Command and consequence*

Josh Parsons

08/10/10

Table of Contents

* Thanks for comments and discussion on this paper to Charles Pigden, Dorothy Grover, Annette Baier, Jonathan McKeown-Green, Denis Robinson, Patrick Girard, Hannah Clark-Younger, Hannah Burgess, Peter Vranas.

1 The problem of imperative consequence

Isn't this an instance of modus ponens?

(A1) Attack at dawn if the weather is fine!	p→q
(A2) The weather is fine.	p
Therefore (A3) Attack at dawn!	q

Isn't this an instance of disjunctive syllogism?

(S1) Attack at dawn unless the weather is fine!	p∨q
(S2) The weather is not fine.	¬q
Therefore (S3) Attack at dawn!	p

Isn't this an instance of the fallacy of affirming the consequent?

(F1) Attack at dawn if the weather is fine!	p→q
(F2) Attack at dawn!	q
*Therefore (F3) The weather is fine.	p

Arguments (A) and (S) are valid; argument (F) is invalid. (Invalid, or apparently invalid, arguments are noted in this paper by adding an asterisk to their conclusions). They are not, however, like the usual examples of such arguments. Their premises and conclusions include sentences in the *imperative* mood ("Attack at dawn!"), as well as the *indicative* ("The weather is fine.")

This poses a problem – the *problem of imperative consequence*. An argument is usually said to be valid iff it is truth-preserving – iff it cannot be that all its premises are true and its conclusion false. But imperatives (it is normally thought) are not truth-apt. They are not in the business of saying how the world is, and therefore cannot either succeed or fail in doing so.[1] The normal criterion of validity cannot be applied to arguments like (A), (S), and (F); or if we insist on applying it, it says that all such arguments, even (F), are valid, since their premises cannot all be true.

To solve the problem, we need to find a new criterion of validity, and I aim to propose such a criterion. I do this in section 3. In section 2, I discuss what I take to the standard approach to the problem, and say why it can't possibly work. Before I do that, however, some remarks on what it would take to solve the problem.

1.1 Three virtues in a criterion of validity

The criterion I want must have the following three virtues. It must be: *general* – applicable to arguments consisting entirely of indicatives, entirely of imperatives, or of a mixture; *conservative* – in agreement with the usual definition given above for the special case of arguments consisting of all indicatives;[2] and finally *adequate* – the criterion must not say of a valid argument that it is invalid, or of an invalid argument that it is valid.

1 One solution to the problem is to claim that imperatives are in fact truth-apt. For discussion of this view, see section 4.5.

2 Note that by "conservative" here, I do not mean what is often meant by this word in discussions of the "conversativeness of logic" in the related literature on "ought" and "is" - for example (Pigden:1989).

How can we decide whether a criterion is adequate? To decide that, we need an independent test for validity – one that is itself general and conservative (if it were not general it could not test a general criterion; if not conservative, it would convict a conservative criterion as inadequate) and yet one that does not itself constitute a criterion (or else my task of finding an adequate criterion would be circular). Here is such a rough-and-ready test: if someone utters all of the premises and yet refuses to also utter the conclusion of an argument (not just because they have a sore throat, but because they don't want to say *that*), would they be making a logical mistake? If yes, the argument is valid, otherwise invalid.[3]

By this standard, (A), for example is valid: if I have commanded you to attack at dawn if a condition is met, and asserted to you that that condition is met, what I have done is tantamount to a command that you attack at dawn. So if I refused to command you to attack at dawn, I am making a logical mistake; what I have already done is tantamount to what I refuse to do. In contrast (F) is invalid: in commanding (F1) I commanded you to attack at dawn if the weather is fine; I went on to command you to attack at dawn (whatever the weather). I would be making no logical mistake if I did not believe that the weather was fine. I would be breaking Grice's conversational maxims by commanding both (F1) and (F2), perhaps, but to do that is not to make a logical mistake.

This test for validity does not constitute a criterion because neither "making a logical mistake" nor "refuses to utter" are very precise concepts. My purpose here isn't to give a criterion of validity, but to elicit intuitive judgements of validity that we can use to calibrate a more exact criterion.

Why aim for a conservative criterion, given that I also aim for an adequate one? For two reasons. First, because conservativeness is a guide to adequacy across a large range of cases, and it is easier to check that a criterion is conservative than to expose to it to a large number of examples of arguments consisting only of indicatives. Second, because I would like to avoid the task of fixing controversial problems with the classical definitions of validity. For example, given the definition given above (an argument is valid iff it cannot be that all its premises are true and its conclusion false), every argument with inconsistent premises is valid (for the premises cannot all be true). This is a problem, but it is not *my* problem, and solving it has little to do with the nature of the imperative mood. For similar reasons, I will be treating the English "if...then" conditional as a material conditional. This view has its problems, but it enables me to focus on the problem before us.

2 The standard view

The *standard view* of imperatives, as I will call it, is very widely held. It stems from the plausible idea that imperatives and indicatives differ in the kind of speech act they are typically used to make. Imperatives are used to make *commands*; indicatives are used to make *assertions*.[4] While imperatives do not have truth values or truth-conditions, they do have something analogous: compliance values and

3 The rough-and-ready test cannot be applied to arguments whose conclusions have conventional implicatures that the premises lack. A person might, without making a logical mistake, be willing to assert (H1) "She was poor and honest" but unwilling to also assert (H2) "She was poor but honest", even though (H1) logically entails (H2). Fortunately none of the examples we will need to apply it to have that feature.

4 I use the term "command" here in a technical sense which is broader than the usual non-technical one. Imperative sentences can be used to make speech acts that are not military-style orders: requests, suggestions, and the like. For the purposes of this paper, these are all to count as "commands" in my technical sense. J. L. Austin introduced the term "exercitive" for a broad category of speech acts including commands (Austin:1962, p. 155) ; but exercitives are a much broader category than I intend command to be, including bequeathings, finings, and that Austinian standby, namings of ships after dictators.

compliance-conditions. An indicative is true iff the world is how it says it is; an imperative is *complied with* iff the world is how it commands that it be. Consider the following two sentences:

1) You attack at dawn.
2) Attack at dawn!

(1) is true iff you attack at dawn; (2) is complied with iff you attack at dawn.[5]

In formal semantics, the truth-conditions of an indicative are often reified – treated as a abstract object that a semantic theory may talk about. A usual way to do this is to regard truth-conditions as sets of possible worlds,[6] where the truth-conditions of a sentence S is the set of worlds at which S is true. So for example, the truth-conditions of (1) will be the set of worlds at which you attack at dawn. We can do something analogous for imperatives. Let the compliance-conditions of an imperative S be the set of worlds at which S is complied with. So the compliance-conditions of (2) will also be the set of worlds at which you attack at dawn.

It will be useful to have a more general label that covers both the relationship between an indicative and its truth-conditions (on the one hand) and the relationship between an imperative and its compliance-conditions (on the other). Let us use "*content*" for this. So the content of (1) is the set of worlds at which you attack at dawn; the content of (2) is the very same set. I will also use "*proposition*" to mean the kind of thing that is the content of a sentence, according to the standard view. For the purposes of this paper, propositions are sets of worlds.

So, imperatives and indicatives both have propositions as their contents; and an imperative and an indicative can even have the same proposition as their respective contents, as (1) and (2) do. What then, is the difference between an imperative and an indicative? (1) and (2) are alike in content; how do they differ? The difference is a difference of *force* – a difference of the kind of speech act (1) and (2) are used to make. (1) has *assertoric* force – it is used to make assertions – (2) has *imperatival* force – it is used to make commands.What is the difference between assertoric and imperatival force? To put it another way, what is the difference between an assertion and a command? In the literature, this is sometimes described as a difference of "direction of fit". An assertion aims to describe the world; a command to change it.[7]

2.1 Content-validity

The idea that both imperatives and indicatives have propositions as their contents gives us a very

5 This example (like others in this paper) brushes over a difficulty in determining the tense and aspect of clauses in the imperative mood. (2) speaks of an attack to take place in the future, so perhaps (1) should be "You will attack at dawn." I have ignored this because to consistently get it right would involve distracting and controversial excursions into the syntax and semantics of tense in natural language. For example, in (A1), "the weather is fine" speaks of the future, in spite of being in the present tense; this could be regarded as an implicit anaphora, with (A1) abbreviating "Attack at dawn if the weather is fine *at that time*!" To sort this out for every example in the text would be well beyond the scope of this paper.

6 Or, to accommodate "de se" content, sets of centered worlds, where a centered world is a possible world together with a time, a speaker, and perhaps other parameters. For the purposes of this paper, I will be speaking of propositions as "sets of worlds", leaving it open whether "world" means possible world or centered world.

7 I have more to say about the differences between assertions and commands in section 3.5. To foreshadow: on my view assertions and commands differ in what mental states they express – assertions express beliefs, commands express intentions – and it is these mental states that differ in direction of fit.

powerful tool to solve the problem of imperative consequence. For propositions themselves can stand in logical relationships. A collection of propositions $p_1...p_n$ *entail* a proposition q iff there is no possible world that is a member of each of $p_1...p_n$ and not a member of q. This is simply an exotic way of restating the traditional criterion of validity: a possible world is a member of a proposition iff that proposition is true at that world, so the definition of entailment above is just saying that $p_1...p_n$ *entail* q iff it can't be that each of $p_1...p_n$ is true and q false.

> An argument is *content-valid* iff the contents of its premises jointly entail the content of its conclusion.

We might boldly conjecture that an argument is valid iff it is content-valid; and indeed all of the examples given in the introduction would confirm this. Content-validity is a general criterion of validity, because both imperatives and indicatives have contents; and it is conservative (in the sense given above) because it is equivalent to the traditional definition of validity in the special case of arguments consisting only of indicatives.

The same idea also enables us to solve an important subsidiary problem: how can there be logically complex imperatives, given that imperatives are not truth-apt. The classical logical connectives ("and", "or", negation, and the conditional) are normally thought to be truth-functional. For example, "and" pastes two subsentences together to make a complex sentence that is true iff both of the subsentences are true. But "and" cannot be doing that in "Attack at dawn and take no prisoners!" for neither "Attack at dawn!" nor "Take no prisoners!" are truth-apt.

The solution to this problem is similar to the definition of content-validity given above. The classical logical connectives can be thought of as operators on contents in a way that produces the truth-functional behaviour described above as a special case. So, for example, "and" pastes two sub-sentences together to make a complex sentence whose content is the set of all worlds that are members of the contents of both of the sub-sentences.

It's worth giving a precise statement of how this works for each of the classical logical connectives, as this will make a difference to the argument later. Writing "V(φ)" for "the content of φ":

$V(\ulcorner \varphi \wedge \psi \urcorner)$	is the set of all worlds w such that $w \in V(\varphi)$ and $w \in V(\psi)$
$V(\ulcorner \varphi \vee \psi \urcorner)$	is the set of all worlds w such that $w \in V(\varphi)$ or $w \in V(\psi)$
$V(\ulcorner \neg \varphi \urcorner)$	is the set of all worlds w such that $w \notin V(\varphi)$
$V(\ulcorner \varphi \rightarrow \psi \urcorner)$	is the set of all worlds w such that $w \notin V(\varphi)$ or $w \in V(\psi)$

We have covered a lot of ground in a short section here, so a summary is appropriate. We have seen a way of construing indicatives and imperatives as both having propositions as their contents, and we have seen a criterion of validity (content-validity) that pays attention only to the content of a sentence (and not to its force). We have also seen a way of understanding the logical connectives so that they too pay attention only to the content of a sentence. We might start to feel attracted to the following slogan: "logical relations are relations between propositions, not sentences". Validity is a logical relation; so validity must be something like content-validity, a relation that only pays attention to the propositions expressed by the sentences that make up an argument, not to any other features of those sentences.

2.2 Fallacies of mood

Unfortunately content-validity is not sufficient for validity. There are intuitively invalid arguments that are content-valid. Here are two examples:

(Z1) You attack at dawn.
*Therefore (Z2) Attack at dawn!

(Y1) Attack at dawn if the weather is fine!
(Y2) Let the weather be fine!
*Therefore (Y3) You attack at dawn.

The rough-and-ready test shows these to be invalid. In the case of (Z) if I predict that you will attack at dawn, but refuse to command so, I am making no logical mistake. In the case of (Y), I want you to attack if the weather is fine; and I further command that the weather be fine. I am making no logical mistake, however, if I refuse to predict that you will attack at dawn. I might simply doubt that you (or the weather!) will carry out my orders.

(Y) and (Z) both also have the feature that if the imperative mood were removed from them – that is, if each imperative clause was replaced by the corresponding indicative – the result would be a valid argument:

(Z1) You attack at dawn.	⇒ You attack at dawn.
*Therefore (Z2) Attack at dawn!	⇒ Therefore, you attack at dawn.
(Y1) Attack at dawn if the weather is fine!	⇒ You attack at dawn if the weather is fine.
(Y2) Let the weather be fine!	⇒ The weather is fine.
*Therefore (Y3) You attack at dawn.	⇒ Therefore, you attack at dawn.

I call arguments such as these, which are invalid, but can be transformed into valid arguments by changing only the mood, *fallacies of mood*. The existence of these arguments shows straightforwardly that validity is not content-validity. If we are to give an adequate criterion of validity, it will have to somehow pay attention to more than the mere contents of the sentences that make up an argument.

2.3 Codifying the fallacies of mood

The existence of fallacies of mood shows that content-validity is not sufficient for validity. We might continue to hope, however, that it is necessary. If that's so, then we should be able to adequately define validity by saying that a valid argument is one that is content-valid, and which is not a fallacy of mood. For this to work, we need a criterion for an argument's being a fallacy of mood which does not itself appeal to validity (that would make our proposed definition of validity circular). A natural thought is that what makes an argument a fallacy of mood is (what else?) the mood of the premises and conclusion. What we need to discover is what patterns of moods in premises and conclusions make a content-valid argument valid. I call this the programme of *codifying the fallacies of mood*.

As an aid to doing this, I will introduce the following notation, derived from (Smart:1984, p. 16). Let the *logical form* of a sentence be a notation of the form $\alpha(\varphi)$ where α is the letter I for an imperative, or

the letter A for indicative, and where φ is a translation of the content of that sentence into a suitable logical notation (for purposes of this paper, propositional logic). Let the *argument form* of an argument be an assignment of logical forms to its premises and conclusion using a consistent vocabulary to represent the contents.

As an example, here is argument (A) with its argument form shown at right:

(A1) Attack at dawn if the weather is fine!	I(p→q)
(A2) The weather is fine.	A(p)
*Therefore (A3) Attack at dawn!	I(q)

Looking at the form of (A), we can see that it is formally valid (and thus content-valid) because the contents of the premises (p→q, p) jointly entail the content of the conclusion (q). We can also see all of the information that could possibly be relevant to determining whether it is a fallacy of mood. For according to the standard view, all that mood does is to determine whether a sentence commands or asserts its content; and that information is shown (by the letters outside the parenthesis) in the argument form. A codification of the fallacies of mood would consist of a specification of which of these argument forms are fallacies of mood.

2.4 Hare's criterion

Here is an example of an attempt to codify the fallacies of mood, drawn from the work of R.M. Hare (Hare:1952, §2.5):

> An argument is *Hare-valid* iff *(a)* it is content-valid, and, *(b)* if the conclusion is indicative, then the contents of the indicative premises jointly entail the content of the conclusion, and, *(c)* if the conclusion is imperative, then the premises include at least one imperative.

Hare's criterion deals correctly with all the arguments we have considered so far. Clauses (b) and (c) are an attempt to codify the fallacies of mood – according to Hare's criterion, a fallacy of mood is an argument that fails to satisfy one or other of (b) or (c).

To understand how this works, let's look again at (Y) and (Z), with their argument forms:

(Z1) You will attack at dawn.	A(p)
*Therefore (Z2) Attack at dawn!	I(p)
(Y1) Attack at dawn if the weather is fine!	I(p→q)
(Y2) Let the weather be fine!	I(p)
*Therefore (Y3) You attack at dawn.	A(q)

The criterion rightly judges (Y) and (Z) invalid. (Y) fails to satisfy clause (b): it has an indicative conclusion, but the conclusion does not follow from its indicative premises alone (it has none). (Z) fails to satisfy clause (c): it has an imperative conclusion, but no imperative premises. It also rightly judges (A) valid.

Unfortunately, some arguments are Hare-valid but not intuitively valid by the rough-and-ready test. Notice, for example, that any argument that fails to satisfy clause (c) can be made to satisfy it by adding

an imperative premise, even if that premise is irrelevant to the argument. Argument (X), below, results from adding such a premise to argument (Z). (X) is surely as invalid as (Z) is, but (X) is Hare-valid.

(X1) The weather is fine.	A(p)
(X2) Attack at dawn!	I(q)
*Therefore (X3) Let the weather be fine!	I(p)

We might attempt to patch Hare's criterion against this counterexample in the following way. Say that a premise is *non-redundant* iff removing that premise results in an argument that is content-invalid.

> An argument is *revised Hare-valid* iff *(a)* it is content-valid, and, *(b)* if the conclusion is indicative, then the contents of the indicative premises jointly entail the content of the conclusion, and, *(c)* if the conclusion is imperative, then the non-redundant premises include at least one imperative.

A more controversial counterexample – this time to clause (b) – is due to Peter Geach (Geach:1958, p. 52). The general idea is this: (A) is an instance of modus ponens, allowing us to derive an imperative from an indicative and a conditional. If that's valid, shouldn't the corresponding modus tollens argument be valid, deriving an indicative from an imperative and a conditional? Suppose that the king says to one of his retainers, "If you are a faithful subject, rise up, Sir George; but do not rise, stay on your knees, fellow!" The king is implying that George is not a faithful subject; Geach's suggestion is that he is doing so by offering a valid argument to that conclusion:

(G1) If you are a faithful subject, rise up, Sir George!	I(p→q)
(G2) But do not rise, fellow!	I(¬q)
*Therefore (G3) You are not a faithful subject.	A(¬p)

(G) is not Hare-valid; it fails to satisfy clause (b). But Geach urges that it is valid; and this judgement receives some support from our rough-and-ready test. The king would be making a logical mistake if he commanded both (G1) and (G2), but yet believed that George was faithful.

It's not clear what to do about (G). There's no reason to suppose that a Hare-like criterion could not be devised that would count (G) as valid. Alternatively, it could be argued that (G) is not, in fact, valid. To do either of these, however, would merely be to prolong the agony for Hare's programme – as I will now argue, there are counterexamples to every possible Hare-like criterion; the fallacies of mood cannot be codified in any way that is consistent with the standard view.

2.5 The uncodifiability of the fallacies of mood.

The counterexamples discussed in the previous section show that a certain criterion of validity – Hare's criterion – is not adequate. I will show that there can be no adequate criterion of validity consistent with the standard view. No counterexample alone can show this – what I have is a recipe for generating counterexamples – a recipe that always succeeds, if the standard view is correct.

I can show you two arguments that are of the same form (according to the standard view, in the sense of "form" described in section 2.3) but of which one is intuitively valid, and the other intuitively invalid (both of which intuitions are backed up by the rough-and-ready test described in section 1.1). Because the arguments are of the same form, any criterion that is consistent with the standard view must treat them the same – must say either that both are valid or both invalid. Either way, that criterion –

whatever it is – suffers a counterexample.

Consider argument (B) below:

(B1) Flee for your life only if our plans are betrayed!

(B2) You flee for your life.

*Therefore (B3) Let our plans be betrayed!

(B) is invalid. Let us apply the rough-and-ready test. In uttering (B1) I tell you to flee for your life only if our plans are betrayed. I then (B2) predict that you will flee for your life, affirming the antecedent of (B1). But that is not tantamount to a command on my part that our plans be betrayed. I would not be making a logical mistake if I was not willing to command this. I might simply pessimistically believe that you will fail to follow my orders.

Notice that (B1) is an "only if" conditional. I take it that "only if" conditionals are the converses of the corresponding "if" conditionals: "It is raining only if the streets are wet" is equivalent to "The streets are wet if it is raining".[8] So, writing p for the proposition that you flee for your life, and q for the proposition that our plans are betrayed, the form of (B) is as shown below:

(B1) Flee for your life only if our plans are betrayed!	I(p→q)
(B2) You flee for your life.	A(p)
*Therefore (B3) Let our plans be betrayed!	I(q)

But then (B) is of the same form as (A), which we thought was valid. Here is (A), with its form again:

(A1) Attack at dawn if the weather is fine!	I(p→q)
(A2) The weather is fine.	A(p)
*Therefore (A3) Attack at dawn!	I(q)

Since the argument forms contain (if the standard view is true) all the information required to determine whether an argument is a fallacy of mood, and both (A) and (B) are content-valid, then either both (A) and (B) are valid, or both are invalid. But (A) is valid, and (B) is invalid, so the standard view is false.

Argument (B) is the "evil twin" of an argument we had already agreed was valid. But not only do valid arguments containing "if" have evil twins containing "only if", valid arguments containing "only if" have evil twins containing "if", as well. Remember argument (Y), which was an example of a fallacy of mood:

8 More generally, in an "if"-conditional, the subordinate clause is the antecedent and the main clause the consequent; in an "only if" conditional, the subordinate clause is the consequent and the main clause the antecedent. The subordinated clause is the one that appears immediately after the word "if". In English (unlike propositional logic) the order of clauses in a conditional is irrelevant to which is the antecedent and which the consequent; also irrelevant is whether the word "then" appears. Thus:

"If p, then q" is formalised...	p→q
"p if q" ...	q→p
"p only if q" ...	p→q

Subordinated clauses cannot occur in the imperative mood, which fact is responsible for the widespread but mistaken view that conditionals cannot have an imperative antecedent. It is possible to put an imperative into the antecedent of a cxonditional by using "only if", as I have done here.

(Y1) Attack at dawn if the weather is fine!	I(p→q)
(Y2) Let the weather be fine!	I(p)
*Therefore (Y3) You attack at dawn.	A(q)

(Y) is the evil twin of the valid argument (I), below:

(I1) Flee for your life only if our plans are betrayed!	I(p→q)
(I2) Flee for your life!	I(p)
*Therefore (I3) Our plans are betrayed.	A(q)

(I) is intuitively valid. If I commanded both (I1) and (I2), I would be making a logical mistake if I did not believe that our plans had been betrayed. (Y), as we said before is invalid. But both are of the same argument form according to the standard view, so (I) is valid if and only if (Y) is.

Hare's criterion (for what it's worth) wrongly categorises (B) as valid, and (I) as invalid. But even if the criterion were patched, to, for example, categorise (I) as valid, that could only be at the cost of wrongly categorising (Y) as valid. There is no patch possible: any criterion that is consistent with the standard view, and thus with the conception of argument forms described in section 2.3 must be inadequate.

3 Preposcription semantics

What is the difference between (A) and (B), or between (I) and (Y)? I think that they are different in form (and the problem with the standard view is that it has no room for the difference in question). The difference is very clearly marked in English: in (A1), the consequent, but not the antecedent, of the conditional is in the imperative mood; in (B1), the antecedent, but not the consequent, of the conditional is in the imperative mood. The same is true of (I) and (Y).

Suppose we represented the imperative mood as a sentential operator, written !, and capable of taking narrow scope with respect to a conditional.[9] Then we could write the forms of (A) and (B) as follows:

(A1) Attack at dawn if the weather is fine!	p→!q
(A2) The weather is fine.	p
*Therefore (A3) Attack at dawn!	!q

9 Treating the imperative as a sentential-like operator was first proposed by Hofstadter and McKinsey (Hofstadter/McKinsey:1939). They however, do not allow this operator to take narrow scope with respect to a normal conditional. This is because their imperativising operator is not, strictly speaking, sentential – it does not produce a sentence, taking sentences as parameters (as the logical connectives of the propositional calculus do) but produces an "imperative" taking a "sentence" as a parameter. For Hofstader and McKinsey imperatives are not sentences, and formulae in which they are used as such are not well formed. To deal with cases like (A1), they introduce a new set of operators that produce imperatives taking imperatives as parameters. This represents the logical vocabulary of natural languages ("if" "or" "and" etc.) as ambiguous in a way that it is not. The first treatment of the imperative mood as a genuine sentential operator was given by Chellas (Chellas:1969). He treats the imperative mood analogously to the obligation operator of a deontic logic and gives a possible worlds semantics. The price Chellas pays for this is that imperatives, in his formal language, are truth-apt! Chellas does not claim that natural language imperatives are truth-apt, though – this is simply a logically irrelevant difference between natural language and his formal system, he says (Chellas:1969, p. 3). I aim to do better: like Chellas, my imperativising operator is genuinely sentential; like Hofstadter and McKinsey, my semantics does not (except in weird corner cases) assign any imperative a truth-value.

(B1) Flee for your life only if our plans are betrayed!	!p→q
(B2) You flee for your life.	p
*Therefore (B3) Let our plans be betrayed!	!q

The difference shown here between the forms of (A1) and (B1) is justified by the difference noted above having to do with their English grammar. In (A1), p stands for "the weather is fine", which occurs in the indicative mood, and should be outside the scope of the imperative operator. In (B1) p stands for "you will flee for your life", which occurs in imperative mood and should be inside. As before, I assume that "only if" conditionals are the converses of "if" conditionals. The difference in validity between (A) and (B) is elegantly explained by this move. (A) is an instance of modus ponens, and so must be valid; (B) is not, and so its validity is up for grabs.[10]

The trouble with this is that it is inconsistent with the philosophical underpinnings of the standard view. Logical operators, such as the conditional, are, according to the standard view, operators on propositions; and propositions are not in themselves either imperatival or assertoric – the difference between assertion and command is a difference of force, not content. It doesn't make sense, therefore, for a conditional to have an imperatival consequent. The only way the standard view can deal with (A1) is to suppose that the occurrence of the imperative mood anywhere in a sentence marks that whole sentence as imperatival. That is just what we supposed implicitly in the previous section, and it led us straight into the problem of (A) and (B).

My proposal is to retain as much as possible of the standard view, while rejecting only those parts that led to trouble. I will retain: (a) that the content of an indicative is its truth conditions; (b) the content of an imperative is its compliance conditions; (c) classical logic; (d) that logical operators are operators on content; (e) that indicatives are truth-apt but imperatives are not. I will reject: (f) that the difference between assertion and command is a difference of force. In my view, the difference between assertion and command is a difference of content; contents may be intrinsically assertoric or intrinsically imperatival, or neither; and the imperative mood is a logical operator, operating on contents, and therefore able to take narrow scope.

In order to do this, we need to replace propositions as the contents of sentences. Propositions – sets of possible worlds, recall – aren't fine grained enough to act as contents under the new regime, for a given proposition doesn't carry enough information to determine whether it is assertoric or imperatival. I will replace them with another kind of set-theoretic construction out of possible worlds, called a preposcription.[11]

---

10 This syntactic move, of treating the imperative mood as a sentence operator, does not solve the problem of imperative consequence by itself – a semantic treatment of this operator is also needed. As well as the philosophical problems mentioned in the text, below, there is the purely formal problem that some argument forms containing the imperative mood – !(p∧q) therefore !p, for example – are intuitively valid but "up for grabs" in the same sense that the intuitively invalid (B) is – they are not instances of a classically valid argument form.

11 The semantic theory that follows owes a debt to Allan Gibbard's norm expressivism (Gibbard:1990, chapter 5). Like me, Gibbard is concerned to combine non-cognitivism (in his case, about morality) with the thesis that sentences that are not truth-apt can entail one another; like me he does this by using a possible worlds semantics on which the points of evaluation are not possible worlds (as they would normally be), but pairs of a world with something else. The key formal difference between Gibbard's approach and mine (setting aside the philosophical difference that Gibbard is interested in moral statements and I am interested in imperatives) is in what that "something else" is: for Gibbard it is a "system of norms"; for me, a second world. That makes it possible for me to give a formal analysis of the imperative mood, by making it a logical operator that looks at the internal structure of the points of evaluation.

3.1 What is a preposcription?

A preposcription is a set of *pairs* of possible worlds. Suppose, for the sake of example, that there are just four possible worlds:

- w_{fa} in which the weather is fine and you attack at dawn
- w_{f*} in which the weather is fine and you do not attack at dawn
- w_{*a} in which the weather is not fine and you attack at dawn
- w_{**} in which the weather is not fine and you do not attack at dawn

A handy of visualising a preposcription, then, is in a matrix[12] such as the one below:

	w_{fa}	w_{f*}	w_{*a}	w_{**}
w_{fa}	✓	✓	✓	✓
w_{f*}	✓	✓	✓	✓
w_{*a}				
w_{**}				

"The weather is fine."

Matrices like this one represent sets of pairs of worlds. Each cell represents the pair whose first member heads its row and whose second member heads its column. The ticked cells are in the set represented; the others are not.[13] The label says that this is the content of "The weather is fine". I will explain why shortly.

Think, for the moment, of a preposcription as a representation of either the truth conditions of an assertion or the compliance conditions of a command. The *first* member of each pair (which is to say, the worlds shown at the *left* of the matrix, heading the *rows*) has to do with truth conditions; the *second* member of each pair (which is to say, the worlds shown at the *top* of the matrix, heading the *columns*) has to do with compliance conditions.

The content of an assertion, such as "The weather is fine" is the set of all pairs such that that sentence is true at the *first* member of the pair; that is, its matrix consists of a row of all ticks alongside each world at which the assertion is true – as shown in the matrix above. The content of a simple command, such as "Attack at dawn", on the other hand, is the set of all pairs such that that sentence is complied with at the *second* member of the pair.

That is, writing V(φ) for "the content of φ":

12 Matrices like these were introduced by Robert Stalnaker (Stalnaker:1984), as a way of representing "propositional concepts". Though Stalnaker's propositional concepts are formally similar to my preposcriptions, they play a quite different role in his paper. Also, Stalnaker's propositional concepts are functions from pairs of worlds to truth values, and so he writes T and F, where I would put a tick or a blank, respectively. Though I have no objection to trading in sets of things for functions from those things to boolean values, it would be quite misleading to identify tick and blank with truth and falsity, as we will see below.

13 These matrices are consistent with there being more than four worlds – you can either make the idealising assumption that only these four worlds are possible, or think of each column and row as abbreviating an infinite number of columns or rows containing the same pattern of ticks and blanks.

V(φ), where φ is an simple indicative, is the set of all pairs (w,w') such that φ is true at w.

V(φ), where φ is a simple imperative, is the set of all pairs (w,w') such that φ is complied with at w'.

"Simple" here means not containing any propositional logical connective. I am assuming here that all simple indicative sentences are assertions, with truth-conditions, and that all simple imperatives are simple commands with compliance conditions – features (a) and (b) of the standard view.

So the content of (A3) "Attack at dawn!" is the set of all (w,w') such that you attack at dawn in w'. This is shown in the following matrix:

	w_{fa}	w_{f*}	w_{*a}	w_{**}
w_{fa}	✓		✓	
w_{f*}	✓		✓	
w_{*a}	✓		✓	
w_{**}	✓		✓	

"Attack at dawn!"

Notice that the matrix for "The weather is fine" has all its columns the same – each row is either all ticks or all blanks; and the matrix for "Attack at dawn!" has all rows the same – each column is either all ticks or all blanks. This is the feature that distinguishes an assertoric from an imperatival preposcription:

A preposcription p is *assertoric* iff for each pair (w,w'), (w,w') ∈ p iff (w,w) ∈ p.
(i.e. iff every column of the matrix for p is identical).

A preposcription p is *imperatival* iff for each pair (w,w'), (w,w') ∈ p iff (w',w') ∈ p.
(i.e. iff every row of the matrix for p is identical).

Assertoric preposcriptions can be said to be true or false, either relative to a world, or absolutely:

p is *true at a world* w iff every pair of worlds (w,w') ∈ p.
(i.e. iff the matrix for p has a row of all ticks headed by w)

p is *false at a world* w iff every pair of worlds (w,w') ∉ p.
(i.e. iff the matrix for p has a row of all blanks headed by w)

p is *true* iff it is true at the actual world.

p is *false* iff it is false at the actual world.

So the preposcription expressed by "The weather is fine" is true at w_{fa} and w_{f*} (the worlds where the weather is fine) and false at w_{*a} and w_{**} (the world where the weather isn't fine) just as we would like. These definitions apply to all preposcriptions, by the way, imperatival ones included. Imperatival preposcriptions, however, are neither true nor false[14] – there is no row in the matrix for "Attack at

14 With two exceptions: the set of all pairs of worlds (the matrix with all cells ticked) and the empty set (the matrix with all cells blank). These preposcriptions are both imperatival and assertoric – as makes sense, because both can be expressed either by a simple indicative ("The weather or not fine." / "The weather is fine and not fine." respectively) and or by a simple command ("Attack or do not attack!" / "Attack and do not attack!"). So without exception, every preposcription that is imperatival but not assertoric is neither true nor false.

dawn" with either all ticks or all blanks – so there is no world at which it is either true or false. This agrees with the tradition that holds that imperatives are not truth-apt[15] – feature (e) of the standard view.

3.2 Conditional commands

The sharp-eyed reader will have noticed that it follows from the definitions above that some preposcriptions are neither assertoric nor imperatival. These preposcriptions are not the content of any assertion, nor the content of any command. To explain why this is, I must digress – away from preposcriptions and back to the problems with the standard view.

According to the standard view, every command is a command to the effect that a certain proposition be true. I will call such commands *simple commands*; so, according to the standard view, every command is a simple command. Now, if this is so, then a conditional imperative, such as (A1), must be a command to the effect that a conditional be true. That is to say, on the standard view, (A1) is equivalent to (A1') below:

A1) Attack at dawn if the weather is fine!	$p \rightarrow !q$
A1') Let this be true: you attack at dawn if the weather is fine!	$!(p \rightarrow q)$

Since (B1) is a conditional imperative as well, on the standard view, it is equivalent to (B1'):

B1) Flee for your life only if our plans are betrayed!	$!p \rightarrow q$
B1') Let this be true: you flee for your life only if our plans are betrayed!	$!(p \rightarrow q)$

If that were all true, then it would have the effect of collapsing the distinction of form I tried to draw between arguments (A) and (B). For then both would be of the form $!(p \rightarrow q)$, p, therefore !q. Not every command, therefore, is a simple command. In particular, conditionals in which the imperative mood takes narrow scope are not simple commands, but irreducibly *conditional commands*.[16]

The view that conditional imperatives are conditional commands, and not simple commands that a conditional be true, can be independently motivated. Consider the following example:

3) Run if you see the monster!

4) See the monster only if you run!

On the standard view, these are equivalent – both of them are ways of commanding that the conditional "if you see the monster, then you run" be true. I think that they are different: (3) asks you to conform your running to your seeing; (4) asks to conform your seeing to your running. (3) commands you as regards running, but leaves it up to you as regards seeing; (4) commands you as regards seeing, but leaves it up to you as regards running. I think that (3) and (4) are both conditional commands, and are

15 It may be objected that I am here confusing truth-aptness with bivalence – with being either-true-or-false (perhaps, for example, "This man is bald", said of a borderline case of baldness, is truth-apt, but not bivalent). I agree that truth-aptness is different from bivalence. To say that a sentence is not truth-apt is to say more that than it is neither true nor false. I'm not quite sure how to describe the difference except by giving contentious examples like the borderline case one. The crucial point of reply to this objection, though, is that my semantics is supposed to be *consistent* with imperatives not being truth-apt. It doesn't have to *entail* that they are not truth-apt. Whatever else a sentence must do to avoid being truth-apt, it must at least be neither true nor false, and I deliver that.

16 The idea that conditional imperatives express conditional commands is much discussed in the literature on indicative conditionals. See for example (Edginton:1995, pp. 287-290). There is a classic discussion in (Dummett:1981, pp. 339-341).

semantically different both from each other and from (5), below:

5) Let it be that either you do not see the monster, or you run!

(5) is a simple command – indeed, it is equivalent, on the standard view, to the simple command that the proposition "If you see the monster, then you run" be true. It leaves it up to you whether to run or to take care not to see the monster.

There are many different ways of incorporating conditional commands into a semantics of imperatives.[17] I suggest that we generalise the concept of compliance conditions slightly: conditional commands have *variable compliance conditions*. What they tell you to do depends on how things are. The compliance conditions of (3) depend on whether you see the monster: if you do, then to comply you must run; if you do not, then (3) is trivially complied with. Where we thought of a compliance condition as a set of possible worlds, you can think of a variable compliance condition as a function from possible worlds to compliance conditions. The content of (3) is the function that takes each world in which you see the monster to the set of worlds in which you run, and takes every world in which you do not see the monster to the set of all worlds.

This takes us back to preposcriptions. The preposcriptions that are neither assertoric nor imperatival are the contents of imperative conditionals, and encode their variable compliance conditions. Here, for example, is the preposcription expressed by (A1):

	w_{fa}	w_{f*}	w_{*a}	w_{**}
w_{fa}	✓		✓	
w_{f*}	✓		✓	
w_{*a}	✓	✓	✓	✓
w_{**}	✓	✓	✓	✓

"Attack at dawn if the weather is fine!"

(A1), like (3), has variable compliance conditions – what it tells you to do depends on how things are. This is represented by the matrix shown above for (A1): each row represents a different possible compliance condition; (A1) varies its compliance conditions depending on which world is actual; the compliance conditions that are active at a world are those shown in the row headed by that world.

3.3 Logical connectives

The classical logical connectives – "and", "or", "not", and the conditional – are operators on preposcriptions. They work in a similar way to the way that they did on the standard view: a pair should be a member of the content of a conjunctive sentence iff it is a member of the contents of both conjuncts, and so on: This means that the logical connectives have the same effect on the truth-conditions of assertions and on the compliance-conditions of simple commands that they did on the

17 A (seemingly) much simpler proposal than mine, for example, would hold that conditional commands are dyadic speech acts, each having two propositional contents: the compliance conditions of the command (e.g. that you attack at dawn), and the antecedent conditions under which the command is active (that the weather is fine). This, however, fails to generalise to more complicated cases, in which imperative conditionals are cojoined, disjoined, or used with "otherwise", "only if" or "unless".

standard view: a conjunction of two assertions is true at the worlds where both of its conjuncts are true; a conjunction of two commands is complied with at the worlds where both of its conjuncts are complied with; and so on.

$V(\ulcorner \varphi \wedge \psi \urcorner)$ is the set of pairs of worlds (w,w') such that (w,w') ∈ V(φ) and (w,w') ∈ V(ψ)
$V(\ulcorner \varphi \vee \psi \urcorner)$ is the set of pairs of worlds (w,w') such that (w,w') ∈ V(φ) or (w,w') ∈ V(ψ)
$V(\ulcorner \neg \varphi \urcorner)$ is the set of pairs of worlds (w,w') such that (w,w') ∉ V(φ)
$V(\ulcorner \varphi \rightarrow \psi \urcorner)$ is the set of pairs of worlds (w,w') such that (w,w') ∉ V(φ) or (w,w') ∈ V(ψ)

We can also think of the imperative mood is also an operator on preposcriptions. It converts an assertoric preposcription into an imperatival one that is complied with at all and only those worlds that the original preposcription was true at. It takes the value of each cell across the diagonal of the input matrix, and reproduces those cells across each row of the output matrix.[18] To put it more carefully:

$V(\ulcorner !\varphi \urcorner)$ is that set of pairs of worlds (w,w') such that (w',w') ∈ V(φ)

The diagram below shows the relationship between "You will attack at dawn." (an assertion) and "Attack at dawn!" (the corresponding simple command). I have highlighted the cells across the diagonal:

	w_{fa}	w_{f*}	w_{*a}	w_{**}
w_{fa}	✓	✓	✓	✓
w_{f*}				
w_{*a}	✓	✓	✓	✓
w_{**}				

"You will attack at dawn."

	w_{fa}	w_{f*}	w_{*a}	w_{**}
w_{fa}	✓		✓	
w_{f*}	✓		✓	
w_{*a}	✓		✓	
w_{**}	✓		✓	

"Attack at dawn!"
i.e. "!(You will attack at dawn)"

Entailment can be defined between preposcriptions:

$p_1,...p_n$ jointly *entail* q iff every member of each of $p_1,...p_n$ is also a member of q
(i.e. every cell ticked in the matrices for all of $p_1,...p_n$ is also ticked in the matrix for q)

This means that entailment relations between assertoric preposcriptions agree with the truth-conditional definition of entailment I gave in the context of the standard view: where p and q are both assertoric, p entails q iff it can't be that p is true and q false; where p and q are both imperatival, p entails q iff it can't be that p is complied with and q not complied with.

Now let's look at a conditional command. I suggested above that (A1) is of the form p→!q. We can thus use the rules given above to determine its content. (A1) is a conditional, and its content is shown below together with the contents of its antecedent and consequent:

18 In the terminology of two-dimensional modal logic, the imperative operator is the "Stalnaker dagger" (Stalnaker:1984 p. 82) applied to preposcriptions.

	W_{fa}	W_{f*}	W_{*a}	W_{**}
W_{fa}	✓	✓	✓	✓
W_{f*}	✓	✓	✓	✓
W_{*a}				
W_{**}				

"The weather is fine." (p)

	W_{fa}	W_{f*}	W_{*a}	W_{**}
W_{fa}	✓		✓	
W_{f*}	✓		✓	
W_{*a}	✓		✓	
W_{**}	✓		✓	

"Attack at dawn!" (!q)

	W_{fa}	W_{f*}	W_{*a}	W_{**}
W_{fa}	✓		✓	
W_{f*}	✓		✓	
W_{*a}	✓	✓	✓	✓
W_{**}	✓	✓	✓	✓

"Attack at dawn if the weather is fine!" (p→!q)

Notice that the preposcription matrix for (A1) that we arrived at by applying the semantic rules for the connectives is the same as the one that we arrived at in section 3.2 by the informal method of filling in the matrix using our intuitive grasp of (A1)'s variable compliance conditions.

3.4 Solving the problem of imperative consequence

We can now return to the problem of imperative consequence itself. If the contents of sentences are preposcriptions, then validity can be safely regarded as content-validity:

An argument is *valid* iff its the contents of its its premises jointly entail the content of its conclusion.

This criterion is general, because it applies across to arguments containing any mixture of imperative and indicative sentences. It is conservative, because, in the case of an argument consisting of just indicatives, all the preposcriptions in question are assertoric, so the requisite entailment relation will hold just if it can't be that the premises are true and the conclusion false. And finally, it is adequate with regard to all of the examples of this paper.

Let us check argument (A) for validity (I have here highlighted the cells that are ticked for both premises):

	W_{fa}	W_{f*}	W_{*a}	W_{**}
W_{fa}	✓		✓	
W_{f*}	✓		✓	
W_{*a}	✓	✓	✓	✓
W_{**}	✓	✓	✓	✓

(A1) "Attack at dawn if the weather is fine!"

	W_{fa}	W_{f*}	W_{*a}	W_{**}
W_{fa}	✓	✓	✓	✓
W_{f*}	✓	✓	✓	✓
W_{*a}				
W_{**}				

(A2) "The weather is fine."

	W_{fa}	W_{f*}	W_{*a}	W_{**}
W_{fa}	✓		✓	
W_{f*}	✓		✓	
W_{*a}	✓		✓	
W_{**}	✓		✓	

(A3) "Attack at dawn!"

Argument (A) is valid: there is no cell not ticked for (A3) that is ticked for both (A1) and (A2).

Now consider argument (B):

B1) Flee for your life only if our plans are betrayed! !p→q
B2) You flee for your life. p
B3) * Therefore, let our plans be betrayed! !q

Since the subject matter has changed, we need a new stock of worlds with which to label the matrices. To make this easier, I will adopt the convention of writing "w_{pq}" for "worlds where p and q" are true, "w_{p*}" for "worlds where p is true and q is false", and so on. The matrix for (B1) is then as follows:

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}	✓	✓	✓	✓
w_{p*}			✓	✓
w_{*q}	✓	✓	✓	✓
w_{**}			✓	✓

(B1) "Flee for your life only if our plans are betrayed!" (!p→q)

The specific pattern of ticks and blanks shown here were derived using the rules given above. Like (A1), (B1) is neither an assertion nor a simple command. It is a conditional command, and we can use the idea of variable compliance conditions to understand what it is saying. (B1) has the trivial compliance conditions for worlds where its consequent is true (i.e. worlds where our plans are betrayed); for worlds where its consequent is false, its compliance conditions are are those of the simple command "Do not flee for your life". In short, (B1) is equivalent to "If our plans are not betrayed, then do not flee for your life!"

We can now check the argument (B) for validity:

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}	✓	✓	✓	✓
w_{p*}			✓	✓
w_{*q}	✓	✓	✓	✓
w_{**}			✓	✓

(B1) !p→q

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}	✓	✓	✓	✓
w_{p*}	✓	✓	✓	✓
w_{*q}				
w_{**}				

(B2) p

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}	✓		✓	
w_{p*}	✓		✓	
w_{*q}	✓		✓	
w_{**}	✓		✓	

(B3) !q

(B) is invalid: the highlighted cells are ticked for both premises and blank for the conclusion. I invite you check for yourself the arguments (I), (Y), (Z), (D) and (F). You will see that (I) and (D) are valid, and the others are invalid, just as our intuitions judge, guided by the rough-and-ready test. The criterion of validity I have just described is adequate, conservative and general.

3.5 Imperassertion

I solved the the problem of imperative consequence by treating assertions and commands as differing in content, rather than in force. I said much about their contents – that those contents are preposcriptions, that they can be used to solve the problem of imperative consequence, and so on – but I have not yet said much about force. What is it that assertions and commands have in common? Here is my answer to that: You can think of an assertion as an attempt to influence or constrain your interlocutors' beliefs; and a simple command as an attempt to influence or constrain your interlocutors'

intentions.[19] An *imperassertion* is a generalisation of both, so that both assertions and commands are special cases of imperassertion.

If that were the whole story imperassertion would perhaps appear an objectionable gerrymander – a disjunction of assertion and command designed to create the illusion that two quite different types of speech act have something in common. But it is not: there are *proper imperassertions* – imperassertions that are neither assertions nor simple commands – and among these are the conditional commands discussed in section 3.2. So imperassertion is not a disjunctive gerrymander: it is more like a continuum with assertion and simple command at the extreme ends. And the idea that assertions and commands are on a continuum helps us to make sense of a phenomenon that would otherwise be quite difficult, namely, conditional commands.

An imperassertion, in general, is an attempt to influence or constrain your interlocutor's beliefs and intentions *together*. Some of the imperassertions I may make to you constrain your intentions only in a trivial way, and constrain your beliefs in the same way no matter what your intentions – those are assertions. Others constrain your beliefs only in a trivial way and constrain your intentions in the same way no matter what your beliefs – those are simple commands. The interesting cases of imperassertion aim to entangle your beliefs and intentions, so that certain combinations of belief and intention are ruled out. Take for example, the conditional command (A3) “Attack at dawn if the weather is fine”. On my view, this is a proper imperassertion that aims to entangle your beliefs concerning the weather with your intentions concerning attacking – that rules out the combination of a belief that the weather is fine with the lack of an intention to attack at dawn.

Suppose you are a soldier, and that you already believe that the weather will be fine. Then, when your commanding officer says to you “Attack at dawn if the weather is fine!”, her imperassertion and your beliefs are jointly inconsistent with your lacking the intention to attack at dawn. So, if the imperassertion succeeds in its communicative purpose, given your background beliefs, you will form the intention to attack at dawn. The proviso “if the imperassertion succeeds” is important – people are not hydraulic mechanisms to be shoved around by speech acts.

What about the case where you are a soldier, you have no beliefs about the weather, and your commanding officer says to you “Attack at dawn if the weather is fine!”? In this case what must happen is that you form some kind of lasting disposition to intend to attack, should you come to believe that the weather is fine. It is this disposition that I referred to above as “entangled” beliefs and intentions.

We can view a preposcription as a representation of a person's beliefs, intentions, and any entanglements thereof. Earlier I said that a preposcription is a set of pairs of worlds, of which the first member has to do with truth, and the second to do with compliance; now we will think of the first

19 I am thinking here of intentions as propositional mental states analogous to beliefs, and not reducible to a combination of desire and belief. Such a view of intentions has been developed at length by Michael Bratman. (Bratman:1999, especially chapters 2 and 3)
Also, I am departing from the tradition that holds that commands express, not intentions, but desires, as held, for example, by (Searle:1969, pp. 64-67). Consider the pacifist general, who combines no desire that any warlike acts occur with the intention to follow his own orders and command an attack. It seems to me that the pacifist general combines a lack of desire that an attack occur with an intention that an attack occur. This shows both that the intention that φ can occur without the desire that φ, and that commands express intentions, not desires, for when the general commands his soldiers “Attack!”, he does not express a desire that they attack, for he has no such desire.

member of the pair as having to do with belief; and the second to do with intention. Let us call a person P's complete belief-intention state (all their beliefs, intentions, and entanglements thereof), construed as a preposcription, that person's *belief-intention-set*. We can then recover P's beliefs and intentions from P's belief-intention-set in the following way:

A world w is *doxastically possible* for a person P iff there is some (w,w') in P's belief-intention-set.
(i.e. the matrix for P's belief-intention-set has a tick in the w row)

A world w' is *intentionally possible* for a person P iff there is some (w,w') in P's belief-intention-set.
(i.e. the matrix for P's belief-intention-set has a tick in the w' column)

P *believes* that φ iff φ is true at all worlds doxastically possible for P.

P *intends* that φ iff φ is true at all worlds intentionally possible for P.

When someone successfully makes an imperassertion to you, they constrain your belief-intention-set by causing you to exclude from it all pairs that were not members of the content imperasserted.[20] Let's look at an example, using the sentences (A1) and (A2), and the stock of worlds that were suitable for representing their contents. Suppose you start off maximally open-minded, with no non-trivial beliefs or intentions one way or the other concerning either the weather or attacking, and no entanglements between them, either. The only belief-intention-set consistent with such open-mindedness is the tautologous preposcription; the one that contains every pair of every world.

Then suppose that someone successfully imperasserts (A1) “Attack at dawn if the weather is fine” to you. Your belief-intention-set should now contain the intersection of your previous belief-intention-set with the content of (A1). Since your previous belief-intention-set was the set of all pairs of worlds, that means that your belief-intention-set is now identical to the content of (A1):

	w_{fa}	w_{f*}	w_{*a}	w_{**}
w_{fa}	✓		✓	
w_{f*}	✓		✓	
w_{*a}	✓	✓	✓	✓
w_{**}	✓	✓	✓	✓

Your belief-intention-set after (A1) is successful on you.

Consider what this means for your beliefs and intentions. You are still maximally open-minded as regards belief alone, for there is a tick in every row. You are also still maximally open-minded as regards intention alone, for there is a tick in every column. But your beliefs and intentions are now entangled – changes to which worlds are doxastically possible for you could change which worlds count as intentionally possible for you (or vice versa).

Suppose that someone now successfully imperasserts (A2) “The weather is fine” to you. Your belief-intention set should now contain the intersection of the content of (A2), for which see section 3.1, with the preposcription shown above; that is:

20 Here I again draw on Robert Stalnaker's excellent article (Stalnaker:1984); the account of imperassertion here advocated is intended to be a generalisation of his account of assertion. Note that I am using Stalnaker's work in two quite different ways in this paper: here, as a model for a theory of imperassertion; and throughout section 3 as the source of the matrix technique for visualising preposcriptions.

	w_{fa}	w_{f*}	w_{*a}	w_{**}
w_{fa}	✓		✓	
w_{f*}	✓		✓	
w_{*a}				
w_{**}				

Your belief-intention-set after (A1) and (A2) are successful on you.

In this state, you believe that the weather is fine (no tick in any row headed by a a world at which the weather isn't fine); you intend that you attack at dawn (no tick in any column headed by a world at which you don't attack); and your beliefs and intentions are no longer entangled.[21]

This concludes my positive account of imperatives, and my solution to the problem of imperative consequence. In the remainder of this paper I reply to a few objections, and contrast my solution to some of its rivals.

4 Objections and alternatives

4.1 Questions

The indicative mood is used to make assertions; imperative to make commands. But these two are not the only moods in English. Can preposcription semantics be extended to include the interrogative, used to make questions?

In short, no. Questions are very different from both assertions and commands. There is nothing analogous to truth conditions or compliance conditions that can play the role of the content of a question. Questions demanding a yes or no answer could have "affirmation conditions" perhaps – circumstances under which the answer "yes" is correct – but this approach doesn't apply to "who" "what", "why", or "how" questions. Also, there are no conditional questions. A question can involve a conditional, as in "If Oswald didn't kill Kennedy, who did?" But the effect of that is not to make a conditional question but to introduce some content that the recipient is supposed to presuppose when coming up with his or her answer.

My approach works for assertions and commands because they operate in a parallel way, expressing beliefs and intentions, respectively. But there is no type of mental state – "querying"? – expressed by all and only questions. Assertions and commands are more like each other than either are like questions.

21 I earlier glossed "entanglement" as the possibility of a change in doxastic possibility constituting a change in intentional possibility or vice versa. More precisely, a preposcription is *disentangled* iff, in its matrix, every non-blank row carries the same pattern of ticks, and every non-blank column carries the same pattern of ticks. A preposcription is *entangled* iff it is not disentangled. Blanking a row (or a column) of a disentangled preposcription cannot have the effect of blanking a column (or row, respectively) unless it blanks the entire matrix.

4.2 Ross's paradox

My semantics shares with the standard view an objection due to Alf Ross (Ross:1944), connected with argument (R), below:

(R1) Post the letter!

Therefore, (R2) Post the letter or burn the letter!

(R) is valid on my criterion of validity, as it is on Hare's criterion (in both cases, because the conclusion is complied with at every world at which the premise is). But, Ross insists, it is "obvious that this inference is not immediately conceived to be logically valid". (Ross:1944 p.38) Ross does not say why this is obvious, but perhaps he was thinking of a case where you hear (R1), infer (R2) from it, and, noticing that complying with (R2) is consistent with burning the letter, burn it, thereby violating (R1).

It seems to me that Ross's paradox is not specific to imperatives; it is really an objection to classical logic. Consider the following argument, (P), which consists only of indicatives and is classically valid:

(P1) The letter is important.

Therefore, (P2) The letter is important or the letter is a piece of rubbish.

Whatever features (R) has that make it seem invalid are shared by (P). Suppose you are tidying up my office and taking away the scrap paper to be burned. You find the letter in the middle of a pile of similar rubbish. I say (P1), you infer (P2), and, noticing that (P2) is consistent with the letter being rubbish, take it away to be burned.

Should this make us think that (P2) is invalid? No: just because I asserted (P1), and (P1) validly entails (P2), it does not follow that you should behave as if I asserted (P2) instead. The same is true for (R): just because I commanded (R1), and (R1) validly entails (R2), it does not follow that you should behave as if I commanded (R2). Ross's paradox gives us no reason to deny that (P) is valid; therefore, it gives us no reason to deny that (R) is.

It may be thought that (R) is invalid because (R2) gives you permission to burn the letter, while (R1) does not. There is something to this idea, but again (P) is a parallel case. (P2) tells you only that the letter is important or rubbish; but there would be no point in my telling you that if I believed that the letter was not rubbish. So it's reasonable for you to infer from my utterance of (P2) that I do not believe that the letter is not rubbish. Similarly, (R2) tells you only that you are to post or burn the letter (and not that you are permitted to burn it); but there would be no point in my telling you so if intended you not to burn it. So it is reasonable for you to infer from my utterance of (R2) that I do not intend that you not burn the letter. In both cases, an explicit disjunction carries a conversational implicature that its disjuncts do not, and this is reponsible for the illusion that the arguments are not valid.[22]

4.3 Some (surprisingingly) invalid argument forms

Argument (S) is an instance of disjunctive syllogism, and is valid on my semantics:

(S1) Attack at dawn unless the weather is fine! $!p \vee q$

22 J.J.C. Smart (Smart:1984, pp. 17-18) also holds that Ross's paradox turns on a conflation of conversational implicature with logical entailment.

(S2) The weather is not fine. ¬q

Therefore (S3) attack at dawn! !p

The following argument is not valid, however:

(P1) Attack at dawn or stand on your head! !p∨!q

(P2) You do not stand on your head. ¬q

*Therefore (P3) Attack at dawn! !p

The difference, of course, besides the slight change of subject matter, is in the use of the imperative mood in the first premise: "stand on your head" is imperative, "the weather is fine" indicative. The invalidity of (P) can be seen from the following matrices (I have highlighted the cells that are ticked for the premises and blank for the conclusion):

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}	✓	✓	✓	
w_{p*}	✓	✓	✓	
w_{*q}	✓	✓	✓	
w_{**}	✓	✓	✓	

!p∨!q

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}				
w_{p*}	✓	✓	✓	✓
w_{*q}				
w_{**}	✓	✓	✓	✓

¬q

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}	✓	✓		
w_{p*}	✓	✓		
w_{*q}	✓	✓		
w_{**}	✓	✓		

!p

Shouldn't (P) be valid though? Afterall, if you already know that you will not stand on your head, shouldn't you try to comply with (P1) by attacking? Maybe: but that doesn't show that (P) is valid. Employ the rough and ready test: for (P) to be valid, it would have to be that someone who commanded (P1) and asserted (P2) but did not wish to command (P3) was making a logical mistake. But it there is no logical mistake in my commanding you to attack or stand on your head, wishing you to make up your own mind which (and thus not being willing to command you to attack), and in believing that you will choose to attack rather than perform the headstand.

It does not matter, by the way, whether we take the scope of the imperative mood in (P1) to be narrow (as I did above) or wide. On my semantics, !(p∨q) is equivalent to !p∨!q. Also, since → is the material conditional, !(p→q) is equivalent to !(¬p∨q); so the argument form below,

!(p→q)

p

Therefore, !q

is also not generally valid. This form is of interest, because it is the form that the standard view believes (A) to have (the first premise being the simple command that a conditional be true, and the second premise being the assertion of its antecedent).

Another controversial argument is (O), below:

(O1) Send up a flare if you capture the enemy leader! p→!q

(O2) Capture the enemy leader! !p

*Therefore (O3) Send up a flare! !q

(O) differs formally from (A) in that its second premise is in the imperative mood, rather than the indicative. It is shown to be invalid by the following matrices:

	W_{pq}	W_{p*}	W_{*q}	W_{**}
W_{pq}	✓		✓	
W_{p*}	✓		✓	
W_{*q}	✓	✓	✓	✓
W_{**}	✓	✓	✓	✓

p→!q

	W_{pq}	W_{p*}	W_{*q}	W_{**}
W_{pq}	✓	✓		
W_{p*}	✓	✓		
W_{*q}	✓	✓		
W_{**}	✓	✓		

!p

	W_{pq}	W_{p*}	W_{*q}	W_{**}
W_{pq}	✓		✓	
W_{p*}	✓		✓	
W_{*q}	✓		✓	
W_{**}	✓		✓	

!q

Shouldn't (O) be valid? After all, if I tell you to send up a flare under some circumstances, and then tell you to make sure those circumstances obtain, aren't I implicitly telling you to send up a flare? No. Again apply the rough-and-ready test: there's no logical mistake in my commanding you to send up a flare if you capture the enemy leader, commanding you to capture the enemy leader, and not wanting to command you to send up a flare no matter what. On the contrary, to command you to send up a flare no matter what would defeat the purpose of making the conditional command (O1), which is to get you to conform your flare sending activities to whether you succeed in carrying out (O2).

4.4 Compliance gaps

Conditional commands, introduced in section 3.2, played an important role in motivating my solution to the problem of imperative consequence. I used the apparent difference between (3) and (4) to argue that imperative conditionals express conditional commands. Another argument – the *compliance gap* argument – to the same conclusion appeals to our intuitions about the compliance-conditions of imperative conditionals.

3) Run if you see the monster! p→!q

4) See the monster only if you run! !p→q

(3) is complied with if you see the monster and you run: the condition was satisfied and you did what the command said. (3) is violated if you see the monster and do not run: the condition was satisfied and you failed to do what the command said. But what if you do not see the monster? Then, says the argument, (3) is not complied with, for it did not tell you to do anything; but (3) was not violated either, for the same reason.

This suggests a gappy, or three-valued approach to compliance values and compliance-conditions. Let us say that an sentence is *avoided* iff it is neither complied with nor violated. The diagram below summarises the intuitions given above concerning the gappy compliance-conditions of (3) (writing C for comply, V for violate, A for avoid):

	W_{pq}	W_{p*}	W_{*q}	W_{**}
	C	V	A	A

"Run if you see the monster!" (p→!q)

In contrast, if we view (4) as forbidding you from seeing the monster should you not run, then (4) is

complied iff you do not run and not see the monster; violated iff you do not run and do see the monster, and avoided otherwise:

	w_{pq}	w_{p*}	w_{*q}	w_{**}
	A	V	A	C

"See the monster only if you run!" (!p→q)

This shows that (3) and (4) are not equivalent, for they have different compliance conditions. They cannot both be equivalent, therefore, to the simple command that the conditional "If you see the monster, then you run" be true. At least one of them is a conditional command. Therefore there are conditional commands.

I did not use this argument earlier for two reasons. First, because I am not in the business of conceptually analysing "compliance" and "truth" here. In my usage, compliance is a semi-technical concept, like truth, and part of the regimentation is that conditionals like (3) are to be regarded as complied with when their antecedents are false. The compliance gap argument pumps our intuitions in a direction that is contrary to what I propose to mean by "compliance". Second, because the compliance gap argument makes salient a very different approach to the problem of imperative consequence. On that approach, we would develop a three-valued imperative logic, with a non-classical (but extensional) treatment of the logical connectives.[23] This contrasts with my view, which combines a classical treatment of the logical connectives with an intensional imperativising operator.

I said above that I agree with the standard view in thinking that (3) is complied with if you do not see the monster. Here is a definition of compliance for preposcriptions that delivers that result:

p is *complied with at a world* w iff the pair (w,w) ∈ p.

p is *violated at a world* w iff the pair (w,w) ∉ p.

p is *complied with* iff p is complied with at the actual world.

p is *violated* iff p is violated at the actual world.

Besides lacking gaps, this definition has the consequence that true assertoric preposcriptions are complied with, and false ones violated. In fact, these are two faces of the same problem. The conditionals that the gappy view would count as avoided are precisely those that my semantics counts as true. My conditionals are material conditionals, so p→!q is true in the worlds where its antecedent is false and !p→q is true in the worlds where its consequent is true; those are just the worlds where the gappy view wanted to count them as avoided (check the tables above).

This suggests a way to recapture a sense of "compliance" that allows for gaps. To avoid confusion I will call the gappy notion of compliance "strong compliance":

p is *strongly complied with at a world* w iff p is complied-with-at w and is not true-at w.

p is *strongly violated at a world* w iff p is violated-at w and is not false-at w.

p is *strongly complied with* iff p is strongly complied with at the actual world.

p is *strongly violated* iff p is strongly violated at the actual world.

23 As developed by, for example, Peter Vranas (Vranas:2008).

Take (A1) for example:

	w_{fa}	w_{f*}	w_{*a}	w_{**}
w_{fa}	✓		✓	
w_{f*}	✓		✓	
w_{*a}	✓	✓	✓	✓
w_{**}	✓	✓	✓	✓

(A1) "Attack at dawn if the weather is fine!"

If the weather isn't fine, some would say, (A1) is not complied with – "there wasn't anything to comply with". On the other hand, it isn't violated either. I can recapture this gappy notion of compliance in the following way: (A1) is strongly complied with at w_{fa}, strongly violated at w_{f*}, and neither at the other two worlds.

4.5 Cognitivism

In section 3.5, I generalised Stalnaker's account of assertion to include commands. I am not the only one to have tried this. In (Lewis:1979), David Lewis added commands to Stalnaker's theory not by generalising assertions to include commands, but by assimilating commands to assertions. On Lewis's view, the command that such-and-such is equivalent to a report that such-and-such is commanded. The ingenious feature of Lewis's view is how he accounts for the fact that a mere report of a command can have the effect of commanding someone – this is done by a "rule of accommodation", an independently motivated principle of conversational dynamics. Considerations of space forbid a full discussion of Lewis's views here – but I would like to briefly discuss the general idea that forms the basis of Lewis's view of commands, that every command is equivalent to some assertion, and more specifically that a command that φ is equivalent to the report that I command that φ. Let us call this position *cognitivism*.

The term "cognitivism" is used in meta-ethics, to apply to theories of moral discourse on which moral statements are fact-stating, express beliefs, and are truth-apt. According to cognitivism, as we are now using the term, imperatives have all three of these features. "Attack at dawn!", according to cognitivism, states a fact about who commands what; expresses my belief that I am about to command you to attack at dawn; and is true iff I succeed in doing so. My own view (and the standard view) deserve to be called "non-cognitivist": "Attack at dawn!" does not state any fact; it expresses an intention, not a belief; and it cannot be either true or false.

Cognitivists have not, so far as I am aware, tackled the problem of imperative consequence. If they did, however, they would find that they had an elegant solution to it: to determine whether an argument containing imperatives is valid, simply translate all imperatives into the indicatives that report that the right sort of command has been made, and then check for validity in the traditional way.[24] To take a simple example:

24 There is also a non-cognitivist solution to the problem of imperative consequence that is parasitic on this cognitivist solution. This is the view that an argument is valid iff the result of translating its premises and conclusion into indicatives in the way that a cognitivist would is valid. The parasitic non-cognitivist denies that the translation preserves the meaning of premises and conclusion, but asserts that it preserves both validity and invalidity. This solution however suffers from the second and third objections to cognitivism, below.

C1) Attack at dawn and take no prisoners!

C2) Therefore, attack at dawn!

Argument (C) becomes (C') below:

C1') I command that you attack at dawn and take no prisoners.

C2') I command that you attack at dawn.

(C') seems to be valid, according to the traditional test. It can't be true that I commanded that you attack at dawn and take no prisoners, without it being true that I commanded that you attack at dawn. (We have to take "I command that..." as an indirect speech report, so that the truth of (C2') does not require that I ordered you to "attack at dawn" in just those words and no more. If that is impermissible, so much the worse for cognitivism).

The cognitivist's criterion of validity also says sensible things about the fallacies of mood. Take (Z) as an example:

Z1') You attack at dawn.

Z2') *Therefore, I command that you attack at dawn.

(Z') seems invalid, counter-exampled by the possibility that you attack without my commanding you to.

The cognitivist can also afford a theory of conditional commands. Conditional commands are, as on my view, expressed by narrow scope imperatives (see section 3.2). It is particularly easy, however, for the cognitivist to make sense of an imperative within the scope of a conditional, since on her view, imperatives are truth-apt. Here is the cognitivist's translation of argument (A):

A1') If the weather is fine, then I command that you attack at dawn.

A2') The weather is fine.

A3') Therefore, I command that you attack at dawn.

(A') is surely valid, and even an instance of modus ponens just as I claimed that (A) was. Moreover, the cognitivist does not fall into the fallacies that caused trouble for the standard view. Here's the cognitivist's translation of (B):

B1') If I command that you flee for your life, then our plans are betrayed.

B2') You flee for your life.

B3') * Therefore, I command that our plans be betrayed.

(B') is counter-exampled by the possibility in which our plans are betrayed, and you flee for your life, but in which I do not make any commands at all. That is a situation in which both its premises are true, and its conclusion false.

With all the virtues of cognitivism, why do I not accept it? For three reasons, here presented in order of increasing importance. First, because it is cognitivism. Some philosophers find it hopelessly implausible that imperatives are truth-apt. There are other theoretical costs to accepting cognitivism too: meta-ethical prescriptivism becomes impossible; we must explain how it is that a report of a command could have the conversational effect that a command does (not that cognitivists cannot do this, but they may incur further theoretical costs doing so – not every philosopher of language accepts

Lewis's theory of conversational score, with its principle of accomodation).

Second, the cognitivist's criterion of validity makes valid some arguments which would at least not strike an neutral observer as valid. Consider argument (S):

S1) Attack at dawn!

S2) *Therefore, someone commands something.

Here is the cognitivist's translation, (S'):

S1') I command that you attack at dawn.

S2') Therefore, someone commands something.

(S') is valid, so if cognitivism is true, then (S) is valid. But (S) seems invalid to me. If you think that (S) is valid, I invite you to compare it to argument (T):

T1) The weather is fine.

T2) *Therefore, someone asserts something.

(T) is invalid according to the traditional test for validity (counter-example: the weather is fine but no-one is around to comment on it). But (T) and (S) seem as though they ought to be parallel cases. Therefore, (S) is invalid too, and so cognitivism must be false.

Third, the cognitivist has some trouble with inconsistent commands. People can command inconsistently, and when they do, their commands are inconsistent with one another. However, it is not inconsistent to report that someone (even oneself) has commanded inconsistently. Therefore commands are not equivalent to reports of commands. Let's use an example to make the problem clearer – (6) and (7) are an inconsistent set of commands:

6) Attack!

7) Do not attack!

The cognitivist holds that these commands are equivalent to the reports (6') and (7') below:

6') I command that you attack.

7') I command that you do not attack.

It is not inconsistent, however, to report that an inconsistent set of commands has been made. It is possible to command inconsistently, out of forgetfulness, or perversity, or out of the desire to make a philosophical point: hey you, reader of this paper, attack and do not attack! I may now fairly and without (further) contradiction report to you what I have just done: I commanded that you attack and I commanded that you not attack. So reporting is all the the set of (6') and (7') does; so (6') and (7') are not an inconsistent set. If cognitivism is true, then (6) is equivalent to (6') and (7) to (7'), so if cognitivism is true, then (6) and (7) are not an inconsistent set. But (6) and (7) are an inconsistent set, so cognitivism is false.

4.6 Geach's argument

Geach's argument, (G), requires some further discussion. (G), I suggest, has the following form:

G1) If you are a faithful subject, rise up, Sir George! $p \rightarrow !q$

G2) But do not rise, fellow! $!\neg q$

G3) Therefore, you are not a faithful subject. $\neg p$

Earlier, I implied that (G) was an instance of modus tollens (and so indeed it appeared against the background of the standard view). The argument form I assign to (G) above, however, is *not* modus tollens: for it to be modus tollens, (G2) would have to be the negation of the consequent of (G1); that is, (G2) would have to have the form $\neg!q$, rather than $!\neg q$. Had I presented (G) as an instance of modus tollens, I would have courted the following objection: "(G2) expresses a simple command to the effect that George not rise; it does not express the negation of a simple command (whatever that may be) to the effect that George rise."

I am prepared to let this objector have their way. (G) is not an instance of modus tollens. It is, however, valid, as can be seen from the matrices below:

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}	✓		✓	
w_{p*}	✓		✓	
w_{*q}	✓	✓	✓	✓
w_{**}	✓	✓	✓	✓

(G1) $p \rightarrow !q$

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}		✓		✓
w_{p*}		✓		✓
w_{*q}		✓		✓
w_{**}		✓		✓

(G2) $!\neg q$

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}				
w_{p*}				
w_{*q}	✓	✓	✓	✓
w_{**}	✓	✓	✓	✓

(G3) $\neg p$

You would not be making a grave mistake, however, if you thought that (G) was an instance of modus tollens. For on my semantics, $\neg!q$ and $!\neg q$ are logically equivalent, as can be seen in the matrices below. The matrix for $\neg!q$ has each cell ticked iff that cell is blank for $!q$:

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}	✓		✓	
w_{p*}	✓		✓	
w_{*q}	✓		✓	
w_{**}	✓		✓	

$!q$

	w_{pq}	w_{p*}	w_{*q}	w_{**}
w_{pq}		✓		✓
w_{p*}		✓		✓
w_{*q}		✓		✓
w_{**}		✓		✓

$\neg!q$

The matrix for $!\neg q$ on the other hand has each row identical to the diagonal of the matrix for $\neg q$:

	W_{pq}	W_{p*}	W_{*q}	W_{**}
W_{pq}				
W_{p*}	✓	✓	✓	✓
W_{*q}				
W_{**}	✓	✓	✓	✓

¬q

	W_{pq}	W_{p*}	W_{*q}	W_{**}
W_{pq}		✓		✓
W_{p*}		✓		✓
W_{*q}		✓		✓
W_{**}		✓		✓

!¬q

As you can see, this produces the same matrix in either case. This is because ! is a "rearranging operator" – it just rearranges the cells of the input matrix, copying the cells on the diagonal into each row. ¬, on the other hand, inverts cells without moving them. It doesn't matter whether we copy first and then invert, or whether we invert first and then copy; either way we end up with the same matrix.

This means that the negation of a simple command that p is equivalent to a simple command that not p. I find this very congenial, for two reasons. First, because it is easy to mistake (G) for modus tollens (explanation: its premises and conclusion are semantically equivalent to the premises and conclusion of a modus tollens argument); second, because there appears to be no way in English to draw the scope distinction between ¬!q and !¬q (explanation: there is no significant distinction to draw).

4.7 Permissions

You might have thought that a theory that is (in part) about commands would also be about permissions. Though English has no special mood to do so, it's possible to permit someone to do something, by saying for example "I hereby permit that you attack at dawn!" You might also have thought that the negation of an imperative should be a permitting – an act of permission. My semantics, however, has no way to represent the content of an act of permission; and it's not the case that a negated command is an act of permission, because, as I pointed out in the section above on Geach's argument, the negation of the command that p is the command that not p. Is this a problem?

I think not. I would like to have a theory of acts of permission, but it is not easy to devise one – not for anyone! So I think it is fair to devise a theory of commands and leave acts of permission to be worked on later. Why is it so hard to devise a theory of acts of permission? Because it's not clear what they do. A command makes certain actions or states of affairs (which may previously have been permissible) impermissible, and which actions or states of affairs it makes impermissible depends on its content. One might think, then, that an act of permission makes actions or states of affairs (which may previously have been impermissible) permissible. If that were so, then acts of permission would be closely analogous to commands. But it is not so.

Suppose I say "Attack!" and then say "I hereby permit that you not attack!" If the simple-minded theory of acts of permission suggested above were true, then the second utterance would undo the effect of the first, leaving you able to comply with my commands by not attacking. But that's not what happens: by uttering both those sentences I command you inconsistently, so that there is nothing you can do to comply with my commands. So acts of permission do *not* make actions or states of affairs permissible!

Perhaps acts of permission are more like promises not to make a certain command than they are like commands. In Austin's terminology (Austin:1962, lecture 12), they are commissive, binding the speaker, not exercitive, advocating a course of action. But then it's no objection to my theory of commands that it doesn't include a theory of acts of permission.

References

Austin, J. L. (1962). How to do things with words. Oxford University Press.
Bratman, Michael (1999). Intention, plans, and practical reason. CSLI.
Chellas, Brian F. (1969). The Logical Form of Imperatives. Perry Lane Press.
Dummett, Michael (1981). Frege: Philosophy of Language. Duckworth.
Edgington, Dorothy (1995). On Conditionals. *Mind*, *104*, 235-329.
Geach, P. T. (1958). Imperative and deontic logic. *Analysis*, , 49-56.
Gibbard, Allan (1990). Wise Choices, Apt Feelings. Hardvard University Press.
Hare, R. M. (1952). The Language of Morals. Oxford University Press.
Hofstadter, Albert, and McKinsey, J. C. C. (1939). On the Logic of Imperatives. *Philosophy of Science*, *6*, 446-457.
Lewis, David K. (1979). Scorekeeping in a Language Game. *Journal of Philosophical Logic*, *8*, 339-359.
Pigden, Charles (1989). Logic and the Autonomy of Ethics. *Australasian Journal of Philosophy*, *67*, 127-151.
Ross, Alf (1944). Imperatives and Logic. *Philosophy of Science*, *11*, 30-46.
Searle, John (1969). Speech Acts. Cambridge University Press.
Smart, J. J. C. (1984). Ethics, Persuasion, and Truth. Routledge.
Stalnaker, Robert C. (1984). Assertion. , , 78-95.
Vranas, Peter C. (2008). New foundations for imperative logic. *Noûs*, *42*, 529-572.