

# Prospects for a Cognitive Norm Account of Logical Consequence

Thomas N. P. A. Brouwer

**Abstract:** When some P implies some Q, this should have some impact on what attitudes we take to P and Q. In other words: logical consequence has a normative import. I use this idea, recently explored by a number of scholars, as a stepping stone to a bolder view: that relations of logical consequence can be identified with norms on our propositional attitudes, or at least that our talk of logical consequence can be explained in terms of such norms. I investigate the prospects of such a *cognitive norm account* of logical consequence. I go over the challenges involved in finding a plausible bridge principle connecting logical consequence to cognitive norms, in particular a biconditional principle that gives us not only necessary but sufficient conditions for logical consequence in terms of norms on propositional attitudes. Then, on the assumption that an adequate norm can be found, I consider what the philosophical merits of such a cognitive norm account would be, and what theoretical commitments it would generate.

**Keywords:** logical consequence, normativity, propositional attitudes, expressivism

## Introduction

It is natural to think that when some P implies some Q, that fact should have some impact on our attitudes to P and Q. At a minimum: if I accept P but reject Q, my attitudes are not as they should be, in an objective sense of 'should'. This natural thought has motivated a small but significant literature in the philosophy of logic, exploring the idea that there is a characteristic sort of norm that relations of logical consequence give rise to. In this paper I want to build on that literature to explore the prospects of a bolder hypothesis: that relations of logical consequence not only give rise to cognitive norms, but that they *are* cognitive norms – or, at least, that our talk of logical consequence is to be explained in terms of such norms. That view has some challenges to overcome, I'll argue, but there is some mileage in it.<sup>1</sup>

## I. Looking for a bridge principle

The aforementioned literature on the normative import of logical consequence has focused on finding a suitable 'bridge principle', one which takes us from claims about logical relations between propositions to claims about how our propositional attitudes ought to be. Significant contributions to this research programme are MacFarlane (unpublished), Restall (2005) and Field (2009, forthcoming) and, from a more sceptical angle, Harman (1986). Here is a first stab at a bridge principle:

(1)  $P_1, P_2, \dots, P_n$  logically imply Q  $\rightarrow$  if one accepts  $P_1, P_2, \dots, P_n$ , one ought to accept Q.

Basically, the logical implication gives rise to an instruction to expand our acceptances in a certain way, given which things we've already accepted.<sup>2</sup> Before we discuss the merits of this principle, let me get a few clarifications out of the way.

---

<sup>1</sup>This paper was prepared within the 2013–15 AHRC project *The Metaphysical basis of Logic: the Law of Non-Contradiction as Basic Knowledge* (grant ref. AH/K001698/1) I would like to thank Ole Hjortland for first giving me the idea for this paper.

<sup>2</sup> Throughout this paper I will use the phrase 'B is a consequence of A' interchangeably with 'A implies B'.

First off, I use “ $P_1$ ”, “ $P_2$ ”, “ $P_n$ ” and “ $Q$ ” as names of propositions. The relation of logical consequence can also be understood as a relation between sentences, so it might be worth pointing out why I use the letters in this way. One obvious reason is that  $P_1$ ,  $P_2$ ,  $P_n$  and  $Q$  are to be the objects of propositional attitudes and, as the term suggests, those are usually taken to be propositions. Another reason is that if logical entailments are ultimately cognitive norms, then presumably the reason to care about logical consequence is to find out what is required of us, cognitively. If so, the requirements must be pre-existent, not artefacts of our logical theorising. In other words, the relation of logical consequence we are looking at is not that of some particular formal system – it is rather the thing that such systems aim to capture. As such, it seems more appropriate to treat it as a relation between propositions, a relation which can then be modelled, in formal systems, by a relation between sentences.

Secondly, I am not going to be greatly concerned with the nature of the (bi)conditionals appearing in these bridge principles. I am going to suppose that they at least (i) detach and (ii) contrapose. Apart from that I take no view on the matter.

Now, what are the merits of principle (1)? It seems to get something right, but not all that it should. To begin with, it might be thought to be too strong. One reason to think so is that it imposes obligations to believe even when the implication relation between  $P_1$ ,  $P_2$ , ...,  $P_n$  and  $Q$  is very difficult to find out about. Presumably some things we already accept either imply Goldbach’s Conjecture or imply its negation. Principle (1) tells us we are presently under an obligation to accept one or the other, though we can’t tell which. That may appear overly strong.

It is not clear whether it is indeed too strong. One thing to observe is that the ‘ought’ in principle (1) is intended to be an objective one. Failure to comply with it may well be excusable and even to be expected in cases where our epistemic situation is less than ideal. Also, there is a certain amount of interpretation possible with regard to the force of the ‘ought’. One might understand it in a very weak sense, as saying that we are ‘committed’ to accepting  $Q$  by what we already accept, where commitment could be such a deflated notion that we may well expect to sometimes end up committed to accepting things which we have no real business accepting. However, while it may be understood that way, it is implicit in the literature on the cognitive import of logical consequence that something more robust and genuinely demanding is intended. And that strikes me as right – commitment in this deflated sense doesn’t seem all that worth caring about. But even on a robust reading of the ‘ought’, it is not clear to me that it is too strong, given that some failures to comply might be eminently excusable, such as in the aforementioned case of Goldbach’s Conjecture.

Principle (1) has another clear failing, though. As Broome (1999) has pointed out, since logical consequence is very often thought to be reflexive, we get the result that I ought to accept anything I do in fact accept. And that seems inappropriate in cases where I have no business accepting something in the first place. This also shows us a way in which (1) is too limited: it only gives us instructions about how to expand our acceptances. And finding out that some  $P$  implies some  $Q$  can have other upshots too: if I happen to think that  $Q$  is a ridiculous thing to accept, the entailment shows me I shouldn’t accept  $P$  either. Principle (1) does not capture that.

### 1.1. Improving the principle

One may fix this latter failing by giving the ‘ought’ wide scope:

- (2)  $P_1, P_2, \dots, P_n$  logically imply  $Q \rightarrow$  one ought to be such that one if one accepts  $P_1, P_2, \dots, P_n$ , one accepts  $Q$ .

This gives us the right sort of bi-directionality. If  $P_1$  through  $P_n$  are good things to accept, then the right thing to do is to accept  $Q$  as well. But if  $Q$  is a bad thing to accept, not all of  $P_1$  through  $P_n$  should be

accepted. However, insofar as (1) was overly demanding, (2) is no better. Macfarlane considers restricting the antecedent to known entailment:

- (3) One knows that  $P_1, P_2, \dots, P_n$  logically imply  $Q \rightarrow$  one ought to be such that one if one accepts  $P_1, P_2, \dots, P_n$ , one accepts  $Q$ .

But he rejects this option, correctly in my opinion, because it muddles up the order of explanation between knowledge of entailments and the normative import of entailments. According to (3), logical entailments require things of me insofar as I know about them. But it seems, rather, than I have an interest in knowing about logical entailments precisely because they require things of me. In that case, these requirements had better be there already, before I find out about them. Furthermore, (3) might just be too weak altogether. In 'easy' cases of entailment, I do not get excused from complying with the norm simply by failing to spot the entailment. It's possible, though, to explain such cases by appealing to a further, separate obligation to know about certain entailments.

## 1.2. A proscriptive norm

Another different to deal with the alleged over-demandingness of (2) is to change the content of the obligations themselves:

- (4)  $P_1, P_2, \dots, P_n$  logically imply  $Q \rightarrow$  one ought not accept  $P_1, P_2, \dots, P_n$  whilst rejecting  $Q$ .

This principle, suggested by MacFarlane and defended in Restall (2005), changes the obligations from obligations to accept to obligations not to reject, and on this principle, logical entailments merely prohibit certain combinations of attitudes.<sup>3</sup> Whereas principles (1) – (3) give us prescriptive norms, (4) gives us proscriptive norms. If one worried about cases like Goldbach's Conjecture, these are now less worrisome: the fact that Goldbach's Conjecture is entailed by what I accept (let's say) does not require any action on my part. It merely requires that I do not reject Goldbach's Conjecture, which was already the sensible thing to do in my epistemic situation.

Principle (4) may strike some as overly weak, in that the requirements it gives rise to can be respected by simply suspending judgment on all matters. But I believe that the merely proscriptive nature of those norms is entirely appropriate to the role that logic plays in our epistemic lives. While we may well think that we are sometimes obliged to accept something, this can be understood as arising from factors extrinsic to logic. One might think, for instance, that there is a value to having true beliefs (and false disbeliefs) and so we have compelling reason to expand our beliefs and disbeliefs. Logic then merely gives us guidance on how to go about it. Such a division of labour gives us much more room to manoeuvre, too: for while one might believe that true belief is generally valuable, one might instead believe that its value depends on the subject matter (in certain matters it is not worth my while to expand my opinions) or on further extrinsic factors (e.g. true belief is only of full value if combined with justification, or somesuch). Principle (4) leaves such matters open, and that seems appropriate. It is the principle I favour, on balance.

This is not nearly the end of the matter. One might worry that even the relatively weak (4) still gives the wrong result with regard to cases like the Preface Paradox. In such cases it seems appropriate that I accept a bunch of propositions, yet reject their conjunction, even though the latter is entailed by the former. There are various ways to go; one might claim that the appropriate attitude to the conjunction is suspension of belief rather than rejection, which is compatible with (4). Or one might just reject adjunction.

---

<sup>3</sup> It should be noted that Restall defends a version of this principle which allows for multiple conclusions – i.e. some things entail some others when one ought not accept all of the former while rejecting all of the latter. I believe there is something to be said for a multiple-conclusion formulation, but in the present context it would be a distraction, so I restrict myself to single-conclusion principles.

MacFarlane (ibid.) favours diagnosing the paradox as a case where objective epistemic obligations come into conflict with subjective epistemic obligations; principle (4) gives us the objective one, the conjunction-rejecting attitude to the paradox is the one we are subjectively obliged to take. But probably the most popular line on the preface paradox is to treat it in terms of partial belief: the appropriate stance is one in which one has a high but lower-than-one credence in the conjuncts, and a low but higher-than-zero credence in the conjunction. Given enough conjuncts, that seems to give the right result. Principle (4) is in principle open to being generalised to the case of partial belief, and Field (2009 & forthcoming) attempts to do this. I will not explore this option, however, as it would introduce many complications not germane to our topic.

It should also be noted that there are further options for bridge principles that might in principle be considered. One might choose not to cast the principles in terms of obligations at all; there are other normative notions around. One might cast them in terms of reasons to accept or not reject, an option MacFarlane considers, or in terms of the epistemic value of accepting or rejecting something, or in terms of what is correct to accept and reject. I think, all things considered, that casting them in terms of obligations is our best bet, but other options deserve examination. However, for the remainder of this paper, I will be building on principle (4).

## 2. Getting logical consequence from cognitive norms

Let us assume, for the sake of argument, that principle (4) gives us a plausible account of the normative import of logical consequence. There is a considerable step from there to the much bolder claim I announced above, that logical entailments are to be identified with norms on our attitudes. Can bridge principle (4), or any of the others for that matter, bear that weight? If it is to do so, it had better be the case not only that every case of logical entailment has the normative upshot (4) predicts, but also that there are no cases in which such norms arise without it being a case of logical entailment. In other words, principle (4) should plausibly give us not only a necessary condition for logical entailment but a sufficient condition as well. We ought to be able to replace the conditional in (4) with a biconditional:

(5)  $P_1, P_2, \dots, P_n$  logically imply  $Q \leftrightarrow$  one ought not accept  $P_1, P_2, \dots, P_n$  whilst rejecting  $Q$ .

Is (5) plausible? Not as it stands; it gives us false positives. Plausibly, one ought not accept that the sky is green. Also plausibly, one ought not reject that snow is white. Combining the two, one ought not accept that the sky is green whilst rejecting that snow is white. But there is not the merest whiff of logical consequence to be detected here. We have, then, a false positive.

### 2.1. Necessitation

A fix immediately suggests itself: we should necessitate the consequent:

(6)  $P_1, P_2, \dots, P_n$  logically imply  $Q \leftrightarrow$  necessarily, one ought not accept  $P_1, P_2, \dots, P_n$  whilst rejecting  $Q$ .

Now the false positive is gone. But there are other ones much like it which meet this modalised requirement. Plausibly one ought not accept that the moon is red and green all over, and since it is necessarily not red and green all over, one necessarily ought not accept it – bearing in mind that we are speaking of objective epistemic obligations. It follows that necessarily, one ought not accept that the moon is red and green all over whilst rejecting that  $Q$ , for any  $Q$ . But surely this necessary falsehood does not logically entail everything. A further fix seems to be required.

### 2.2. Incoherence

Restall (2005) suggests that the cognitive requirements that logical entailments impose on us arise from the need to avoid a kind of incoherence in our mental states. Here's how one might build that in:

(7)  $P_1, P_2, \dots, P_n$  logically imply  $Q \leftrightarrow$  one ought not accept  $P_1, P_2, \dots, P_n$  whilst rejecting  $Q$ , on pain of incoherence.

The thought is as follows: when some  $P$  entails some  $Q$ , it's because either accepting  $P$  would land one in a state of incoherence, or rejecting  $Q$  would do so, or the combination of the two would do so. If  $P$  were a contradiction, one could see how it would do so, and similarly if  $Q$  were a tautology. Hence contradictions imply everything and tautologies follow from anything. But the claim that the moon is red and green all over is not a contradiction, necessarily false though it might be, and so it is not incoherent to accept it, even though one ought not do so. Thus, to accept it whilst rejecting some other proposition is also not incoherent (unless the other proposition is a tautology).

Principle (7) is promising, but as it stands it is not quite satisfactory. It appeals to the notion of incoherence, and that in itself is not an unreasonable thing to do – we do have such a notion. But the problem is whether the notion of incoherence that is available to us pre-theoretically is circumscribed enough to get the job done here. And if it is not, can we specify the more precise notion of incoherence in a way that does not involve a prior account of what logically entails what?

First off, it is not clear to me that the claim that the moon is red and green all over is not incoherent, on a pre-theoretic notion of incoherence. Its truth is unimaginable, so the attempt to do so at the very least engenders some confusion in us. It is not clear to us what believing that the moon is red and green all over would amount to. So we cannot assume that incoherence, on a pre-theoretic understanding, is sufficient for being logically at fault. Secondly, it seems fair to me to describe a subject as incoherent whose beliefs are radically unconnected, i.e. seemingly chosen at random. But to believe entirely unconnected things is not sufficient for being logically at fault – as long as none of the things one believes are inconsistent.

What we need, then, is a more precise and independently motivated notion of incoherence. Restall does sketch one. It boils down to the idea that incoherence is, in the first instance, a state where one accepts and rejects the same proposition – as state in which one, as it were, disagrees with oneself. Other states can then also be regarded as incoherent if they implicitly commit one to taking opposite attitudes to some particular proposition. For instance, the state of accepting  $A \& B$  and rejecting  $A$  is incoherent because, in some sense, to accept  $A$  is to commit to accepting  $A$ , and hence to be committed to accepting something one rejects.

The state of accepting and rejecting the same proposition strikes me as a distinctive form of incoherence, and to that extent I think Restall's suggestion is attractive. What is not obvious is how the extended notion of 'implicit' incoherence is to be unpacked satisfactorily, i.e. without relying on the fact that  $A \& B$  logically implies  $A$ . One way to do this, perhaps, is to claim that the very meaning of a logical connective like "&" is bound up with commitments or obligations to take (or not take) certain attitudes to its connectands, so that (for instance) by asserting a conjunction one expresses a felt obligation to not reject its conjuncts. I believe such a view has its attractions.<sup>4</sup>

### 2.3. Logical form

Another tactic for disposing of the false positive is to appeal explicitly to considerations of logical form. Here is how the principle might run:

---

<sup>4</sup> I have explored the merits of such a view my PhD thesis (2013).

(8)  $P_1, P_2, \dots, P_n$  logically imply  $Q \leftrightarrow$  for any  $\Phi_1, \Phi_2, \dots, \Phi_n$  and  $\Psi$  of the same logical form as  $P_1, P_2, \dots, P_n$  and  $Q$ , one ought not accept  $\Phi_1, \Phi_2, \dots, \Phi_n$  whilst rejecting  $\Psi$ .

The thought would be as follows. Perhaps we necessarily ought not accept that the moon is red and green all over, but there is nothing about the logical form of that claim that makes it so. It is not the case, for instance, that we necessarily ought not accept that the moon is rocky and cold all over. This is different with contradictions: the very reason they are unacceptable (*pace* dialethic views on contradictions) is that they have a particular logical form, and *mutatis mutandis* the same goes for logical tautologies.

So far I have been speaking solely about propositions, but the prototypical bearers of logical form are sentences. There is thus a theoretical choice at this point; either we understand principle (8) as somehow bringing sentences into the picture, or we adopt some account of propositions one which they can be attributed logical form. The latter is not at all unprecedented: views that treat propositions as structured entities can talk of the logical form of propositions, and such views can be independently motivated, for instance by the need to deal with the phenomenon of hyperintensionality (see for instance Soames (1987)).

If we do not want to treat propositions as bearers of logical form, we could understand principle (8) with reference to the logical form of sentences. Let us imagine a hypothetical language  $L$ , the consequence relation of which (a relation between sentences) perfectly models the 'true' consequence relation (a relation between propositions). When principle (8) speaks of propositions 'of the same logical form as' some proposition  $P$ , take these to be the propositions expressed by sentences of the same form as any sentence of  $L$  that expresses  $P$ . I say 'any', because if propositions are understood to lack logical form and hence as coarse-grained, there will be multiple sentences of  $L$  that express  $P$ .

But whichever way talk of logical form in principle (8) is understood, it had better be the case that the notion of logical form employed does not implicitly involve a prior conception of logical consequence. Otherwise, the account of logical consequence which we are exploring will end up being parasitic upon another. One option that fits with the spirit of the proposal we are examining is to understand the logical form of a sentence/proposition directly in terms of what the assertion/acceptance of it would require of us as regards attitudes to other propositions (particularly the ones occurring in the sentence/proposition). In other words, something similar to the view floated above, in the context of Restall's notion of logical incoherence.

#### 2.4. Generic consequence

An entirely different response to the false positives we've been discussing is to embrace them. Perhaps they are cases of consequence – just not cases of logical consequence. This would amount to shifting the goalposts: to claim that the more interesting notion or our account to capture is a more general notion of consequence, of which logical consequence is just one species.<sup>5</sup> This is not a crazy view. One might think that this more general kind of consequence is really of more basic interest to us than logical consequence: after all, what we primarily want is for our reasoning to take us from things worth believing to (only) other things worth believing. The logical validity of such reasoning, one might think, is only of interest to us only secondarily: insofar as it confers particularly high grade of reliability on the reasoning.

There are a few ways to go. One might hold that the proper notion of generic consequence is captured by principle (6). That is: some propositions generically imply some other proposition when, of necessity, one ought not accept the former whilst rejecting the latter. However, once one has abandoned logicity as a target, it might be hard to motivate caring in particular about those inferences which are backed up by necessitated norms. One might as well embrace all of the purported false positives and go back

---

<sup>5</sup> I would like to thank Göran Sundholm and Dorothy Edgington for drawing my attention to this option.

to principle (5). After all, to reason from 'Fido is a dog' to 'Fido is a mammal' is a perfectly good thing to do, even though neither logical form nor any deep kind of necessity backs up the goodness of the inference.

## 2.5. Relevance

Nevertheless, even if one is willing to abandon logical consequence as one's particular target, one might be bothered by another aspect of the false positives discussed above: the fact that they constitute failures of relevance. To reason from the 'grass is red' to 'snow is white' is indeed bad – one might think – but not so much because the latter does not follow logically from the former, but rather because the latter has nothing to do, content-wise, with the former.

Relevance is usually discussed by logicians as a possible necessary condition on logical consequence. But there is no reason why one could not care about relevance in its own right, quite apart from whether it contributes to the *logicality* of a consequence relation. One might simply desire from one's inferences that they preserve truth or some other privileged status, and in addition that they be relevant. One might think that inferences meeting the standard of relevance are by that token more reliable, in a way different from but comparable with the way that logical validity makes inferences more reliable.

There are various ways one might build a condition of relevance into principle (5), and inspiration for these might be sourced from the various approaches to relevant logic that are already around in the literature. A popular view is that relevant consequences should at least exhibit *variable sharing*: the premises and the conclusion have some proper or improper part in common. This however, can only be built into the principle if one has a notion of logical form available, which one may not. A better option might be to go directly for a condition of a more semantic nature. At a first approximation, we might formulate it in terms of aboutness:

- (9)  $P_1, P_2, \dots, P_n$  imply  $Q \leftrightarrow$  one ought not accept  $P_1, P_2, \dots, P_n$  whilst rejecting  $Q$  and every proposition among  $P_1, P_2, \dots, P_n$  and  $Q$  is about the same thing or things as at least one other proposition among  $P_1, P_2, \dots, P_n$  and  $Q$ .

The notion of some proposition being about some thing or things would then of course need theoretical unpacking, but this is already an area of active research (see for instance Yablo 2014). It should be noted too that if one is not attracted to the goal-post-shifting tactic, and one wants to pursue a notion of properly logical consequence, there is no reason why one could not use a relevance condition in addition to any of the other conditions above, albeit at the price of further complicating one's bridge principle.

## 3. Features of a cognitive norm account

The preceding section leaves a lot open, and is meant to. I have indicated a number of directions along which one might pursue a biconditional principle that captures necessary and sufficient conditions for logical or generic consequence. I will be happy if I have shown that there is at least some hope for finding such a principle. In the remainder of this article I will assume, for the sake of argument, that we have found ourselves a biconditional principle that does the work required. Let us see what a theory of logical consequence built on top of such a principle would be like.

### 3.1. Identity theory versus expressivism

The most straightforward version of a cognitive norm account of logical consequence will simply identify the relation of logical consequence between some  $P$  and  $Q$  with the norm that our preferred biconditional associates with it. Relations of logical consequence just *are* norms, on this version of the view. Call it the 'identity theory'. Is it an appealing view?

On the identity theory, to assert the existence of a relation of consequence is to assert the existence of a norm. There are advantages to that. Assuming that we were in any case committed to some norm arising from the consequence relation, we have now effected a theoretical reduction of two entities to one. And we have rid ourselves of the burden of explaining why relations of consequence have a bearing on our propositional attitudes – it is their very nature to have a bearing on propositional attitudes, given the sorts of things they are.

But one might also have metaphysical misgivings at this point. Normativity is and has long been a source of philosophical puzzlement. How can such things exist in a world that at bottom is supposed to consist of cold hard facts? How can an 'is' ever yield an 'ought'? If the identity theory identifies an entity which up until now was not a source of metaphysical puzzlement (though perhaps of other kinds of puzzlement) with an entity that is, then we might be worse off than before.

Norms are indeed ill-understood, but that is no reason for despair, for many important notions are. Meta-ethicists are on the case, and hopefully someday they will tell us how it is that norms fit in a world of cold hard facts. But, that being said, it might be that the task of explaining how cognitive norms exist can be avoided altogether. Hartry Field, who as far as I am aware is the only one presently defending a cognitive norm account of logical consequence (2009 & forthcoming), does not defend the identity theory. In fact he is a primitivist about logical consequence, and does not think it admits of any analysis. But he thinks our talk of logical consequence can be explained in an expressivist manner, in terms of cognitive norms.

An expressivist version of the cognitive norm account would say that, whenever we assert that some P implies some Q, we are expressing a felt obligation to have certain attitudes to P and Q. Perhaps, along the lines of principle (5), we feel that we ought not accept P whilst rejecting Q – and nor should anyone else. The necessary and sufficient conditions that principles (5) – (9) give us are, on this view, not conditions for the obtaining of a consequence relation, but conditions for the assertibility of a consequence claim. Talk of norms does not vanish on the expressivist version, but appears in a different, less metaphysically robust way. Statements of logical consequence do not state the existence of a norm, but express our commitment to a norm. A norm one so commits to need not have any mind-independent status – it can be a mere fiction, and nevertheless do all the work it needs to.

That, at least, is the hope. It is not altogether obvious that metaphysical misgivings about normativity will not reappear somewhere else in the picture. One might be tempted to say, for instance, that the cognitive norms that we are inclined to impose on ourselves are not entirely accidental. There are better and worse norms that one might impose on oneself, and that fact demands explanation of a metaphysical sort. And perhaps there our prior puzzlement will reassert itself. Perhaps – but it may be that a sufficiently serious expressivist will dismiss such further questions about better and worse norms as symptoms of an insufficiently exorcised realism. One can indeed – they would say – affirm that some norm is better than another; but to do so is nothing more than to commit to the former and not the latter.

The expressivist version of the account also helps with another potential worry about the identity theory. When we assert that some P implies some Q, we do not think that we are talking about obligations or attitudes. We think we are talking about whatever P and Q are about, and nothing else. Yet if the identity theory is correct, to talk about a relation of logical consequence is *ipso facto* to talk about obligations and attitudes. Isn't that the wrong result? To be honest, I think the force of this worry should not be overstated. It is perhaps a surprising result that talk of consequence is talk of cognitive norms, but the identity theory is an analysis of logical consequence, and analyses are allowed to sometimes be surprising. We did not think we were talking about obligations and attitudes, but perhaps it has turned out we were.

Nevertheless, if one does feel that the identity theory does too much violence to our intuitions of aboutness, the expressivist theory does better. On that view, talk of consequence isn't about anything more than the things that are stated to follow from each other. Something about obligations and attitudes is expressed, but not by means of a descriptive claim about obligations and attitudes. It is like saying 'hooray': doing so does not describe my elation – it expresses it.

Although it is not clear that the identity theory cannot overcome the above challenges on its own, I do think that the expressivist version of the cognitive norm account of consequence is preferable overall. It also strikes me – though I could not say precisely why – as more elegant.

### **3.2. Whither truth?**

One striking feature of a cognitive norm account of (logical) consequence is that it does not appear to mention truth at all, despite the fact that this is a notion usually closely associated with consequence. Is this a feature or a bug? It might appear as a bug if one thinks that an adequate theory of consequence should have a close connection between truth and consequence drop out automatically. But it is not clear to me that the connection ought to be quite so automatic.

Let's first observe that we may still secure a general connection between truth and consequence by the expedient of equating the obligatoriness of accepting a proposition with its truth, and the obligatoriness of rejecting it with its falsity. True, we cannot allow truth and falsity to be the notions that *explain* why some things ought to be accepted and others ought to be rejected, for if we did that our cognitive norm account would effectively collapse into a truth-theoretic account of consequence – and such accounts are already widely available. So perhaps this will not satisfy the friend of truth, because they might demand an explanatory role for truth vis-à-vis consequence.

But perhaps they should not be so quick to do so. First off, many philosophers are attracted to deflationary accounts of truth, on which the entire content of the truth predicate is given by a principle of disquotation, and the notion of truth has no robust metaphysical nature of any sort. But if one is attracted to that sort of account, one should be hesitant about according truth any explanatory role vis-à-vis consequence or indeed anything else: can something with no nature of its own explain the nature of anything else? The cognitive norm account allows truth to drop out of the picture, as perhaps it should.

Secondly – and somewhat speculatively – there might be discourses on which consequence has a grip but truth does not. Suppose we think that moral discourse is not truth-apt. We might think that, and still feel that some moral claims should be accepted and others rejected, perhaps because they are appropriate in some non-alethic fashion. There is no reason to think that considerations such as consistency, coherence and relevance could not be just as applicable to a non-truth-apt discourse as to a truth-apt one, and if so, relations of consequence – conceived as cognitive norms – would obtain in these discourses. It would be good, then, if our account of consequence was not inextricably bound up with truth.

### **3.3. Non-classical logics and structural rules**

Would a cognitive norm account of (logical) consequence be ecumenical with regard to the wide variety of formal systems of logic that are out there, or would it be compatible with only some? This depends on the conditions that our preferred biconditional principle imposes on the consequence relation. As it happens the principles considered above apply without any issues to a wide variety of logics, including classical, paracomplete and at least the more standard paraconsistent logics. This is not surprising: the main contributors to the literature on the normative import of logical consequence, on whose work we have built here, are keenly aware of non-classical logics, and have designed their principles appropriately.

The matter becomes trickier once we consider sub-structural logics, logics which omit one or more of the classic structural properties of the consequence relation: transitivity, reflexivity, monotonicity, permutability and contractibility. Restall (ibid.) has suggested that principle (4) does impose these properties on the consequence relation, but it seems that it does not really do so without the help of some further assumptions. Since we have mostly explored norms based on principle (4), it will be instructive to look at how much it tells us about the consequence relation. As a reminder, here is the principle:

(4)  $P_1, P_2, \dots, P_n$  logically imply  $Q \rightarrow$  one ought not accept  $P_1, P_2, \dots, P_n$  whilst rejecting  $Q$ .

First we can observe that the principle does not immediately give us transitivity. Suppose that I ought not accept A whilst rejecting B, and that I ought not accept B whilst rejecting C. With regard to the first requirement, I can oblige by accepting A and suspending judgment on B. With regard to the second, I can oblige by suspending judgment on B and rejecting C. I thus may end up accepting A and rejecting C in full compliance with (4); thus it does not tell me that I ought not accept A whilst rejecting C. Nevertheless, transitivity might be secured by means of a further, quite plausible norm: that there are no propositions such that I ought not accept them and ought not reject them. This norm is to be firmly distinguished from the less plausible norm that there are no propositions such that I neither ought to accept them nor ought to reject them. The latter rules out, draconically, that there could be propositions concerning which I have no obligations; the former, more reasonably, rules out that there are propositions concerning which I am not allowed any attitudes. With this latter norm in place, the counterexample to transitivity above will be ruled out. A defender of non-transitive logic wanting to adopt the cognitive norm account therefore had better reject this additional norm, which they might motivate by claiming that paradoxical propositions, or whatever else they regard as proper loci for failures of transitivity, are indeed such that we are obliged not to have any attitudes towards them.

On the very plausible assumption that I ought not accept and reject P, for any P, principle (4) tells me that any P implies itself. In other words, reflexivity is fairly easily obtained, and quite hard to get rid of.

With regard to monotonicity (aka weakening), permutability (aka exchange) and contractibility (aka contraction) we might think that they are quite easily obtained. With regard to monotonicity: if it is the case that I ought not accept some P whilst rejecting some Q, then presumably I also ought not do that whilst doing any other particular thing, like accepting R or playing cricket. With regard to permutability: since there is no apparent sense in which my obligations with regard to some propositions in any way involve any ordering on them, there is also no sense in which a change in their order could affect my obligations. With regard to contractibility: there seems to be no difference between me merely accepting P and me accepting P and accepting P, and so it seems there could be no difference between the obligation to accept P and the obligation to do so twice.

We are assuming too much here, though. Implicitly, the reasoning I used to obtain monotonicity assumed that the obligations I have with regard to P and Q arise only from the nature of P and Q, i.e. that they are somehow internal to P and Q. If so, then it is indeed clear that my performing some further action cannot change that. But if my obligation to not accept P whilst rejecting Q arose specifically because of something else I was doing or (let's say rejecting R) then it is conceivable that my obligations with regard to P and Q might change if I stopped doing this further thing (rejecting R). So then it does not follow from my obligation to not accept P whilst rejecting Q that I have an obligation to not accept P and accept R whilst rejecting Q. Accepting R might be just the thing for getting me out of my obligation vis-à-vis P and Q. The plausibility of monotonicity thus depends on whether we think that our cognitive obligations concerning any given bunch of propositions do indeed depend only on the nature of those propositions. That claim strikes me as plausible, but it might be doubted, and it should be doubted by a proponent of non-monotonic logic.

Similarly, with regard to permutability, though we do not think that propositional attitudes, of themselves, have any ordering, one could imagine such orderings obtaining relative to some external factor (for example time). If my obligations involving some propositional attitudes depend in part on that external factor (i.e. the temporal order in which I adopted them) then a change in ordering could mean a change in obligations. So to uphold permutability is implicitly to deny the involvement of any such external factor. Similarly with contractibility: on the face of it there is no difference between me accepting P and doing so twice, but if there is some external factor relative to which my attitudes can be distinguished in a more fine-grained way, there may be such a difference after all. To uphold contractibility is thus to deny the relevance of any such external factor.

To sum up: it seems that all the structural rules can be called into question within the context of a cognitive norm account of (logical) consequence, though some more easily than others. A cognitive norm account need not restrict us to some particular subset of the available logics.<sup>6</sup>

### 3.4. Logical pluralism

At the beginning of this article I said that a cognitive norm account of (logical) consequence involves the thought that there is some particular relation of consequence that exists independently of our formal systems, a relation between propositions, which we try to model in those formal systems by means of a relation between sentences. One might assume this involves thinking that there is a single ‘true’ logic out there, i.e. logical monism. If so, then a cognitive norm account is of no use to those attracted to logical pluralism, the view that there is a plurality of logics, all of which are in some manner ‘in good standing’. The assumption is natural, but I don’t think a cognitive norm account is committed to monism in quite such a rigid fashion.

What a cognitive norm account is strictly committed to is the claim that there is one and only one consequence relation out there which corresponds to the norm which governs our propositional attitudes. It leaves open the possibility that there are other relations among propositions which, somehow or other, also deserve to be called consequence relations, and which are significant to us for some different reason. If we imagine for the moment that it is the classical consequence relation which captures our general cognitive obligations, it could for instance be that some other, stricter logic captures our obligations with regard to some particular intellectual pursuit or some particular subject matter, such as mathematics. A certain amount of pluralism is thus compatible with a cognitive norm account.

## 4. Conclusion

In this article I have sketched a possible view of logical consequence. It is a view, or family of views, on which consequence relations are to be identified with norms on our propositional attitudes, or, at least, our talk of consequence relations is to be explained in terms of such norms. I have explored some of the challenges involved in getting the basic machinery of such an account off the ground: to wit, a plausible biconditional bridge principle that connects consequence claims with normative claims. And I have explored, to some extent, what a cognitive norm account of (logical) consequence would be like and what further theoretical commitments it would generate.

I am by no means convinced that a cognitive norm account gets at the truth of things. No one should be; it is too early to tell. But I think that it may get something important right, and I think that there is enough mileage in the view to make it worth exploring in greater detail.

---

<sup>6</sup> Which is not to say that no version of it will. It is entirely possible that some ways of embellishing principle (5) will indeed restrict us to a narrower set of logics. Imposing a relevance condition, for instance, will compromise monotonicity.

Thomas N. P. A. Brouwer

Research Fellow

Northern Institute of Philosophy

University of Aberdeen

Old Brewery, High Street, Aberdeen AB24 3UB

tnpabrouwer@gmail.com

### **Bibliography**

- Broome, J. (1999). Normative Requirements. *Ratio*, 12, 398–419.
- Brouwer, T. N. P. A. (2013). *The Metaphysical Commitments of Logic*. PhD thesis, University of Leeds.
- Field, H.H. (2009). What is the Normative Role of Logic. *Proceedings of the Aristotelian Society*, Suppl. Vol. 83, 251–68.
- Field, H.H. What Is Logical Validity? Forthcoming in: Caret, C., Hjortland, O. (eds.), *Foundations of Logical Consequence*. Oxford: Oxford University Press.
- Harman, G. (1986). *Change in View*. Cambridge, Mass.: M.I.T. Press.
- MacFarlane, J. (unpublished). In What Sense (If Any) Is Logic Normative for Thought? Retrieved from <<http://johnmacfarlane.net>>.
- Restall, G. (2005). Multiple Conclusions. In: Hajek, P., Valdes-Villanueva, L., Westerstahl, D. (eds.), *Logic, Methodology and Philosophy of Science: Proceedings of the Twelfth International Congress (189-205)*. London: Kings' College Publications.
- Soames, S. (1987). Direct Reference, Propositional Attitudes and Semantic Content. *Philosophical Topics* 15, 47–87.
- Yablo, S. (2014). *Aboutness*. Princeton: Princeton University Press.