the new context that he ought to do it (or may do it), even though this was not true before. We may also assume that the change in context is *determinate* - which is to say that it must still hold, in the new context, that each deontic statement is either true or false. But this means, on pain of inconsistency, that some other deontic statements must change from true to false or from false to true; and so the problem is to say which they are (and why).

If we approach the problem from a possible worlds perspective[7], then the set of true and false deontic statements will, in effect, be given by a 'sphere' of ideal worlds. On being told that A is obligatory or permissible, the agent must change the deontic sphere so that A is obligatory or permissible. This means, in the case of obligation, that every A-world must belong to the new deontic sphere and, in the case of permission, that some A-world must belong to the new deontic sphere (even though neither was so before). We presumably want the change to be minimal - we want in some sense that is not altogether clear to minimize the change in the truth-value of the deontic statements; and this presumably translates into the change in the deontic sphere being minimal - we want, again in a sense that is not altogether clear, to minimize the change in which worlds belong to the deontic sphere.

It might be thought that the solution to this problem is straightforward in the case of obligation. For the new sphere can simply be the intersection of the old sphere with the set of A-worlds. In other words, we may restrict the former ideal worlds to the A-worlds. However, there is one special case in which this strategy does not work. For the intersection may be empty, there may be no A-worlds that were previously ideal (even though there are A-worlds); and in this case - if we insist that the deontic sphere be non-empty, i.e. that something be permissible - then it is not at all clear what the change in the deontic sphere should be.

In the case of permission, the problem seems generally hopeless unless one brings some further information to bear upon what the change should be. For let us suppose that A was not originally permissible, i.e. that none of the A-worlds belong to the original deontic sphere. Then, in the absence of any further information, there is no reason to prefer the addition of one of the A-worlds to the sphere as opposed to any other (and at least one must be added if A is to become permissible). One could, of course, add all the A-worlds to the sphere, since this does not discriminate between them, but this would then result, as a rule, in all sorts of monstrous worlds being rendered permissible. The solution to the problem, in the case of permission, therefore requires some further basis for distinguishing between the worlds that might be added to the original sphere.

Let us now consider the same problem from the truthmaker perspective.[8] In this case, the deontic base will not be a deontic sphere, the set of ideal worlds, but a code of conduct, the set of ideal actions. The problem then takes the form of how to modify the code of conduct so as to make the given prescription X obligatory or permitted. This means, in the case of obligation, that every action sanctioned by the new code of conduct should contain an X-action and, in the case of permission, that every X-action should be contained in an action sanctioned by the new code of conduct and, in the case of unbounded obligation, that both requirements should be met.

It is perhaps rather odd to update with a bounded statement of obligation. If I tell the agent that he is obliged to wear a jacket or a tie, then it would normally be supposed that he was

---

[7]As is done in the paper of Lewis [1979] which introduces the problem.

[8]Yablo [2011] also proposes a solution to the problem from the truthmaker perspective, but he still conceives of the problem in terms of a change to the sphere of permissible worlds and, given this and other differences of framework, it is not altogether clear how his solution relates to mine.

being permitted to do either. So let us first consider the case of unbounded obligation and only then turn to the cases of bounded obligation and permission.

There is a way in which the update problem in this case may be trivial. For it may not be my intention to over-ride the original code of conduct. Suppose the original code of conduct contains the action *a* of fighting in the Resistance and suppose that the agent is told 'you ought to stay at home and look after your mother', where this is a matter of performing an action *b* incompatible with *a*. Then the result of the update may simply be a code of conduct consisting of that action $a \sqcup b$ of fighting in the Resistance and staying at home. On this way of looking at the matter, the result C:O^P(X) of updating a code of conduct *C* with the prescription O^P(X) will simply be the conjunction $C \wedge X$ of *C* and *X*, i.e. the code of conduct that consists of all the actions of the form $c \sqcup a$ with *c* in *C* and *a* in *X*. (If codes of conduct are subject to certain closure conditions, such as convexity, it will then also be necessary to subject the resulting code C:O^P(X) to these conditions).

In this case, of course, some of the actions in the new code of conduct may be inconsistent. So let us suppose, as is plausible, that the original code of conduct is completely consistent and that the intent behind the updating is that the new code of conduct should also be completely consistent. How then should the code of conduct be updated? The obvious solution is to apply the consistency filter. The new code of conduct should consist of all those actions in $C \wedge X$ which are consistent.

But if a given action *a* in compliance with X is not compatible with any action sanctioned by the original code of conduct, the resulting code of conduct will sanction no action containing *a* and so X will not be permitted. So what should we do when no action sanctioned by the original code of conduct is compatible with *a*? When presented with a similar problem, under the possible worlds approach, of what to do when the new deontic sphere was empty (no ideal A-worlds), there seemed nothing sensible to say - the update simply failed. But in the present case, we have further resources with which to deal with the problem.

If *a* is already inconsistent, the requirement of complete consistency cannot be preserved and so the update fails, as before. But suppose *a is* consistent. In such a case, what we would like to do is to keep *a* in its entirety and as much of *c* as is compatible with *a*.

To see how this might work, let us suppose that there exists a greatest part c' of *c* compatible with *a*, i.e. c' is compatible with *a* and contains any part of *c* compatible with *a*. In this case, we may let c:a, the result of updating the action *c* with *a*, be $a \sqcup c'$ and then add c:a rather than $a \sqcup c$ to *C*. But even when the greatest part c' does not exist, there may still exist maximal parts c' of *c* compatible with *a*, i.e. parts of *c* compatible with *a* which are not proper parts of any other part of *c* compatible with *a*; and in this case, we can add $a \sqcup c'$ for each of the maximal parts c' to *C*. It is not altogether clear to me how far in this direction we may proceed. But it seems clear that there may be cases in which $a \sqcup c$ is inconsistent and yet there is no reason, on the basis of the given information, for preferring one consistent fusion $a \sqcup c'$, for c' a non-null part of *c*, over any other. In such a case, we should add $a \sqcup c'$ to *C* for *each* non-null part c' of *c* compatible with *a* (or, if there is no such c', we should simply add *a*).

In any case, let us suppose, if only for the sake of simplicity, that an update of the form c:a always exists. Instead of taking the update C:O^P(X) to be $C \wedge X = \{c \sqcup a\text{: } c \in C \text{ and } a \in X\}$, we can take it to be $\{c\text{:}a\text{: } c \in C \text{ and } a \in X\}$, or perhaps better still if we want to make a minimal change to X, we might take $C{:}O^{P}(X)$ to be $\{c{:}a{:}\ c \in C, a \in X$ and either $c \sqcup a$ is consistent or there exists no $c \in C$ for which $c \sqcup a$ is consistent$\}$.

We may now apply a similar strategy for defining updates with bounded statements of obligation or with permission statements. For a bounded statement of obligation, it is not necessary that each action in compliance with X be permissible and so we might simply take $C{:}O(X)$ to be the set $\{c \sqcup a{:}\ c \in C, a \in X$ and $c \sqcup a$ is consistent$\}$. However, the resulting code may be empty; and so if we wish to insist that the code be non-empty, we should take $C{:}O(X)$ to be $\{c \sqcup a{:}\ c \in C, a \in X$ and $c \sqcup a$ is consistent$\}$ when this set is non-empty, as before, and otherwise take it to be $\{c{:}a{:}\ c \in C$ and $a \in X\}$, since in this case there is no reason to prefer any particular $c \in C$ to any other.

Updating with a statement of permission can be regarded as a special case of updating with an unbounded statement of obligation, given the equivalence of P(X) to $O^{P}(\top \vee X)$. Thus we will have $C{:}P(X) = C{:}O^{P}(\top \vee X) = C \cup C{:}O^{P}(X)$. In this case, we will wish to retain each member of the original code of conduct C, since we are not *required* by P(X) to perform any of the actions in compliance with X. But we will also wish each action a in compliance with X to be contained in an action sanctioned by the new code of conduct and this will then be guaranteed by the presence of $C{:}O^{P}(X)$.[9]

We see that we can make much greater progress on how to perform a deontic update when we work with a code of conduct rather than with a sphere of worlds. Part of the reason for this is that the updating is ultimately done at the level of the actions themselves and it is much easier, given the mereological structure of actions, to see how an action should be modified rather than a set of worlds. But another reason is that we are working with a strict notion of permission. When something is not strictly permitted it is relatively easy to see how, through updating, it might become strictly permitted. Under the possible worlds approach, on the other hand, we work with a weak notion of permission; and it is much harder to see how what is not weakly permitted might, through updating, become weakly permitted. From this point of view, the possible worlds approach fails to take full advantage of the performative character of permission statements. For in permitting something it thereby becomes strictly permitted; and by only taking account of the weak content of the performative utterance, the possible world approach can take no advantage of this relatively straightforward form of updating. Our own approach, by contrast, updates on what is strictly permitted. It can still achieve a change in what is weakly permitted, but only indirectly, as a consequence of a change in the code of conduct.

Formal Appendix

Recall from part I that imperative formulas are constructed from imperative atoms α_1, α_2, ... by means of negation ($\neg$), conjunction ($\wedge$), disjunction ($\vee$) and the verum constant $\top$. We now take O(X) and P(X) to be the atomic deontic formulas, for X any imperative formula; and arbitrary deontic formulas, in their turn, are constructed in the usual way from the atomic deontic formulas by means of the usual truthfunctional connectives. We use α, β, γ and the like for

[9] A related account of updating with permissions is to be found in Mastop [2005], p. 110.

arbitrary imperative atoms, X, Y, Z and the like for arbitrary imperative formulas, and P, Q, R and the like for arbitrary deontic formulas.

Recall the definitions of a state space and of a modalized state space from Part I. A state *s* of a modalized space $\boldsymbol{M} = (S, S^{\Diamond}, \sqsubseteq)$ is said to be a *world-state* if it is consistent and if any consistent state is either a part of *s* or incompatible with *s*; and the space ***S*** itself is said to be a *W-space* if every consistent state of ***S*** is part of a world-state. An action space is just a state space under another name.

A *normative action space **A*** is a structure of the form $(A, C, \sqsubseteq)$, where $(A, \sqsubseteq)$ is an action space and *C* (code of conduct) is a non-empty convex subset of *A*. A *normative action model **M*** is a structure $(A, C, \sqsubseteq, |\bullet|)$, where $(A, C, \sqsubseteq)$ is a normative action space and $|\bullet|$ is a bilateral valuation of the usual sort (and similarly when modalized models are in play).

Relative to a normative action model $(A, C, \sqsubseteq, |\bullet|)$, we stipulate the following truth-theoretic clauses for the various deontic operators:

(i) $\models$ O[X] if *C* subsumes *X*;
(ii) $\models$ P[X] if *X* subserves *C*;
(iii) $\models$ O^P[X] if *C* subsumes *X* and *X* subserves *C*.

For the modal construal of these operators, we must suppose we are working within a modalized space $(A, C, \sqsubseteq, A^{\Diamond})$. Relative to such a space, we say that every subset *Y* of *A modally subserves* the subset *X* if every $b \in Y$ is compatible with an $a \in X$ and that *X modally subsumes* (*necessitates*) *Y* if any consistent extension $a^+ \sqsupseteq a$ of an action $a \in X$ is compatible with an action $b \in Y$. The corresponding clauses, relative to a modalized model $(A, C, \sqsubseteq, A^{\Diamond}, |\bullet|)$ are then:

(i) $\models$ O[X] if *C* modally subsumes *X* (the standard definition)
(ii) $\models$ P[X] if *X* modally subserves *C*
(iii) $\models$ O^P[X] if *C* modally subsumes *X* and *X* modally subserves *C*.

References

Aloni , M. [2007] 'Free Choice, Modals, and Imperatives', Natural Language Semantics 15:65-94.

Aloni M.D., Ciardelli & I.A. [2013] 'A logical account of free choice imperatives', in 'The Dynamic, Inquisitive, and Visionary Life of φ, φ?, and ◇φ -- A Festschrift for Jeroen Groenendijk, Martin Stokhof, and Frank Veltman' (eds., M.D. Aloni, M. Franke).

Anderson A. R., [1966] 'The Formal Analysis of Normative Systems' in 'The Logic of Decision and Action' (ed. N. Rescher), Pittsburgh: Univ. of Pittsburgh Press [strong permission]

Asher N., Bonevac D., [2005] 'Free Choice Permission is Strong Permission', Synthese 145: 303-323.

Barker C., [2010] 'Free Choice Permission as Resource Sensitive Reasoning', Semantics and Pragmatics 3:1-38.

Fine K., [2014] 'Permission and Possible Worlds', Dialectica 68.3, 317-336.

Fine K. [2015] 'Angellic Content', Journal of Philosophical Logic, I-28 (2015).

Fox, Danny [2007] 'Free choice disjunction and the theory of scalar implicature', in Uli Sauerland and Penka Stateva (eds.) 'Presupposition and Implicature in Compositional Semantics' 71–120, New York: Palgrave Macmillan.

Hilpinen R., [1982] ‘Disjunctive Permission and Conditionals with Disjunctive Antecedents’, in ‘Intensional Logic: Theory an Applications’ (eds. I. Nimiluoto I. & Saarinen E., Acta Philosophica Fennica 35, 175-94. [strong permission]

Hilpinen, R. (ed.), [1971] Deontic Logic: Introductory and Systematic Read-ings, Dordrecht: Reidel, 1971.

Kamp [1978] for distinction between performative and descriptive uses of permission statements.

Kanger, S., [1957] ‘New Foundations for Ethical Theory: Part 1', duplic., 42pp., reprinted in Hilpinen [1971].

Kratzer, Angelika [1991] ‘Modality’, in Arnim von Stechow, Dieter Wunderlich (eds.). ‘Semantik: Ein internationales Handbuch der zeitgenössischen Forschung’, Berlin: De Gruyter. 639–650. [inconsistent obligations]

Lewis D., [1979] ‘A Problem about Permission’, in: Saarinen E. et al (eds) ‘Essays in Honor of Jaakko Hintikka’, Reidel: Dordrecht, 163-75; reprinted in Lewis [2000], 20-33

Lewis, D. [2000] ‘Papers in Ethics and Social Philosophy’, Cambridge studies in Philosophy, Cambridge: Cambridge University Press.

Mastop , R., [2005] ‘What Can You Do: Imperative Mood in Semantic Theory’, Doctoral Dissertation, University of Amsterdam.

Nute D., [1985] ‘Permission’, Journal of Philosophical Logic 16, 169-90.

Portner P (2005) The semantics of imperatives within a theory of clause types. In: Watanabe K, Young RB (eds) Proceedings of SALT 14, CLC Publications, New York, NY

Portner P (2007) Imperatives and modals. Nat Lang Semantics 15:351–383 [imperative/deontic inference]

Portner P (2009) Modality. Oxford University Press, Oxford

Portner P (2010) Permission and choice. Manuscript, Available from semanticsarchive.net

Portner P, Zanuttini R (1999) ‘The notion of clause type: Some considerations from Exclamatives’, in Smith C (ed) Spoken and written texts, Workshop Proceedings, 29–31 Jan 1999. Texas [to do list, updated using a preference ordering] Linguistics Forum, Austin

Schwager M., [2005] ‘Interpreting Imperatives’, Ph. D. dissertation, Univ. of Frankfurt.

van Fraassen [1973] ‘Values and the Heart’s Command’, Journal of Philosophy 70.1, 5-19 (inconsistent norms)

von Wright, G. H., [1970] [strong permission]

von Wright, G. H. [1999] “Deontic logic: A personal view”,. Ratio Juris 12: 26–38 [for non-prohibition does not imply permission].

Yablo S., [2011] ‘A Problem about Permission and Possibility’ in ‘Epistemic Modality’ (eds. A. Egan, B. Weatherstone), Oxford: Oxford University Press.